Welfare effects of ex-ante bias and tie-breaking rules on Observational Learning with Fake Agents

Pawan Poojary and Randall Berry
ECE Department, Northwestern University, Evanston, IL 60208, USA
Email: pawanpoojary2018@u.northwestern.edu, rberry@ece.northwestern.edu

***Abstract*—Networks that provide agents with access to a common database of the agents' actions enable an agent to easily learn by observing the actions of others, but are also susceptible to manipulation by "fake" agents. Prior work has studied a model for the impact of such fake agents on ordinary (rational) agents in a sequential Bayesian observational learning framework. That model assumes that ordinary agents do not have an ex-ante bias in their actions and that they follow their private information in case of an ex-post tie between actions. This paper builds on that work to study the effect of fake agents on the welfare obtained by ordinary agents under different ex-ante biases and different tie-breaking rules. We show that varying either of these can lead to cases where, unlike in the prior work, the addition of fake agents leads to a gain in welfare. This implies that in such cases, if fake agents are absent or are not adequately present, an altruistic platform could artificially introduce fake actions to effect improved learning.**

***Index Terms*—Information cascades, herding, Bayesian optimality, ex-ante bias.**

I. Introduction

In many scenarios, agents seek to learn from observations of other agents' actions. Such learning is facilitated by networks that enable an agent to connect to a common database for making decisions. Examples of such networks include wireless networks in which devices upload and share their data through a common Access Point, on-line markets, social networks, etc. Bayesian observation learning provides a framework for studying such scenarios. In these models, Bayesian rational agents take actions over time. Each agent updates it beliefs about about the value of a given action based on its observations of previous agents' actions. Initial models for such settings in [1], [2] considered the case where there is a common underlying value for each of two actions, which agents take sequentially after fully observing the prior agents' actions. Additionally, each agent receives an i.i.d. binary-valued signal modeling their ex-ante beliefs about the underlying "state-of-the world" that determines the values of the two actions. A key result is that with probability one, agents will enter into an *information cascade* in which all subsequent agents follow the majority decision of prior agents, regardless of their own signal.

This basic observational learning model has been extended in many directions including relaxing the assumption of i.i.d. binary valued signals [3], assuming that every agent does not observe all previous agents' actions [4], allowing for imperfect observations [5], [6]. Other variations include [7]–[13]. This paper, is motivated by one of these extensions in [6], which considered a model similar to [1], [2] except that with a given probability each agent is a "fake" agent who always reports a favoured action regardless of the other users' actions. This models scenarios in which fake agents report a certain action to encourage other non-fake agents to also adopt that action. In [6], which was extended from [14], it was shown that the presence of such fake agents can in some cases reduce the likelihood of their preferred cascade. Conditioned on the underlying state of the world, these fake agents may even lead to improvement in the *welfare*, i.e., the expected pay-off obtained by the agents. However, when averaged across the state of the world, numerical results in [6] show that the welfare is always reduced by the presence of fake agents. A motivation for this paper is to understand this welfare reduction and how it depends on both the ex-ante beliefs of the agents and the tie-breaking rule used when agents are indifferent between the two actions.

Fake agents degrade the information obtained in some observations. This is reminiscent of Blackwell's theorem on comparing different information structures for single agent decision problems [15], [16]. Blackwell's result shows that if one information structure is a *garbling* of a second, then the former will result in a lower welfare for any decision problem faced by a Bayesian rational agent. A natural question is then: Can the welfare loss observed in [6] be explained through a similar argument extended to this sequential decision making setting? Our results show that this is not the case. Indeed if Blackwell's results directly generalized to this setting, it would imply a welfare loss regardless of the agents' ex-ante beliefs. We show that this is not the case and that when agent's have an ex-ante bias, the presence of fake agents may lead to improved welfare. We also show that even when agents are ex-ante unbiased, if the tie breaking rule assumed in [6] is changed, fake agents can again lead to improved welfare. This is in line with other works that have shown that "better" information in multiplayer games may not always lead to improved pay-offs for players (see e.g. [17]).

The remainder of 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 characterize agents' welfares and identify im-

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

ISBN 978-3-903176-55-3 © 2023 IFIP

Authorized licensed use limited to: Northwestern University. Downloaded on February 13,2024 at 01:33:08 UTC from IEEE Xplore. Restrictions apply.

portant properties exhibited by them. In Section V, we present our Markov chain formulation and devise an iterative method to compute cascade probabilties. Further, in Subsections V-C and V-D, we observe the effects of varying the ex-ante bias and the tie-breaking rule respectively, on the agents' asymptotic welfare. We present our conclusions in Section VI and defer a detailed description of the iterative method to the Appendix.

II. Model

We consider a model similar to [1] in which there is a countable sequence of agents, indexed $i = 1, 2, \ldots$ 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). We assume a non-revealing general prior for the true value, $\mathbb{P}(V = G) = q \in (0, 1)$.

The agents are Bayes-rational utility maximizers where the pay-off received by each agent i, denoted by π_i, depends on its action A_i and the true value V as follows. If the agent chooses N, his payoff is 0. Whereas, if the agent chooses Y, he incurs a cost of $C > 0$ for buying the item and gains an amount that reflects the item's value/utility to its buyer. The buyer gains the amount x if $V = G$ and $-y$ if $V = B$, where $x > C$ and $y \geq 0$. Then, the net pay-off for any agent i is

$$\pi_i = \begin{cases} x - C, & \text{if } A_i = Y \text{ and } V = G, \\ -y - C, & \text{if } A_i = Y \text{ and } V = B, \\ 0, & \text{if } A_i = N. \end{cases} \tag{1}$$

It follows from (1) that the *ex ante* expected pay-off for any agent is $xq - y(1-q) - C$ if it buys the item, and is 0 otherwise. In the case $xq - y(1-q) > C$, an agent *a priori* prefers taking action Y over N, and vice-versa if $xq - y(1-q) < C$. The model in such cases is said to have an ex-ante *bias* between the two actions. On the other hand, if $xq - y(1-q) = C$, then the ex ante expected pay-off for any agent is 0 for either of the actions. Let $q^* \in (0, 1)$ be this unique probability of $V = G$, which is $q^* := (y + C)/(x + y)$.[1] Thus, if $q = q^*$, then to begin with, an agent is indifferent to the two actions. While previous works [1], [5], [6], [10], [12] consider this *unbiased* ex-ante preference of actions, our work extends to the generic case of a possible ex ante bias in the actions.

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 1a, partially reveals the information about the true value of the item through a binary symmetric channel (BSC) with crossover probability $1 - p$, where $1/2 < p < 1$. This 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 V. Each agent i takes a *rational* action A_i that depends on his private signal S_i and the past observations $\{O_1, O_2, \ldots, O_{i-1}\}$ of actions $\{A_1, A_2, \ldots, A_{i-1}\}$. 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. $1 - \epsilon$, where ϵ is common knowledge. 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, if $A_i = N$ then w.p. $1-\epsilon$, the reported action $O_i = N$ and w.p. ϵ, $O_i = Y$. Whereas, if $A_i = Y$ then $O_i = Y$ w.p. 1. Refer to Figure 1b.

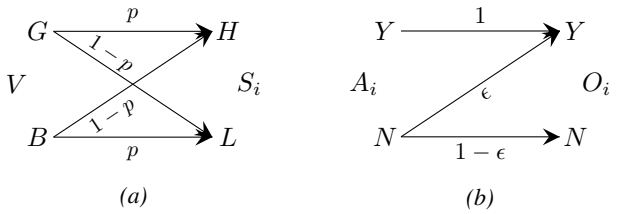

(a) (b)

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 $\mathcal{H}_{n-1} := \{O_1, O_2, \ldots, O_{n-1}\}$. Then, the Bayes' optimal action for every agent n, $A_n \in \{Y, N\}$ is chosen such that it maximizes the expected pay-off given its information set $I_n := \{S_n, \mathcal{H}_{n-1}\}$. Let $\gamma_n(S_n, \mathcal{H}_{n-1}) \triangleq \mathbb{P}(G|S_n, \mathcal{H}_{n-1})$ denote the posterior probability of the item being good, i.e., $V = G$. The expected pay-off E_{A_n} that agent n obtains by taking action A_n, given $\{S_n, \mathcal{H}_{n-1}\}$ is then expressed by:

$$E_{A_n} = \begin{cases} (x - C)\gamma_n - (y + C)(1 - \gamma_n) & \text{if } A_n = Y, \\ 0 = (x - C)q^* - (y + C)(1 - q^*) & \text{if } A_n = N. \end{cases} \tag{2}$$

Here, recall that q^* is the unique probability of $V = G$ for which an agent would be indifferent to the two actions. By comparing E_Y with E_N in (2), it follows that the Bayes' optimal decision rule is:

$$A_n = \begin{cases} Y, & \text{if } \gamma_n > q^*, \\ N, & \text{if } \gamma_n < q^*, \\ T, & \text{if } \gamma_n = q^*. \end{cases} \tag{3}$$

Note that when $\gamma_n = q^*$, $E_Y = E_N$ and so the agent is indifferent between the actions. This conditon is characterized by the two elements of set I_n, namely, the history $\mathcal{H}_{n-1}$ and the private signal S_n favouring opposing actions, equally strongly. In this case, we assume that all agents adhere to a common *tie-breaking* decision rule, denoted by T. In prior works [2], [5], [6], [10], [12], the decision rule in this case follows the private signal S_n, i.e., the agent buys the item only if $S_n = H$. Another choice in this case is to follow the history $\mathcal{H}_{n-1}$, which also means to oppose the private signal S_n, i.e., the agent buys the item only if $S_n = L$. In this paper, we focus on both these choices for breaking ties, which is represented by $T \in \{s \text{ (follow } S_n), h \text{ (follow } \mathcal{H}_{n-1})\}$.

[1] The condiitons $x > C > 0$ and $y \geq 0$ assumed in (1) ensure that $q^* \not\leq 0$ and $\not\geq 1$, as otherwise, all agents prefer a fixed action regardless of their beliefs on the item's true value, V.

ISBN 978-3-903176-55-3 © 2023 IFIP

Authorized licensed use limited to: Northwestern University. Downloaded on February 13,2024 at 01:33:08 UTC from IEEE Xplore. Restrictions apply.

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

We first consider the rule $T = s$. It follows from (3) that, agent n cascades to a Y (N) if and only if $\gamma_n > q^*$ $(< q^*)$ for all $S_n \in \{H, L\}$. The other case being $\gamma_n \geq q^*$ for $S_n = H$ and $\gamma_n \leq q^*$ for $S_n = L$; in which case, agent n follows S_n. To better understand the above cascade conditions, we encapsulate the information contained in the history $\mathcal{H}_{n-1}$ observed by agent n in the term $g_{n-1}(\mathcal{H}_{n-1}) \triangleq (\frac{1-q}{q}/\frac{1-q^*}{q^*})l_{n-1}(\mathcal{H}_{n-1})$, where $l_{n-1}(\cdot) \triangleq \mathbb{P}(\cdot|B)/\mathbb{P}(\cdot|G)$ is the likelihood ratio function of the public history $\mathcal{H}_{n-1}$. Further, we define $\beta_n(\cdot) \triangleq \mathbb{P}(\cdot|B)/\mathbb{P}(\cdot|G)$ as the likelihood ratio of the private signal S_n, where it follows from Figure 1a that $\beta_n(H) = (1-p)/p$ and $\beta_n(L) = p/(1-p)$. Next, using Bayes' rule, we express γ_n in terms of g_{n-1} and β_n as $\gamma_n = 1/\big(1 + \beta_n g_{n-1}\frac{1-q^*}{q^*}\big)$. The inequality $\gamma_n > q^*$ can then be simplified to the form $g_{n-1} < 1/\beta_n$. Thus, for $T = s$, the condition on γ_n for a $Y\,(N)$ cascade translates to $g_{n-1} < 1/\beta_n$ $(> 1/\beta_n)$ for all S_n. This gives bounds on g_{n-1} for a cascade to occur, as stated in the next lemma.

Lemma 1: Under the tie-breaking rule $T = s$ $(T = h)$, agent n cascades to a Y if $g_{n-1} < \frac{1-p}{p}$ $\big(g_{n-1} \leq \frac{1-p}{p}\big)$, cascades to a N if $g_{n-1} > \frac{p}{1-p}$ $\big(g_{n-1} \geq \frac{p}{1-p}\big)$, and otherwise follows its private signal S_n.

For the other rule, $T = h$, similar techniques yields the bounds on g_{n-1}, also stated in Lemma 1. If agent n cascades, then O_n does not provide any additional information about the true value V to the successors over what is contained in $\mathcal{H}_{n-1}$. As a result, $l_{n+i} = l_{n-1}$, which implies $g_{n+i} = g_{n-1}$ for all $i = 0, 1, 2, \ldots$ and hence they remain in the cascade, which leads us to the following property, also exhibited by prior models, e.g. [1], [2], [5]–[7].

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. For every such observation O_i, $i \leq n$, as S_i is conditionally independent of the history $\mathcal{H}_{i-1}$ given V, the likelihood ratio can be updated as: $l_i = (\frac{1-b}{a})l_{i-1}$ if $O_i = Y$, else $l_i = (\frac{b}{1-a})l_{i-1}$ if $O_i = N$. Here,

$$a := \mathbb{P}(O_i = Y|V = G) \text{ and } b := \mathbb{P}(O_i = N|V = B) \quad (4)$$

denote the probabilities that an observation O_i follows the true value V, given that A_i follows S_i. It can be shown from Figures 1a and 1b that in the above case,

$$a = p + (1-p)\epsilon \text{ and } b = p(1-\epsilon). \quad (5)$$

Now, as a result of the updates, l_n can be shown to depend 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, $l_n = \big(\frac{1-p}{p}\big)^{\eta n_Y - n_N}$ where the exponent is the difference between the number of Y's weighted by η and the number of N's. The weight η is given by

$$\eta := \log\Big(\frac{a}{1-b}\Big)/\log\Big(\frac{p}{1-p}\Big). \quad (6)$$

Then, given that agent n has not cascaded, the term g_n can be expressed as: $g_n = l_n\big(\frac{1-p}{p}\big)^{h_0} = \big(\frac{1-p}{p}\big)^{h_n}$, where

$$h_n = \eta n_Y - n_N + h_0, \quad (7)$$

$$h_0 = \Big[\log\Big(\frac{q}{1-q}\Big) - \log\Big(\frac{q^*}{1-q^*}\Big)\Big]/\log\Big(\frac{p}{1-p}\Big). \quad (8)$$

Thus, agents that have not cascaded satisfy the next property.

Property 2: Until a cascade occurs, each agent follows its private signal, and h_n defined in (7) is a sufficient statistic of the information contained in $\mathcal{H}_{n-1}$.

Note in (8) that $h_0 = 0$ only when there is no ex-ante bias between actions, i.e., only when $q = q^*$. Otherwise, $h_0 > 0$ when $q > q^*$, i.e., when action Y is preferred over action N a priori; else $h_0 < 0$ when vice-versa.

Remark 1: The term h_0 in (8) reflects the ex-ante bias between actions, where $h_0 = 0$ implies that there is no ex-ante bias. Otherwise, $h_0 > 0$ implies that a priori, action Y is preferred over action N, and $h_0 < 0$ implies vice-versa.

The model with fake agents in [6] also exhibits Property 2, with h_n as its sufficient statistic, except that it is restricted to $h_0 = 0$. Further, as in [6], if $\epsilon = 0$ (no fake agents) then $a = b = p$ and $\eta = 1$. Whereas, if $\epsilon > 0$ then $\eta < 1$. It then follows from (7) that, due to the presence of fake agents, the information conveyed by a Y in an agent's observation history reduces by a factor of η, whereas the information conveyed by a N remains unaffected. This is because, unlike a N which always comes from an honest agent, a Y incurs the possibility that the agent could be fake. Further, this reduction in information from a Y is exacerbated with an increase in the possibility of fake agents, as η reduces with an increase in ϵ.

Substituting $g_n = \big(\frac{1-p}{p}\big)^{h_n}$ in Lemma 1, it follows that for all times n until a cascade occurs, $h_n \in [-1, 1]$ for $T = s$, and $h_n \in (-1, 1)$ for $T = h$, respectively. Further, it follows from (7) that at all such times n, the update rule for h_n is:

$$h_n = \begin{cases} h_{n-1} + \eta & \text{if } O_n = Y, \\ h_{n-1} - 1 & \text{if } O_n = N. \end{cases} \quad (9)$$

Whereas, for $T = s$, when $h_n > 1$ (< -1), likewise for $T = h$, when $h_n \geq 1$ (≤ -1), a $Y\,(N)$ cascade begins and h_n stops updating (Property 1). Here, h_0 defined in (8), is the fixed initial state of process $\{h_n\}$, since the first agent has no observation history. Now, if $h_0 > 1$ (< -1) for $T = s$, or if $h_0 \geq 1$ (≤ -1) for $T = h$, then a $Y\,(N)$ cascade begins from the first agent itself. In such cases, the history $\mathcal{H}_{n-1}$ for any agent n does not play any role in its decision; this makes the channel in Figure 1b irrelevant. We thus state the next remark.

Remark 2: We assume $h_0 \in (-1, 1)$ as otherwise, for either or both values of T, the presence of fake agents has no effect on rational agents.

Lastly, given the true value V, we denote the probability that a Y (N) cascade begins for process $\{h_n\}$ by $\mathbb{P}^V_{Y\text{-cas}}$ $\big(\mathbb{P}^V_{N\text{-cas}}\big)$. Here, $\mathbb{P}^V_{N\text{-cas}} = 1 - \mathbb{P}^V_{Y\text{-cas}}$ as it can be shown that $\{h_n\}$ eventually enters a cascade w.p. 1.

ISBN 978-3-903176-55-3 © 2023 IFIP

Authorized licensed use limited to: Northwestern University. Downloaded on February 13,2024 at 01:33:08 UTC from IEEE Xplore. Restrictions apply.

IV. Agent Welfare

Let the n^{th} agent's *welfare* refer to its pay-off averaged (in expectation) over $V \in \{G, B\}$, denoted by $\mathbb{E}[\pi_n]$. It can be shown that the *asymptotic welfare*, denoted by Π^T for rule T, relates to the cascade probabilities of process $\{h_n\}$ as:

$$\Pi^T := \lim_{n\to\infty} \mathbb{E}[\pi_n] = (x-C)q\mathbb{P}^G_{Y\text{-cas}} - (y+C)(1-q)\mathbb{P}^B_{Y\text{-cas}}. \quad (10)$$

Now, the bounds on h_0 (in Remark 2) ensure that it takes at least one time-step, starting from state h_0, to begin a cascade. This implies that the first agent always follow its private signal, and hence obtains the welfare:

$$\mathbb{E}[\pi_1] = F := (x-C)qp - (y+C)(1-q)(1-p). \quad (11)$$

In fact, F defined in (11) refers to the welfare for any agent n, if A_n always follows S_n disregarding the optimal decision rule in (3), i.e., $E[\pi_n|A_n$ always follows $S_n] = F$, for all n. The next property shows that $\mathbb{E}[\pi_n]$ is monotonic in the agents' indices.

Property 3: Given any T and h_0, the welfare of each agent is at least equal or greater than the welfare of its predecessors. Thus, $\mathbb{E}[\pi_n] \geq F$ and is non-decreasing in n.

Proof: Consider two consecutive agents, $n-1$ and n with their respective information sets I_n and I_{n-1}. Under the informational equivalance of their private signals: S_{n-1} and S_n, we have $I_{n-1} \subset I_n$. Property 3 is then proved by applying the celebrated Blackwell's Theorem on comparing information structures [15], which implies that it is sufficient to show that the signals from observing the smaller set I_{n-1} are obtained as a stochastic *mapping* (garbling) of the signals from the larger set I_n. Let $\bar{I}_n$ and $\bar{I}_{n-1}$ be the n and $(n-1)$-length random vectors corresponding to the observations sets I_n and I_{n-1} respectively, such that the two vectors share the first $n-1$ elements. Then, the desired mapping is given by $\bar{I}_{n-1} = G\bar{I}_n$, where G is a $(n-1)\times n$ diagonal matrix. Then, Blackwell's result for the corresponding optimal welfares states that $\mathbb{E}[\pi_n] \geq \mathbb{E}[\pi_{n-1}]$; and $\mathbb{E}[\pi_n] \geq F$ for all n follows from (11). ∎

V. Markovian Analysis of Cascades

In this section, we analyse the process $\{h_n\}$, given V, for the probability of cascades. We consider the tie-breaking rule $T = s$ for the sake of analysis. Now, it follows from Section III that conditioned on V, the process $\{h_n\}$ is a discrete-time Markov process taking values in $[-1, 1]$ before getting absorbed into the N *cascade region* (< -1) or the Y *cascade region* (> 1). Specifically, eq. (9) shows that, given V, $\{h_n\}$ is a random walk (r.w.) that starts from state h_0 and moves to the right by η w.p. $\mathbb{P}(O_n = Y|V)$ or to the left by 1 w.p. $\mathbb{P}(O_n = N|V)$ until a cascade occurs, where these probabilities are defined in terms of a and b in (4). Figure 2 depicts this random walk, where $p_f \triangleq \mathbb{P}(O_n = Y|V)$ denotes the probability of a Y being observed given V, when any agent n follows S_n. We have from (4): $p_f = a$ for $V = G$, whereas $p_f = 1-b$ for $V = B$.

Note that in special cases such as when h_0 and η satisfy $(1-|h_0|)/\eta = t$, $|h_0|/\eta = v$ for some $t, v \in \{0, 1, 2, \ldots\}$, the process $\{h_n\}$ is equivalent to a Markov chain with finite state-space $\mathcal{A} = \{-r-1, -r, \ldots, -1, 0, 1, \ldots, r, r+1\}$, where $r := t+v$ and $-r-1$ and $r+1$ are 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 generic case of non-integer values of $(1-|h_0|)/\eta$ and $|h_0|/\eta$ resulting in $\{h_n\}$ possibly taking countably infinite values in $[-1, 1]$.[2]

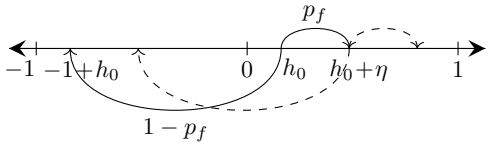

Fig. 2: Partial transition diagram of random walk $\{h_n\}$ given V.

Consider the special case where fake agents are absent ($\epsilon = 0$). Then, for the r.w. depicted in Fig. 2, it follows that $\eta = 1$, and $p_f = p$ if $V = G$, else $p_f = 1-p$ if $V = B$. Fig. 2 then implies that $\{h_n\}$ occupies an equivalent finite state-space $\mathcal{A}$, and thereby has a closed-form expression for $\mathbb{P}^V_{Y\text{-cas}}$ which for varying values of h_0 and $T \in \{s, h\}$ are tabulated below.

h_0	T	$\mathcal{A}$	$\mathbb{P}^V_{Y\text{-cas}}$
$(0,1)$	s,h	$\{-2,-1,0,1\}$	$p_f/[p_f+(1-p_f)^2]$
$(-1,0)$	s,h	$\{-1,0,1,2\}$	$p_f^2/[p_f^2+(1-p_f)]$
0	s	$\{-2,-1,0,1,2\}$	$p_f^2/[p_f^2+(1-p_f)^2]$
0	h	$\{-1,0,1\}$	p_f

TABLE I: Equivalent state-space $\mathcal{A}$ and $\mathbb{P}^V_{Y\text{-cas}}$ of r.w. $\{h_n\}$ under $\epsilon = 0$ for varying values of ex-ante bias h_0 and tie-breaking rule T.

In Table I, for $\mathcal{A}$ in each case, 0 is the initial state and the leftmost and the rightmost states are the N and Y cascade absorption states, respectively. Note that for cases where $h_0 \neq 0$, the tie-breaking rule T does not matter since $\{h_n\}$ for any n never equals 1 or -1, which are the only values at which a tie could occur. Lastly, the asymptotic welfare $\Pi^T(0, h_0)$ is obtained by using the expression for $\mathbb{P}^V_{Y\text{-cas}}$ from Table I, for the corresponding values of h_0 and T, in equation (10).

A. Thresholds $\{\epsilon_t\}$ under a non-negative ex-ante bias

In the case $h_0 = 0$ and $\epsilon = 0$ as in [1], $\eta = 1$ and so a Y cascade requires at least two consecutive Y's. However, when $h_0 > 0$, then starting from h_0, only a single Y triggers a Y cascade. However, sequences having one or more N's, but have not cascaded, still require at least two consecutive Y's for a Y cascade to begin. Now, as ϵ increases and reduces η, a greater number of consecutive Y's (≥ 1 for $h_0 > 0$ and ≥ 2 for $h_0 = 0$) may be required to cause a Y cascade. This

[2]For example, if η was chosen uniformly at random, then almost surely (w.p. 1) it would fall into this case.

ISBN 978-3-903176-55-3 © 2023 IFIP

Authorized licensed use limited to: Northwestern University. Downloaded on February 13,2024 at 01:33:08 UTC from IEEE Xplore. Restrictions apply.

is characterized by first defining an increasing sequence of ϵ-thresholds $\{\epsilon_t\}_{t=0}^{\infty}$, for a fixed signal quality p.

$$\epsilon_t = \frac{\alpha - \alpha^{\frac{1-h_0}{t}}}{\alpha^{\frac{1-h_0}{t}+1} - 1}, \quad \text{for } t = 1, 2, \ldots, \tag{12}$$

and $\epsilon_0 = 0$, with $\alpha := \frac{p}{1-p}$. Here, for $t \geq 1$, ϵ_t is such that at $\epsilon = \epsilon_t$, $\frac{1-h_0}{\eta} = t$. Now, starting from state h_0, $t+1$ consecutive Y's would be needed to begin a Y cascade only when $\frac{1-h_0}{t+1} < \eta \leq \frac{1-h_0}{t}$. This inequality when simplified in terms of ϵ implies that $\epsilon \in [\epsilon_t, \epsilon_{t+1})$ where ϵ_t is the t^{th} threshold, defined in (12).

Remark 3: For an ex-ante bias of $h_0 \geq 0$, if $\epsilon \in [\epsilon_t, \epsilon_{t+1})$, then starting from state h_0, $t+1$ consecutive Y's start a Y cascade. Here, integer t satisfies $t = \left\lfloor \frac{1-h_0}{\eta} \right\rfloor$.

The thresholds $\{\epsilon_r\}_{r=1}^{\infty}$ derived in [6], which does not consider any ex-ante bias, then become a special case of the thresholds $\{\epsilon_t\}$ in (12) when $h_0 = 0$.

B. Y cascade probability, $\mathbb{P}^V_{Y\text{-cas}}$

As depicted in Figure 2, the r.w. $\{h_n\}$ typically occupies a countably infinite state-space and does not readily allow for a closed-form solution to $\mathbb{P}^V_{Y\text{-cas}}$. So instead, we develop recursive equations that can compute $\mathbb{P}^V_{Y\text{-cas}}$ with arbitrary precision. These equations are motivated by the construction of an iterative method, which enumerates all possible observation sequences that would lead to Y cascade. This method, first developed in [6] for $h_0 = 0$ and $T = s$, is modified in this paper to also account for a possible ex-ante bias in actions ($h_0 \neq 0$). Refer to Appendix A for a detailed description of this iterative method. Later, in Appendix C, the case of $T = h$ is also considered. We use the iterative method to compute $\mathbb{P}^V_{Y\text{-cas}}(\epsilon, h_0)$ for $V \in \{G, B\}$. Then, substituting these values in eq. (10) yields the asymptotic welfare, $\Pi^T(\epsilon, h_0)$.

C. Improved learning due to fake agents in the presence of ex-ante bias

In [6], it has been numerically shown that with $T = s$, when there is no ex-ante bias ($h_0 = 0$), agents' asymptotic welfare deteriorates for any $\epsilon > 0$. Interestingly, when ex-ante bias between actions does exist ($h_0 \neq 0$), we observe that in certain cases, the presence of fake agents improves agents' asymptotic welfare. To demonstrate this, Figure 3 plots $\Pi^s(\epsilon, h_0)$ with respect to $\epsilon \in (0, 1)$ for two cases, namely, $h_0 = 0$ and $h_0 = 0.047$ (Y preferred over N a priori) under fixed signal quality $p = 0.7$ and cost structure: $x = 1, y = 0, C = 1/2$. Thus, $q^* = 1/2$ and we vary the value of h_0 by only varying the prior q in (8). Figure 3 also contrasts these welfares with the corresponding welfares in the absence of fake agents (computed by using $\mathbb{P}^V_{Y\text{-cas}}$ from Table I in eq. (10)), indicated as baselines. Observe that $\Pi^s(\epsilon, 0) < \Pi^s(0, 0)$ for all $\epsilon \in (0, 1)$, in line with the results in [6]. However, with ex-ante bias $h_0 = 0.047$, we see that there exist a range of values of ϵ at which $\Pi^s(\epsilon, h_0) > \Pi^s(0, h_0)$. An important reason for this ordering is the drop in the baseline welfare $\Pi^s(0, h_0)$ when h_0 is increased from 0 to 0.047. This drop

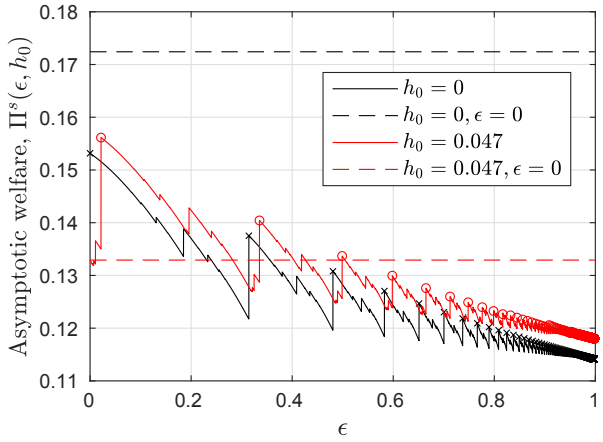

Fig. 3: Asymptotic welfare as a function of ϵ for $p = 0.7$, $T = s$ and indicated values of h_0. Π^s at thresholds $\{\epsilon_t\}$ for $h_0 = 0.047$ are marked by $\circ$, and for $h_0 = 0$ are marked by $\times$.

in turn occurs due to a change in the state-space $\mathcal{A}$ and the corresponding expression for $\mathbb{P}^V_{Y\text{-cas}}$ in Table I, when under $T = s$, the value of h_0 is changed from 0 to any value in $(0, 1)$.

Further, observe in Figure 3 that an abrupt and significant increase in Π^s occurs for $h_0 = 0$ (marked by $\times$) and for $h_0 = 0.047$ (marked by $\circ$) at their respective ϵ-threshold values: $\{\epsilon_t\}_{t=1}^{\infty}$, which are defined by (12). In both cases, the abrupt increase in asymptotic welfare at each of the thresholds can be attributed to a sudden increment (by 1) of the number of consecutive Y's required to trigger a Y cascade, starting from h_0 (Remark 3).

The next theorem shows that even when a priori, action Y is marginally preferred over action N, there exists an interval of ϵ-values in $(0, 1)$ such that the presence of fake agents w.p. ϵ improves the asymptotic welfare.

Theorem 1: Given an infinitesimal ex-ante bias towards action Y, i.e. $h_0 = 0^+$, and a fixed private signal quality $p \in (0.5, 1)$, there exists some $\underline{\epsilon} = f(p) > 0$ for some function f, such that

$$\Pi^s(\epsilon, 0^+) > \Pi^s(0, 0^+), \quad \forall \ \epsilon \in (0, \underline{\epsilon}). \tag{13}$$

Proof: We prove this theorem by showing that the limiting value of the function $\Pi^s(\epsilon, 0^+)$ as $\epsilon \to 0$ exceeds it's value when $\epsilon = 0$. For this, we first find the limiting value of $\mathbb{P}^V_{Y\text{-cas}}$ at $h_0 = 0^+$ and $\epsilon \to 0^+$, which is done in Lemma 2 in Appendix C and is restated here for convenience.

$$\lim_{\epsilon \to 0} \mathbb{P}^V_{Y\text{-cas}}(\epsilon, 0^+) = p_f^2 \frac{1 + (1 - p_f)p_f}{1 - (1 - p_f)p_f}. \tag{14}$$

Next, the value of $\mathbb{P}^V_{Y\text{-cas}}$ as $h_0 = 0^+$ and $\epsilon = 0$ can be obtained from Table I as follows.

$$\mathbb{P}^V_{Y\text{-cas}}(0, 0^+) = \frac{p_f}{p_f + (1 - p_f)^2}. \tag{15}$$

Now, note that at $h_0 = 0^+$, $q = q^*$ which implies that $(x -$

ISBN 978-3-903176-55-3 © 2023 IFIP

Authorized licensed use limited to: Northwestern University. Downloaded on February 13,2024 at 01:33:08 UTC from IEEE Xplore. Restrictions apply.

$C)q = (y + C)(1 - q)$. Thus, eq. (10) for $h_0 = 0^+$ and any ϵ simplifies to:

$$\Pi^s(\epsilon, 0^+) = (x - C)q\big[\mathbb{P}^G_{Y\text{-cas}}(\epsilon, 0^+) - \mathbb{P}^B_{Y\text{-cas}}(\epsilon, 0^+)\big]. \quad (16)$$

Then, the difference in the asymptotic welfares as $\epsilon \to 0$ and at $\epsilon = 0$ can be expressed as:

$$\lim_{\epsilon \to 0} \Pi^s(\epsilon, 0^+) - \Pi^s(0, 0^+) = (x - C)q[\Delta^G - \Delta^B], \quad (17)$$

where Δ^V for $V \in \{G, B\}$ is defined as:

$$\Delta^V := \lim_{\epsilon \to 0} \mathbb{P}^V_{Y\text{-cas}}(\epsilon, 0^+) - \mathbb{P}^V_{Y\text{-cas}}(0, 0^+), \quad (18)$$

$$\overset{(a)}{=} \frac{-p_f(1 - p_f)(1 - p_f^2)}{1 - p_f(1 - p_f)}. \quad (19)$$

Step (a) follows by substituting the expressions obtained in eq. (14) and (15) in eq. (18). Now, with $p_f = p$ for $V = G$ and $p_f = 1 - p$ for $V = B$, we substitute (19) in (17), which simplfies to give:

$$\lim_{\epsilon \to 0} \Pi^s(\epsilon, 0^+) - \Pi^s(0, 0^+) = (x - C)q \frac{p(1 - p)(2p - 1)}{1 - p(1 - p)}, \quad (20)$$

$$> 0, \qquad \forall\, p \in (0.5, 1). \quad (21)$$

The inequality in (21) follows from the fact $x > C$ and $p > 0.5$. Thus, there exists some $\underline{\epsilon} > 0$ such that (13) holds true. ■

D. Improved learning due to fake agents when rational agents follow the history to break ties

In this subsection, we investigate agents' welfare for the case: $T = h$ and $h_0 = 0$. To compute $\mathbb{P}^V_{Y\text{-cas}}$ for this case, the iterative method in Appendix A requires a few modifications to account for the change in tie-breaking rule, i.e., for $T = h$. Appendix B describes them in detail. Now, equipped with these modifications, the iterative method computes $\mathbb{P}^V_{Y\text{-cas}}$, and welfare $\Pi^T(\epsilon, 0)$. Figure 4 plots $\Pi^T(\epsilon, 0)$ with respect to $\epsilon \in (0, 1)$ for $T = s$ and $T = h$, given $p = 0.7$ and cost structure: $x = 1, y = 0, C = 1/2$. Here, as $h_0 = 0$ in (8), $q = q^* = 1/2$. Figure 4 also contrasts these welfares with the respective welfares in the absence of fake agents, indicated as baselines. As in Figure 3, we observe that when $T = s$, $\Pi^s(\epsilon, 0) < \Pi^s(0, 0)$ for all $\epsilon \in (0, 1)$. However, when $T = h$, we see that the inequality gets reversed, i.e., $\Pi^h(\epsilon, 0) > \Pi^h(0, 0)$, for all $\epsilon \in (0, 1)$. This implies that if agents prefer to follow their history to break ties, then presence of fake agents in any amount $\epsilon \in (0, 1)$ improves agents' asymptotic welfare. The next theorem formalizes this property for any p, and more specifically shows that such an improvement in welfare due to any $\epsilon \in (0, 1)$ occurs not only as $n \to \infty$ but also occurs for every agent n.

Theorem 2: Given $T = h$, $h_0 = 0$ and a fixed private signal quality $p \in (0.5, 1)$, agent n's welfare π^h_n satisfies:

$$\pi^h_n(\epsilon, 0) \geq \pi^h_n(0, 0) = F, \qquad \forall\, \epsilon \in (0, 1), n \in \mathbb{N}, \quad (22)$$

and the improvement in welfare, $\Delta_n(\epsilon) := \pi^h_n(\epsilon, 0) - \pi^h_n(0, 0)$ is monotonic and non-decreasing in n. F is defined in (11).

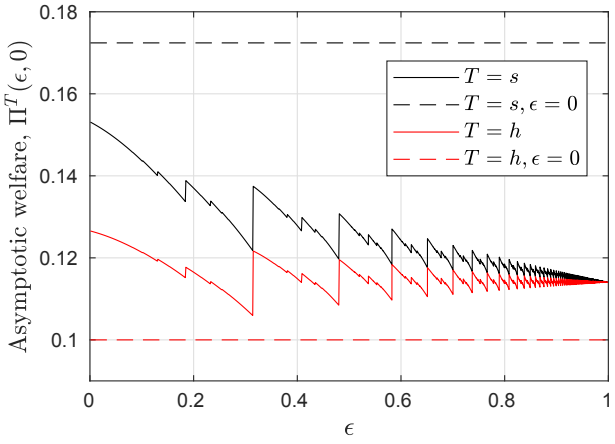

Fig. 4: Asymptotic welfare as a function of ϵ for $p = 0.7$, $h_0 = 0$ and indicated values of T.

Proof: First, note that for $T = h$, $h_0 = 0$, the state-space $\mathcal{A}$ in Table I implies that a cascade begins immediately after the first agent follows its private signal. That is, $A_n = Y(N)$ for all n, if $S_1 = H(L)$, which is informationally equivalent to every agent following its private signal. Thus, $\pi^h_n(0, 0) = F$, for all n. Second, by Property 3, $\pi^h_n(\epsilon, 0) \geq F$ and is monotonic and non-decreasing in n. This concludes the proof. ■

VI. Conclusions

We revisited the model for observational learning with fake agents from [6]. We showed that while in [6], fake agents reduced the welfare of rational agents or equivalently caused poorer learning, this conclusion may not hold when rational agents either have an ex-ante bias or employ a different tie-breaking rule. In particular, we proved that when there is no ex-ante bias as in [6], if rational agents follow their history to break ties, then the presence of fake agents in any amount always leads to improved learning, and thus a higher welfare. On the other hand, under an ex-ante bias, when rational agents follow their private signals to break ties, we show that there exists a range for the amount of fake agents, in which better learning occurs. These cases are other examples of a game setting in which better quality information may lead to lower pay-offs for the players of that game.

In future work, we plan on determining sharp characterizations of when degraded observations lead to lower welfare. Other potential future directions include considering other forms for noisy observations such as those considered in [5], non-Bayesian rationality and random tie-breaking rules.

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] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.

ISBN 978-3-903176-55-3 © 2023 IFIP

Authorized licensed use limited to: Northwestern University. Downloaded on February 13,2024 at 01:33:08 UTC from IEEE Xplore. Restrictions apply.

[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] 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.
[6] P. Poojary and R. Berry, "Observational learning with fake agents," *CoRR*, vol. abs/2005.05518, 2020. [Online]. Available: https://arxiv.org/abs/2005.05518
[7] I. Welch, "Sequential sales, learning, and cascades," *The Journal of finance*, vol. 47, no. 2, pp. 695–732, 1992.
[8] 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.
[9] S.-T. Su, V. G. Subramanian, and G. Schoenebeck, "Social learning with questions," in *Proceedings of the 14th Workshop on the Economics of Networks, Systems and Computation*, 2019.
[10] P. Poojary and R. Berry, "Observational learning with negative externalities," in *2022 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2022, pp. 1495–1496.
[11] I. Bistritz 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.
[12] 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.
[13] Y. Song, "Social learning with endogenous network formation," *CoRR*, vol. abs/1504.05222, 2015. [Online]. http://arxiv.org/abs/1504.05222
[14] P. Poojary and R. Berry, "Observational learning with fake agents," in *2020 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2020, pp. 1373–1378.
[15] M. Leshno and Y. Spector, "An elementary proof of blackwell's theorem," *Mathematical Social Sciences*, vol. 25, no. 1, pp. 95–98, 1992.
[16] D. Blackwell, "Equivalent comparisons of experiments," *The annals of mathematical statistics*, pp. 265–272, 1953.
[17] O. Gossner, "Comparison of information structures," *Games and Economic Behavior*, vol. 30, no. 1, pp. 44–63, 2000.

APPENDIX A
ITERATIVE METHOD TO COMPUTE THE Y CASCADE PROBABILITY $\mathbb{P}^V_{Y\text{-CAS}}$, UNDER $T = s$ AND EX-ANTE BIAS h_0

The iterative method presented here computes $\mathbb{P}^V_{Y\text{-cas}}$ for r.w. $\{h_n\}$, given $T = s$, $h_0 \in (-1, 1)$, by enumerating all possible sequences that can lead to a Y cascade. The enumeration of all such sequences is depicted in a stage-wise manner in Figure 5. First, consider the case $h_0 \geq 0$, for which Stage 0 does not contain any sequence, i.e., r_0 is set to 0. The iterative method thus begins from Stage 1, with the process $\{h_n\}$ starting from state h_0. We initialize Stage 1 with $r_1 = t + 1$. As a result, The first sequence of r_1 consecutive Y's, denoted by Y^{r_1}, clearly enters the Y cascade region (Remark 3), and so $r_1\eta \in [1, 1+\eta]$. The rest of the sequences, each of length $r_1 + 1$, are simply permutations of each other that contain only a single N. This is because two N's or more are not possible without entering the N cascade region. Now, each of these r_1 distinct sequences results in the same net shift of $r_1\eta - 1$, which ends in the region $[0, \eta]$ as we know that $r_1\eta \in [1, 1 + \eta]$. This completes Stage 1. For each Stage n that follows, we define r_n in a manner similar to r_1, as follows:

Definition 2: The integer r_n, for $n = 1, 2, \ldots$ is the number of consecutive Y's required to enter the Y cascade region in Stage n of Figure 5.

The sequences in Stage n are then enumerated exactly as in Stage 1, except that r_n now replaces r_1. This ensures that at the end of each Stage $j = 1, 2, \ldots$, any sequence that of length $n_j = \sum_{i=1}^{j}(1 + r_i)$ that has not yet begun a cascade satisifies: $h_{n_j} \in [0, \eta]$. This resets the r.w. for the next stage. In this manner, all sequences that lead to a Y cascade are enumerated.

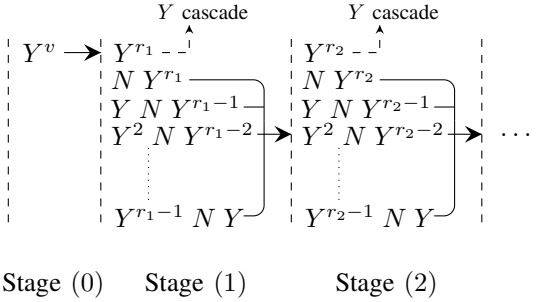

Fig. 5: An enumeration of all possible sequences that would lead to a Y cascade. The term Y^t represents a sequence of t consecutive Y's. The sequence $\{r_n\}$ is defined as per (23) for $h_0 \geq 0$, and as per (24) for $h_0 < 0$. Stage (0) is only applicable when $h_0 < 0$.

To obtain the values of $\{r_n\}$, we first initialize $r_1 = t + 1$. Now, at the end of each Stage j, given that a Y cascade has not yet begun, we know that $h_{n_j} \in [0, \eta]$. From here, it would take either r or $r + 1$ consecutive Y's to enter the Y cascade region, where integer $r := \lfloor 1/\eta \rfloor$. Therefore, $r_n \in \{r, r + 1\}$ for $n = 2, 3, \ldots$. Then, with $r_1 = t + 1$, successive values of r_n for $n = 2, 3, \ldots$ are obtained by applying Definition 2, which can be expressed as follows:

$$r_n = \begin{cases} r, & \text{if } \ h_0 + \sum_{i=1}^{n-1}(r_i\eta - 1) + r\eta > 1, \\ r+1, & \text{o.w.} \end{cases} \tag{23}$$

For the other case: $h_0 < 0$, the only change to the enumerations in Figure 5 is the inclusion of Stage 0, which is the sequence Y^{r_0}. Here, $r_0 = v := \lfloor h_0/\eta \rfloor$, is the least number of consecutive Y's such that $h_v \in [0, \eta]$. This is because a N cannot occur at any $i \leq v$ without causing a N cascade. Moreover, this ensures that at the end of Stage 0, $h_v \in [0, \eta]$, which as stated earlier, is required to reset the r.w. for the next stage. Now, we apply Definition 2 to obtain the values of $\{r_n\}$, . Then, having obtained the values: $\{r_i\}_{i=1}^{n-1}$, r_n as per Definition 2 is given by:

$$r_n = \begin{cases} r, & \text{if } \ h_0 + v\eta + \sum_{i=1}^{n-1}(r_i\eta - 1) + r\eta > 1, \\ r+1, & \text{o.w.} \end{cases} \tag{24}$$

We now proceed to compute $\mathbb{P}^V_{Y\text{-cas}}$ as follows. Let P_n denote the probability that a sequence in Fig. 5 terminates in a Y cascade, but not before the n^{th} stage. The following recursion then holds.

$$P_n = p_f^{r_n}\left[1 + r_n(1 - p_f)P_{n+1}\right], \ \text{ for } n = 1, 2, \ldots \tag{25}$$

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

$$\mathbb{P}^V_{Y\text{-cas}}(\epsilon, h_0) = \begin{cases} P_1, & \text{for } \ h_0 \geq 0, \\ p_f^v P_1, & \text{for } \ h_0 < 0. \end{cases} \tag{26}$$

ISBN 978-3-903176-55-3 © 2023 IFIP

Authorized licensed use limited to: Northwestern University. Downloaded on February 13,2024 at 01:33:08 UTC from IEEE Xplore. Restrictions apply.

Note that the factor p_f^v for $h_0 < 0$ corresponds to Stage 0 being included as a prefix to sequences in the subsequent stages.

Since (25) is an infinite recursion, to compute $\mathbb{P}^V_{Y\text{-cas}}$ in practice, we truncate the process to a finite number of iterations M, as done in [6]. To this end, we first assume that $P_{M+1} = 1$. Next, we use (25) to successively compute P_k while k counts down from M to 1, performing a total of M iterations. We denote the obtained value as ${}^{(M)}\mathbb{P}^V_{Y\text{-cas}}$. The analysis in [6] shows that for $h_0 = 0$, ${}^{(M)}\mathbb{P}^V_{Y\text{-cas}}$ is in fact a tight upper bound to $\mathbb{P}^V_{Y\text{-cas}}$ as $M \to \infty$ for any $p \in (0.5, 1)$ and $\epsilon \in [0, 1)$. Moreover, the difference ${}^{(M)}\mathbb{P}^V_{Y\text{-cas}} - \mathbb{P}^V_{Y\text{-cas}}$ decays to zero at least as fast as $\{0.5^M\}$, in the number of iterations M. With minor modifications to this analysis, it can be shown that the same result extends to the more general case of $h_0 \in [-1, 1]$ considered here. For the plots presented in this paper, we use $M = 10$ to compute $\mathbb{P}^V_{Y\text{-cas}}(\epsilon, h_0)$ for $V \in \{G, B\}$ using the above recursive method, which gives an error of less than 10^{-3}. Then, substituting these obtained values in equation (10) yields the asymptotic welfare, $\Pi^s(\epsilon, h_0)$.

Appendix B
Modifications to the iterative method for $T = h$

In this section, we discuss how the iterative method for computing $\mathbb{P}^V_{Y\text{-cas}}$, outlined in Appendix A, can be modified for $T = h$, but with $h_0 = 0$. Now, the first modification is in (23), which computes the values $\{r_i\}$. Here, the inequality condition to be satisfied for $r_n = r$ changes from > 1 to ≥ 1. This reflects the change in the Y cascade region from > 1 for $T = s$ to ≥ 1 for $T = h$. The second modification is in the enumerations depicted in Figure 5, which is to discard the sub-sequence NY^{r_1} from Stage (1). This is because, with $T = h$, starting from state 0, an N would trigger a N cascade. To reflect this change, the recursion in (25) which computes the values $\{S_n\}$ has to be modified for $n = 1$ as follows:

$$P_1 = p_f^{r_1}\big[1 + (r_1 - 1)(1 - p_f)P_2\big] = \mathbb{P}^V_{Y\text{-cas}}. \qquad (27)$$

Now, we assume $\epsilon \in \{\epsilon : 1/\eta \in \mathbb{Q}\}$ to avoid the case of rational values of $1/\eta$. This ensures that at the start of any later stage, $h_n \neq 0$. This in turn ensures that the N in the sub-sequence NY^{r_j} in Stage (j), for any $j > 1$ does not trigger a N cascade. Else the sub-sequence would also have to be discarded. As a result, recursion in (25) remains unchanged for all $n > 1$. Equipped with the above modifications, the iterative method in Appendix A can now compute $\mathbb{P}^V_{Y\text{-cas}}$, for $V \in \{G, B\}$, when $T = h$ and $h_0 = 0$.

Appendix C

Lemma 2: Given the ex-ante bias h_0 and the fraction of fake agents ϵ, the limiting value of the Y cascade probability, $\mathbb{P}^V_{Y\text{-cas}}$ at $h_0 = 0^+$ and $\epsilon \to 0^+$ is given by

$$\lim_{\epsilon \to 0} \mathbb{P}^V_{Y\text{-cas}}(\epsilon, 0^+) = p_f^2 \frac{1 + (1 - p_f)p_f}{1 - (1 - p_f)p_f}. \qquad (28)$$

Proof: To find the limiting value of $\mathbb{P}^V_{Y\text{-cas}}$ at $h_0 = 0^+$ and $\epsilon \to 0^+$, we first show that in this limiting regime, for all $i \geq 2$, r_i in (23) satisfies: $r_i \to 1$ as $\epsilon \to 0$. To prove this, assume $r_j = 1$ for all $2 \leq j \leq i - 1$ and note that $r_1 = t + 1 = 2$ as $t = 1$. Then it follows from (23) that $r_i = 1$ only if $\eta > \left(1 + \frac{1}{i}\right)^{-1}$. Now as $\epsilon \to 0^+$, $\eta \to 1$ and hence this condition is satisfied. Using this argument inductively shows that $r_i = 1$ for all $i \geq 2$.

Now, with $r_i \to 1$ for all $i \geq 2$ as $\epsilon \to 0$, it thus follows that the recursion in (25) results in the same infinite computation to obtain P_n as for P_{n+1}, for all $i \geq 2$. Thus, all P_n for $n \geq 2$ have the same value which satisfies: $P_n = p_f\big[1 + (1 - p_f)P_n\big]$. Solving this equation for $n = 2$ gives the value for P_2. Using this value in equation (25) for $n = 1$ yields P_1, which from (26) is in fact $\mathbb{P}^V_{Y\text{-cas}}$ and is as follows

$$\lim_{\epsilon \to 0} \mathbb{P}^V_{Y\text{-cas}}(\epsilon, 0^+) = p_f^2 \frac{1 + (1 - p_f)p_f}{1 - (1 - p_f)p_f}.$$

■

ISBN 978-3-903176-55-3 © 2023 IFIP

Authorized licensed use limited to: Northwestern University. Downloaded on February 13,2024 at 01:33:08 UTC from IEEE Xplore. Restrictions apply.