

Impact of Fake Agents on Information Cascades

Pawan Poojary, *Member, IEEE*, and Randall Berry, *Fellow, IEEE*

Abstract—In online markets, agents often learn from other's actions in addition to their private information. Such observational learning can lead to *herding* or *information cascades* in which agents eventually ignore their private information and “follow the crowd”. Models for such cascades have been well studied for Bayes-rational agents that arrive sequentially and choose pay-off optimal actions. This paper additionally considers the presence of *fake agents* that take a fixed action in order to influence subsequent rational agents towards their preferred action. We characterize how the fraction of such fake agents impacts the behavior of rational agents given a fixed quality of private information. Our model results in a Markov chain with a countably infinite state space, for which we give an iterative method to compute an agent's chances of herding and its welfare (expected pay-off). Our main result shows a counter-intuitive phenomenon: there exist infinitely many scenarios where an increase in the fraction of fake agents in fact reduces the chances of their preferred outcome. Moreover, this increase causes a significant improvement in the welfare of every rational agent. Hence, this increase is not only counter-productive for the fake agents but is also beneficial to the rational agents.

Index Terms—Information cascades, herding, Bayesian optimality, Perfect Bayesian Equilibrium (PBE).

I. INTRODUCTION

CONSIDER a new item that is up for sale in a recommendation-based market where agents arrive sequentially and decide whether to buy the item, with their choice serving as a recommendation for later agents (e.g., through a public database). This item has a common quality/utility (say “good” or “bad”) that is unknown to the agents. Each agent then makes a pay-off optimal decision by using its own prior knowledge of the item's quality and by observing the choices of its predecessors. Such models of “observational learning” were first studied by [2], [3], [4] under a Bayesian learning framework wherein each agent has some prior knowledge in the form of a privately observed signal about a pay-off-relevant state of the world generated from a commonly known probability distribution. Agents arrive in an exogenous order, and every agent chooses its action based on its own private signal and observations of past agents' actions. Agents are assumed to be Bayes *rational*, i.e., their actions are optimal with respect to their posterior beliefs of the true state of the

Manuscript received November 23, 2023; revised June 29, 2024; accepted February 19, 2025. Date of publication Month DD, YYYY; date of current version Month DD, YYYY. This work was supported in part by the U.S. National Science Foundation under grants ECCS-2030251, ECCS-2216970-002 and ECCS-2331106. This paper was presented in part at the IEEE International Symposium on Information Theory, Los Angeles, CA, USA, June 2020 [1]. The editor coordinating the review of this manuscript and approving it for publication was Dr.

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

Digital Object Identifier

world given their observations¹. A key result for such models is the emergence of *information cascades* or *herding*, i.e., at some point, it is optimal for an agent to ignore its own private signal and follow the actions of the past agents. Subsequent agents then follow suit due to their homogeneity. As a result, from the onset of a cascade, the agents' actions do not reveal any information conveyed to them by their private signals; hence learning stops.

The model described above has two possible outcomes. First, the agents may end up in a *correct* cascade, that is the information accumulated within the observation history eventually forces all successive agents to take the optimal action with respect to the underlying true state of the world. In this case, *learning* or a socially optimal behaviour is said to be achieved. Second, there is also the possibility that past observations might get accumulated in a manner that forces all successive agents to take the action that, although individually optimal, is in fact socially sub-optimal, which we refer to as a *wrong* cascade. Thus, herding to a wrong cascade would prevent the agents from learning the socially optimal (correct) choice. Moreover, a main result of work in [2], [3], [4] (and others) that assumes homogeneous agents and discrete bounded private signals is that a cascade (correct or wrong) commences within the arrival of a finite number of agents with probability one. This leads to a positive probability that a wrong cascade occurs.

In this paper, we study a Bayesian learning model similar to [2], [3], [4] where in addition to rational agents, we introduce randomly arriving *fake agents* that always take (or fake) a fixed action, regardless of their pay-off, in order to influence the outcome of a cascade. For example, this could model a setting where a poor quality item is posted for sale on a website that records buyers' decisions, and fake agents intentionally buy (or appear to buy) this item to make it seem worth buying to future buyers [5]. The model could also represent a setting where rational actions are manipulated or where fake actions are inserted into a recommendation system [6], [7], [8]. The objective is to study the impact of varying the amount of these fake agents on the probability of their preferred cascade and on the resulting *welfares* (expected pay-offs) of the rational agents. Our main result shows a counter-intuitive phenomenon: the probability with which the cascade preferred by the fake agents occurs is not monotonically increasing in the fraction of fake agents, ϵ . In fact, there exist infinitely many cases where an increase in the fraction of fake agents reduces the chances of their preferred cascade. We identify a sequence of thresholds for ϵ where this phenomenon is most pronounced.

¹From a game theoretic point-of-view, this is a dynamic game with asymmetric information where the agents' optimal actions are a Perfect Bayesian Equilibrium (PBE) of this game.

Moreover, we show that exceeding any such ϵ threshold causes an abrupt improvement in the welfare of every rational agent.

In our model, we assume that the presence of fake agents is common knowledge to the buyers and is accounted for in their actions. This is motivated by several empirical results, such as in [9] and [10], which suggest that over repeated interactions with the buying platform and through word-of-mouth, buyers tend to be conservative with their decisions by taking into account that a certain fraction of positive reviews for an item might be fake or that its sales statistics might be exaggerated.

A. Our Contributions

In this paper, we develop a Markov chain model to represent the process of information accumulation driven by the actions of sequentially arriving agents (either rational or fake), which eventually results in a cascade. To obtain this Markov chain, we identify a sufficient statistic of past observations that determines whether an agent cascades or follows its private signal.

Secondly, we analyse the Markov chain for the probability of cascades as a function of the quality p of the prior knowledge available to every rational agent (i.e., its private signal), and the fraction of fake agents, ϵ . The Markov chain typically occupies a countably infinite state-space, and does not readily allow for a closed-form solution to the cascade probabilities. Instead, we develop recursive equations that can compute the cascade probabilities with arbitrary precision. These equations are motivated by the construction of an iterative method that enumerates all possible action sequences that would lead to a cascade. This iterative method also provides exact probabilities for both correct and wrong cascades to begin within the arrival of a finite number of agents. Computing these cascade probabilities in turn yields the welfare for any arbitrary agent.

Thirdly, we study the effects of varying ϵ on the cascade probabilities and on agents' welfares for a fixed private signal quality, p . Our results highlight the following counter-intuitive phenomenon: there exist an infinite sequence of thresholds $\{\epsilon_r\}_{r=1}^\infty$ where increasing ϵ slightly above ϵ_r causes an abrupt decrease in the probability of the cascade preferred by the fake agents. Thus, a marginal increase in the presence of fake agents beyond any such threshold reduces the chances of their preferred cascade instead of increasing it, which is what they had intended. Further, we analytically show that increasing ϵ just beyond any such ϵ_r also causes an abrupt and significant increase in the welfares of all agents. Therefore, marginally exceeding ϵ_r is not only counter-productive for the fake agents but is also beneficial to the rational agents.

Fourthly, we prove that an agent's welfare improves if it chooses to arrive later in the sequence of arrivals. This result implies that the welfare reaches a limiting maximum value as the agent's arrival index tends to infinity. We refer to this value as the long run or *asymptotic* welfare. The proof involves the application of Blackwell's theorem on comparing information structures [11], [12], which more generally also shows that an agent can never do any better by ignoring any of the past observations or its own signal.

Fifthly, we quantify the cascade probabilities and social welfare as a function of p in the interesting scenario where

the proportion of fake agents approaches unity. We analytically show that even when fake agents have overwhelmed the rational agents: (i) a better signal quality results in higher chances of learning the socially optimal action (correct cascade) and (ii) rational agents continue to benefit from past observations.

Sixthly, we discuss the role of a *Platform Co-ordinator* in improving learning through modifications to the observation database. We analyse and compare the performances of three approaches for improving the agents' welfares, namely (a) increasing the fraction of fake actions, (b) filtering out possibly fake actions and (c) modifying the possibly fake actions. For a wide range of values for p , we observe that for low values of ϵ , Scenario (c) outperforms all other scenarios, providing the best improvement in welfare. As $\epsilon \rightarrow 0$, it entirely mitigates the reduction in welfare caused by the fake agents. Whereas, for high values of ϵ , Scenario (b) provides the best improvement. There also exist several intervals with moderate ϵ -values where Scenario (a) performs the best.

We conclude by showing that the analysis, results and discussions presented in this work readily extend to general priors and agent pay-offs, as long as the ex-ante pay-off is zero. Moreover, our analytic techniques can be easily modified to include a non-zero ex-ante pay-off.

B. Related Work

Our work follows the basic model studied in [2], [3], [4], except we introduce fake agents which arrive at random amidst the sequence of rational agents. Many variations of this basic model have been studied, some of which we review here briefly. The work in [13] relaxed the assumptions of agents' homogeneity and binary valued private signals made in [2], [3], [4]. It showed in particular that allowing for a richer class of signals, such that their likelihood ratio is unbounded, could result in learning to occur with probability one. Our model maintains the assumptions of [2], [3], [4], i.e., homogenous agents and binary valued private signals.

Another change to the basic model is to consider different structures for observing past actions. For example, [14] considers that agents can observe only a subset of the past actions, defined by an underlying network structure. This work finds conditions on the network structure that guarantee asymptotic learning. Whereas [15] makes the network formation endogenous by allowing agents to select their observations at a cost. Another variation, studied in [16], forces a fixed number of initial agents to only observe their private signals. These agents thereby act as "guinea pigs", used to explore the unknown true value. This causes improved welfares of the subsequent agents at the cost of the guinea pigs' welfares. Our work here stays with the original model [2], [3], [4] which assumes that actions are recorded to a common database, thereby allowing each agent to observe all prior actions (although in our case, each action could be either rational or fake).

Closer to our work is the model in [17], which assumes that the recording of actions for subsequent agents is subject to an error that is unbiased towards each of the possible actions. In our setting, an action being either fake or rational depending on the agent-type could equivalently be perceived as an error

while recording a rational action (as in [17]); except that in our case, the error is biased only towards a preferred action.²

There is also a body of work that considers agents similar to our fake agents, who only take a single action regardless of the true state. This includes the *crazy agents* considered in [13], *stubborn agents* in [18] and *zealots* in [19]. While [13] relaxes the assumption of binary signals, it does not consider how changing the fraction of crazy agents affects the cascade probability, which is the main focus of our work. The works in [18] and [19] do not consider learning, instead they model opinion/belief dynamics in the presence of their respective types of abnormal agents. They consider non-Bayesian models for updating agents' beliefs, while our work remains with the Bayesian model. Other types of agents considered in the literature include the *revealers* in [20], *experts* and *laymen* in [21] and *non-myopic* agents in [22].

The learning models considered in this line of work also have ties to early work on sequential detection with finite memory, e.g, [23] and [24]. There, the sequence of i.i.d. signals has to be summarized by a test-statistic of finite cardinality, which gets updated as per a rule designed by a planner with the objective of maximizing the chance of asymptotically learning the true state of the world. Our model does not assume any memory constraints and all past actions are perfectly observed. The more important distinction is that in [23], [24], agents simply follow the rule prescribed by the planner, i.e., agents are not strategic and thus their actions might not constitute an equilibrium. Our work instead considers a setting where there is no planner and agents are strategic and act in their best interests. Thus, agents' actions are in a PBE.

Our work is a substantial extension of [1], where we first proposed the iterative method to compute cascade probabilities and identified the infinite sequence of ϵ -thresholds at which an abrupt reduction in the preferred cascade probability occurs. This paper makes several contributions beyond [1], such as studying the cascade probability for finite agents, identifying important properties exhibited by agents' welfares, and exploring the effects of a Platform Co-ordinator on learning, among others as stated earlier.

C. Organisation

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 present our Markov chain formulation, identify error thresholds and devise an iterative method to compute cascade probabilities. Further, in Subsection IV-C, we quantify the abrupt reduction in chances of a correct cascade at these error thresholds. Subsection IV-D investigates the preferred cascade probability for low values of ϵ and contrasts it to the case when fake agents are absent ($\epsilon = 0$). Section V characterizes agents' welfares and identifies important properties exhibited by them. In Section VI, we investigate learning in the limiting scenario where the proportion of fake agents approaches one.

²This change requires a different analysis approach than that in [17] as the underlying Markov chain now typically has a countably infinite state-space, while in [17] it was finite. We show that this change also yields substantially different outcomes.

Section VII introduces a Platform Co-ordinator and presents approaches for modifying the observations that could improve learning. Section VIII extends this work to general priors and agent pay-offs, while maintaining zero ex-ante pay-off. Lastly, we present our conclusions in Section IX. Detailed proofs, extended analyses and supporting plots are provided in the Appendices of [25].

II. MODEL

We consider a model similar to [2] in which there is a countable sequence of agents, indexed $i = 1, 2, \dots$ where the index represents both the time and the order of actions. Each agent i takes an action A_i of either buying (Y) or not buying (N) a new item that has a true value (V), which could either be good (G) or bad (B). For simplicity, both possibilities of V are assumed to be equally likely.

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 $C = 1/2$ for buying the item and gains an amount that reflects the item's value/utility to its buyer. The buyer loses an amount $y = 0$ if $V = B$ and gains an amount $x = 1$ if $V = G$. The agent's net pay-off is given by

$$\pi_i = \begin{cases} x - C = 1/2, & \text{if } A_i = Y \text{ and } V = G, \\ -y - C = -1/2, & \text{if } A_i = Y \text{ and } V = B, \\ 0, & \text{if } A_i = N. \end{cases} \quad (1)$$

Given the values considered for x , y and C , observe that since V is equiprobable, the *ex ante* expected pay-off for any agent is 0 for either of the actions.³ Thus, to begin with, an agent is indifferent to the two actions.

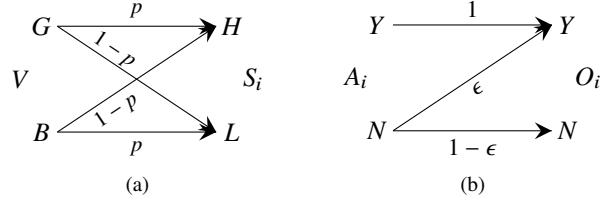


Fig. 1: (a) The BSC through which agents receive private signals. (b) The channel through which agents' actions are corrupted.

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 $\{S_1, S_2, \dots\}$ 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 $\{O_1, O_2, \dots, O_{i-1}\}$ of actions $\{A_1, A_2, \dots, A_{i-1}\}$. Next, we modify the information structure in [2] by assuming that at each time instant, an

³Section VIII generalizes the model to possibly non-uniform priors for V and a general pay-off structure for agents, while still retaining the condition of zero ex-ante pay-off.

agent could either be *fake* with probability (w.p.) $\epsilon \in [0, 1]$ or *ordinary* w.p. $1-\epsilon$, where the value ϵ is assumed to be common knowledge, so that all agents know the probability that any agent is fake but do not know which specific agents are fake. 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 with probability $1 - \epsilon$, the reported action $O_i = N$ and with probability ϵ , $O_i = Y$. Whereas, if $A_i = Y$ then $O_i = Y$ with probability 1. Refer to Figure 1b.

An equivalent model is where action A_i is rational only if agent i is ordinary and is fixed to Y otherwise, while $O_i = A_i$ for all agents. This yields the same information structure as the model above and so the same analysis applies to model the behavior of the ordinary agents. We chose the former model mainly to simplify our notation.

III. OPTIMAL DECISION, CASCADES AND WELFARE

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

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

Note that when $\gamma_n = 1/2$, an agent is indifferent to the two actions. Similar to [17], our decision rule in this case follows the private signal S_n , unlike [2], which employs a randomized tie-breaking rule. Another choice in this case is to follow the history \mathcal{H}_{n-1} , *i.e.*, to take the action that is most observed in the past. Techniques in this paper can be readily adapted to reflect this alternate choice of breaking ties [26].

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

It follows from (2) that, agent n cascades to a Y (N) if and only if $\gamma_n > 1/2$ ($< 1/2$) for all $S_n \in \{H, L\}$. The other case being $\gamma_n \geq 1/2$ for $S_n = H$ and $\gamma_n \leq 1/2$ for $S_n = L$; in which case, agent n follows S_n . A more intuitive way to present this condition is to first express the information contained in the history \mathcal{H}_{n-1} observed by agent n in the form of its public likelihood ratio,

$$l_{n-1}(\mathcal{H}_{n-1}) := \frac{\mathbb{P}(\mathcal{H}_{n-1}|B)}{\mathbb{P}(\mathcal{H}_{n-1}|G)}, \quad (3)$$

and then state it as follows.

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

To prove this lemma, first define agent n 's private likelihood ratio, $\beta_n(S_n) := \mathbb{P}(S_n|B)/\mathbb{P}(S_n|G)$. It follows from Figure 1a that $\beta_n(H) = (1-p)/p$ and $\beta_n(L) = p/(1-p)$. Next, using Bayes' rule, express γ_n in terms of l_{n-1} and β_n as $\gamma_n = 1/(1 + \beta_n l_{n-1})$. As a result, the condition on γ_n for a Y (N) cascade translates to $l_{n-1} < 1/\beta_n$ ($> 1/\beta_n$) for all S_n ; this simplifies to give Lemma 1.

If agent n cascades, then the observation 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, $l_{n+i} = l_{n-1}$ for all $i = 0, 1, 2, \dots$ and hence they remain in the cascade, which leads us to the following property, also exhibited by prior models, *e.g.* [2], [3], [4], [17].

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

$$l_i = \begin{cases} \left(\frac{1-b}{a}\right)l_{i-1}, & \text{if } O_i = Y, \\ \left(\frac{b}{1-a}\right)l_{i-1}, & \text{if } O_i = N, \end{cases} \quad (4)$$

$$a := \mathbb{P}(O_i = Y|V = G) \text{ and } b := \mathbb{P}(O_i = N|V = B). \quad (5)$$

Here, a and b denote the probabilities that an observation O_i follows V if agent i follows its private signal, given $V = G$ and B , respectively. It can be shown from Figures 1a and 1b that in the above case, *i.e.*, when A_i follows S_i ,

$$a = p + (1-p)\epsilon \text{ and } b = p(1-\epsilon). \quad (6)$$

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, $l_n = \left(\frac{1-p}{p}\right)^{h_n}$ where h_n is the difference between the number of Y 's weighted by η and the number of N 's,

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

$$\eta := \log\left(\frac{a}{1-b}\right)/\log\left(\frac{p}{1-p}\right). \quad (8)$$

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

Property 2. *Until a cascade occurs, each agent follows its private signal. Moreover, h_n defined in (7) is a sufficient statistic of the information contained in the past observations.*

Note that if $\epsilon = 0$ (no fake agents) then $a = b = p$ and $\eta = 1$, in which case h_n is the *unweighted* difference, $\eta n_Y - \eta n_N$, which is also the case for the unbiased noise model in [17]. Whereas, if $\epsilon > 0$ then $\eta < 1$. The expression for h_n in (7) shows that, due to the presence of fake agents, the dependence of an agent's decision on a Y in his observation history reduces by a factor of η , whereas the dependence on a N remains unaffected. This is to be expected because, unlike a N which

surely comes from an honest agent, a Y incurs the possibility that the agent could be fake. Further, this reduced dependence on Y is exacerbated with an increase in the possibility of fake agents, as η reduces with an increase in ϵ .

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

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

Whereas, once $h_n > 1$ (< -1), a $Y(N)$ cascade begins and h_n stops updating (Property 1). Note that $h_0 = 0$ since the first agent has no observation history. Now, given the true value $V \in \{G, B\}$, let the probability that a $Y(N)$ cascade begins be denoted by $\mathbb{P}_{Y\text{-cas}}^V$ ($\mathbb{P}_{N\text{-cas}}^V$). Here, $\mathbb{P}_{N\text{-cas}}^V = 1 - \mathbb{P}_{Y\text{-cas}}^V$ as it can be shown that the process $\{h_n\}$ exits the range $[-1, 1]$ w.p. 1. Further, let the n^{th} agent's *welfare* refer to its pay-off averaged (in expectation) over $V \in \{G, B\}$. We show later in Section V that this welfare as $n \rightarrow \infty$ relates to the cascade probabilities of the process $\{h_n\}$ as

$$\Pi := \lim_{n \rightarrow \infty} \mathbb{E}[\pi_n] = \frac{1}{4} [\mathbb{P}_{Y\text{-cas}}^G - \mathbb{P}_{Y\text{-cas}}^B]. \quad (10)$$

IV. MARKOVIAN ANALYSIS OF CASCADES

In this section, we analyse the process $\{h_n\}$, given V , to determine the probability of cascades. It follows from the previous section that conditioned on V , the process $\{h_n\}$ is a discrete-time Markov chain taking values in $[-1, 1]$ before getting absorbed into the *left absorption region* (< -1) causing a N cascade or the *right absorption region* (> 1) causing a Y cascade. More specifically, equation (9) shows that, given V , $\{h_n\}$ is a random walk (r.w.) that starts from state 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 (5). 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 its private signal S_n . We have from (5) that $p_f = a$ for $V = G$, whereas $p_f = 1 - b$ for $V = B$.

Note that in the special case where η , given by (8), satisfies $1/\eta = r$ for some $r = 1, 2, \dots$, the process $\{h_n\}$ is equivalent to a Markov chain with finite state-space $\mathcal{A} = \{-r-1, -r, \dots, -1, 0, 1, \dots, r, r+1\}$, and with $-r-1$ and $r+1$ being absorption states corresponding to N and Y cascades, respectively. More generally, it can be proved that $\{h_n\}$ has a finite state-space for any rational-valued $1/\eta$. In such cases, absorption probabilities can be obtained directly by solving a system of linear equations. In this paper, our main focus is on the more generic case of irrational values of $1/\eta$ resulting in $\{h_n\}$ taking countably infinite values in $[-1, 1]$,⁴ which does not readily allow for a direct solution to the cascade probabilities. This is unlike the unbiased noise model in [17] where the state-space of the Markov chain is always finite.

⁴For example, if η was chosen uniformly at random, then almost surely (w.p. 1) it would fall into this case.

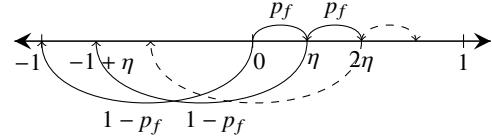


Fig. 2: Partial transition diagram of random walk $\{h_n\}$ given V .

A. Error thresholds

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

Lemma 2. *Let $\alpha = p/(1-p)$. For $r \in \mathbb{N}$, define the increasing sequence of thresholds $\{\epsilon_r\}_{r=1}^{\infty}$, with the r^{th} threshold given by*

$$\epsilon_r = \frac{\alpha - \alpha^{\frac{1}{r}}}{\alpha^{\frac{1}{r+1}} - 1}. \quad (11)$$

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

The proof follows by noting that, if a cascade has not begun immediately after an N is observed, then the rightmost position that the random walk can be in, is in state 0. Starting from state 0 therefore gives a lower bound on the number of consecutive Y 's required to begin a Y cascade. From here, $r+1$ consecutive Y 's would be needed to begin a Y cascade when $\frac{1}{r+1} < \eta \leq \frac{1}{r}$. This inequality implies that $\epsilon \in \mathcal{I}_r = [\epsilon_r, \epsilon_{r+1})$ where ϵ_r is the r^{th} threshold, defined in (11). Here, recall from (8) that η is a function of ϵ and p .

Remark 1. *For $\epsilon \in \mathcal{I}_r$, starting from state 0, $r+1$ consecutive Y 's start a Y cascade. Further, the integer r satisfies $r = \lfloor 1/\eta \rfloor$.*

Figure 3 shows the thresholds ϵ_r varying with p , for different values of r . For a fixed ϵ , we see that as the signal quality (p) improves, more consecutive Y 's are required for a Y cascade to begin. This is because, an increase in p increases the information contained in a Y , but not as much

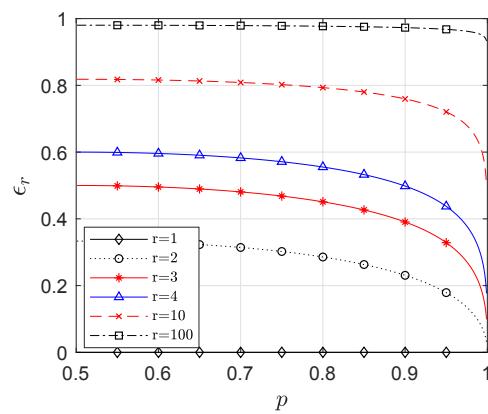


Fig. 3: Thresholds ϵ_r for the indicated values of r versus p .

as the corresponding increase in the information contained in a N . Further, note that as $\epsilon \rightarrow 1$, $r \rightarrow \infty$ which implies that infinitely many consecutive Y 's are required for a Y cascade to begin. Equivalently, the information contained in a Y observation becomes negligible. We further investigate learning in this asymptotic scenario in Section VI.

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

In this subsection, we will compute the probability of absorption of $\{h_n\}$ to a Y cascade (into the right absorption region) given V , namely $\mathbb{P}_{Y\text{-cas}}^V$. Any probability stated in the remainder of this section is assumed to be conditional under V through its dependence on p_f ; we thereby drop V from its superscript for notational convenience. Consider the iterative method depicted in Figure 4 that describes all possible sequences that can lead to a Y cascade. For this process, we assume $\epsilon \in \mathcal{I}_r \setminus \{\epsilon : 1/\eta \in \mathbb{Q}\}$ for some $r = 1, 2, \dots$. We do this to avoid the case of rational values of $1/\eta$ which could result in certain special sequences (having two consecutive N 's) with non-zero probability, that are not enumerated in Figure 4.⁵ For the case of rational-valued $1/\eta$, recall that the absorption probabilities can be obtained directly as solutions to a system of linear equations.

To begin the iterative method, we initialize Stage 1 with $r_1 = r+1$. Now, starting from state 0, consider the sequences shown in Stage 1 of Figure 4. The first sequence of r_1 consecutive Y 's, denoted by Y^{r_1} , clearly enters the right absorption region (Remark 1), 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 left absorption region. Now, each of these r_1 distinct sequences results in the same net right 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. From here, it would take either r or $r+1$ consecutive Y 's to enter the right absorption region. Let this value be denoted by r_2 . The sequences in Stage 2 can then be enumerated exactly as in the first stage, except that r_2 now replaces r_1 . Now, unless there are r_2 consecutive Y 's, the sequences of Stage 2 again end in the region $[0, \eta]$, and then the process continues to the next stage. Here, r_n denotes the number of consecutive Y 's required to enter the right absorption region in the n^{th} stage. In this manner, all sequences that lead to a Y cascade are enumerated.

Let P_n denote the probability of entering the right absorption region given that the sequence has not terminated in a Y cascade before the n^{th} stage. The following recursion then holds.

$$P_n = p_f^{r_n} [1 + r_n(1 - p_f)P_{n+1}], \text{ for } n = 1, 2, \dots \quad (12)$$

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

$$\mathbb{P}_{Y\text{-cas}}(\epsilon) = P_1, \text{ for } \epsilon \in \mathcal{I}_r \setminus \{\epsilon : 1/\eta \in \mathbb{Q}\}. \quad (13)$$

⁵These special sequences can be readily incorporated into the enumerations in Figure 4, but would disrupt the recursive pattern. For simplicity, we avoid such values of ϵ .

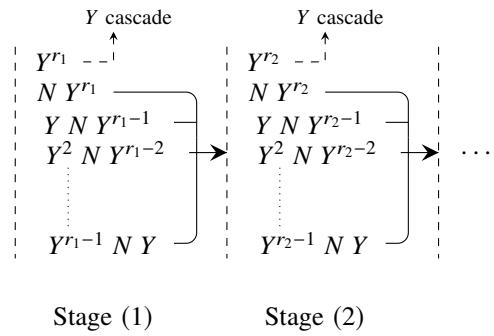


Fig. 4: An enumeration of all possible sequences that would lead to a Y cascade. The term Y^t represents a sequence of t consecutive Y 's. The sequence $\{r_n\}$ is defined as per (14).

Here, while $r_1 = r+1$, successive values of r_i for $i = 2, 3, \dots$ can be obtained from r_1 using the updates:

$$r_n = \begin{cases} r, & \text{if } \sum_{i=1}^{n-1} (r_i\eta - 1) + r\eta > 1, \\ r+1, & \text{o.w.} \end{cases} \quad (14)$$

Since (12) is an infinite recursion, to compute $\mathbb{P}_{Y\text{-cas}}$ in practice, we truncate the process to a finite number of iterations M . To this end, we first assume that $P_{M+1} = 1$. Next, we use (12) 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 $\mathbb{P}_{Y\text{-cas}}^M$. The following theorem shows that $\mathbb{P}_{Y\text{-cas}}^M$ is in fact a tight upper bound to $\mathbb{P}_{Y\text{-cas}}$ as $M \rightarrow \infty$. Moreover, the difference $\mathbb{P}_{Y\text{-cas}}^M - \mathbb{P}_{Y\text{-cas}}$ decays to zero at least as fast as 0.5^M , in the number of iterations M . Refer to Appendix A in [25] for a detailed proof.

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

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

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

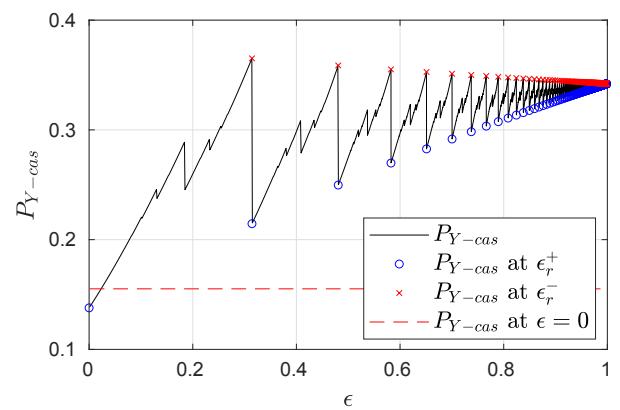


Fig. 5: Probability of Y cascade as a function of ϵ for $V = B$ and $p = 0.7$.

Figure 5 shows a plot of $\mathbb{P}_{Y\text{-cas}}$ with respect to ϵ , for the case $V = B$ with $p = 0.7$. The plot uses $M = 10$ which gives an error of less than 10^{-3} . It can be seen that in the

r^{th} ϵ -interval \mathcal{I}_r , $\mathbb{P}_{Y\text{-cas}}$ increases with ϵ , but with infinitely many discontinuities (where $\mathbb{P}_{Y\text{-cas}}(\epsilon)$ decreases). This is in contrast with the unbiased noise model in [17], where it increases continuously over a similar interval. This distinction is a consequence of the state-space of $\{h_n\}$ being infinite, unlike in [17]. Despite the discontinuities, $\mathbb{P}_{Y\text{-cas}}$ achieves the minimum (marked by \circ) and maximum (marked by \times) values at the edge points of \mathcal{I}_r , i.e., ϵ_r^+ and ϵ_{r+1}^- , respectively. Further, note that the relatively larger drops in $\mathbb{P}_{Y\text{-cas}}$ observed in Figure 5 occur exactly at the threshold points $\{\epsilon_r\}_{r=1}^\infty$. Here, counter to expectation, a slight increase in ϵ beyond ϵ_r causes a significant decrease in the probability of a Y cascade. The same characteristic is exhibited when $V = G$, which is shown in Figure 11, in Appendix B of [25]. In Section IV-C, we show that closed-form expressions for $\mathbb{P}_{Y\text{-cas}}$ as $\epsilon \rightarrow \epsilon_r^+$ and $\epsilon \rightarrow \epsilon_r^-$ can be obtained (refer to (15) and (16)); thereby quantifying its drop at each threshold ϵ_r .

In Figure 12 of Appendix C in [25], we outline an iterative method similar to that in Figure 4, to compute the probability of absorption of $\{h_n\}$ to a N cascade, namely $\mathbb{P}_{N\text{-cas}}$ with arbitrary precision. Having already computed $\mathbb{P}_{Y\text{-cas}}$, this alternate process may seem redundant due to the fact that $\mathbb{P}_{N\text{-cas}} = 1 - \mathbb{P}_{Y\text{-cas}}$. However, later in Subsection IV-E, we will be using both the iterative methods depicted by Figures 4 and 12 to compute exact probabilities of a Y and N cascade respectively, that may begin at or before a finite agent index n . We denote the probability that a Y (N) cascade occurs by the n^{th} agent index by v_n^{*V} (u_n^{*V}). Note that agent $n+1$ always has a non-zero probability $1 - v_n^{*V} - u_n^{*V}$ of following its private signal S_{n+1} , hence $u_n^{*V} \neq 1 - v_n^{*V}$. As a result, u_n^{*V} cannot be inferred by knowing v_n^{*V} , and vice-versa. So, both the iterative methods are required to obtain cascade probabilities for any finite n .

C. Quantifying abrupt changes in $\mathbb{P}_{Y\text{-cas}}$ at error thresholds

We now quantify $\mathbb{P}_{Y\text{-cas}}$ as ϵ tends to each threshold point ϵ_r . For this, we state the following lemma, with the proof provided in Appendix D of [25].

Lemma 3. For all $i \geq 2$, r_i in (14) satisfies: $r_i \rightarrow r$ as $\epsilon \rightarrow \epsilon_r$.

It follows from Lemma 3 that as $\epsilon \rightarrow \epsilon_r$, the recursion in (12) results in the same infinite computation to obtain P_i as for P_{i+1} , for all $i \geq 2$. Thus, all P_i for $i \geq 2$ have the same value which satisfies: $P_i = p_f^r [1 + r(1 - p_f)P_i]$. Solving this equation for $i = 2$ gives the value for P_2 which when used in equation (12) for $n = 1$ yields P_1 , i.e., $\mathbb{P}_{Y\text{-cas}}$. However, note that while solving equation (12), $r_1 = r+1$ for $\epsilon = \epsilon_r^+$ whereas $r_1 = r$ for $\epsilon = \epsilon_r^-$. This corresponds to the following two different values of $\mathbb{P}_{Y\text{-cas}}$ as $\epsilon \rightarrow \epsilon_r$:

$$\mathbb{P}_{Y\text{-cas}}(\epsilon_r^+) = p_f^{r+1} \frac{1 + (1 - p_f)p_f^r}{1 - r(1 - p_f)p_f^r}, \quad (15)$$

$$\mathbb{P}_{Y\text{-cas}}(\epsilon_r^-) = p_f^r \frac{1}{1 - r(1 - p_f)p_f^r}. \quad (16)$$

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

$$\delta_r = (1 - p_f)(1 - p_f^{r+1}). \quad (17)$$

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

The above property implies that, increasing ϵ just above ϵ_r is in fact counter-productive for the fake agents. The intuition underlying Property 3 is as follows. Recall from Lemma 2 that an increase in ϵ from ϵ_r^- to ϵ_r^+ increases the least number of consecutive Y 's required to begin a Y cascade by one. This implies that any observation sequence that ends in a Y cascade for $\epsilon = \epsilon_r^+$ guarantees a Y cascade for $\epsilon = \epsilon_r^-$. However, the opposite is not true. Therefore, increasing the noise value from ϵ_r^- to ϵ_r^+ results in a drop in the Y cascade probability.

Lastly, it can be verified from (17) that as $r \rightarrow \infty$, $\delta_r \rightarrow 0$. This is depicted in Figure 5 where the sequences $\{\mathbb{P}_{Y\text{-cas}}(\epsilon_r)\}$ and $\{\mathbb{P}_{Y\text{-cas}}(\epsilon_r^+)\}$, marked by \times and \circ respectively, converge to a limiting value as $r \rightarrow \infty$. Note that $r \rightarrow \infty$ implies the limiting scenario of $\epsilon \rightarrow 1$. We investigate this scenario further in Section VI.

D. Effect of fake agents on $\mathbb{P}_{Y\text{-cas}}$ at low values of ϵ

Next, we consider the cascade behaviour for low values of ϵ . In the absence of fake agents ($\epsilon = 0$) as in [2], $\eta = 1$. It then follows from Figure 2 that $\{h_n\}$ has a finite state-space $\{-2, -1, 0, 1, 2\}$, and a Y (N) cascade starts when $h_n = 2$ (-2). Here, solving for $\mathbb{P}_{Y\text{-cas}}$ gives

$$\mathbb{P}_{Y\text{-cas}}(0) = \begin{cases} p^2 / [p^2 + (1 - p)^2], & \text{for } V = G, \\ (1 - p)^2 / [p^2 + (1 - p)^2], & \text{for } V = B. \end{cases} \quad (18)$$

In Figure 5 and Figure 11 (Appendix B of [25]), this is indicated as a baseline to compare it with $\mathbb{P}_{Y\text{-cas}}$ for any $\epsilon \in (0, 1)$, when $V = B$ and $V = G$, respectively. Observe that for low values of ϵ , $\mathbb{P}_{Y\text{-cas}}$ is smaller than when fake agents are absent. The next theorem formalizes this property for any p and any V .

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

$$\mathbb{P}_{Y\text{-cas}}(\epsilon) < \mathbb{P}_{Y\text{-cas}}(0), \quad \forall \epsilon \in (0, \underline{\epsilon}). \quad (19)$$

From the point-of-view of the fake agents, the above theorem implies the following property.

Property 4. If fake agents occur with a probability of less than $\underline{\epsilon}$, then the effect that their presence has on the honest buyers is opposite to what they had intended. That is, they reduce the chances of their preferred (Y) cascade instead of increasing it.

Likewise, Theorem 2 implies that if $V = B$, then the honest buyers benefit from the presence of fake agents when $\epsilon < \underline{\epsilon}$. Otherwise, if $V = G$, then they are harmed by such a presence of fake agents.

Proof: Note that as $\epsilon \rightarrow 0$, the limiting value of $\mathbb{P}_{Y\text{-cas}}$ can be obtained from (15) with $r = 1$ and $p_f \rightarrow 1 - p$ for $V = B$, whereas $p_f \rightarrow p$ for $V = G$. This gives

$$\lim_{\epsilon \rightarrow 0} \mathbb{P}_{Y\text{-cas}}(\epsilon) = \begin{cases} (1-p)^2 \frac{1+p(1-p)}{1-p(1-p)}, & \text{for } V = B, \\ p^2 \frac{1+p(1-p)}{1-p(1-p)}, & \text{for } V = G. \end{cases} \quad (20)$$

In Figure 5, this limiting value is marked by \circ at $\epsilon = 0$. By comparing this expression with the one in (18) for the corresponding values of V , we have

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

Thus, there exists some $\underline{\epsilon} > 0$ such that (19) holds true. \blacksquare

From the above proof, it follows that Theorem 2 is a consequence of the discontinuity in $\mathbb{P}_{Y\text{-cas}}$ at $\epsilon = 0$, for any given V . This contrasts with the unbiased noise model in [17], where there is no discontinuity in the cascade probabilities at $\epsilon = 0$ (no noise). Once again, this distinction results from the state-space of $\{h_n\}$ being infinite, unlike in [17].

E. Probability of cascades for finite agent arrivals

In this subsection, we use the iterative methods depicted in Figure 4 and Figure 12 (in Appendix C of [25]) to compute exact probabilities for correct and wrong cascades to begin within the arrival of a finite number of agents. For the random walk $\{h_n\}$ in Figure 2, given the true value $V \in \{G, B\}$, let v_n^{*V} (u_n^{*V}) be the probability of being absorbed by the right (left) absorption region by the n^{th} time-step. As the time-steps of $\{h_n\}$ correspond to the agents' indices, v_n^{*V} and u_n^{*V} are respectively the probabilities that a correct and wrong cascade occurs by the n^{th} arrival, under $V = G$. Vice-versa holds when $V = B$. Now, let v_n^V (u_n^V) be the probability that $\{h_n\}$ enters the right (left) absorption region exactly at the n^{th} time-step.

$$v_n^{*V} = \sum_{i=1}^n v_i^V \quad \text{and} \quad u_n^{*V} = \sum_{i=1}^n u_i^V. \quad (21)$$

Figure 5 shows that all sequences that enter the right absorption region (*i.e.* end in a Y cascade) in Stage (j) terminate with a Y^{r_j} and have the same length: $l_j \triangleq r_j + \sum_{i=1}^{j-1} (r_i + 1)$, for $j = 1, 2, \dots$. This yields the values for $\{v_n\}$ as follows:

$$v_n = \begin{cases} p_f^{r_j} \prod_{i=1}^{j-1} r_i (1-p_f) p_f^{r_i}, & \text{if } n = l_j, \\ 0, & \text{o.w.} \end{cases} \quad (22)$$

In a similar manner, by observing Figure 12 in [25], we see that among all sequences that end in a N cascade in Stage (j), the t sequences that terminate in any allowable permutation of $NY^{t-1}N$ have the same length: $t+1 + \sum_{i=1}^{j-1} (r_i + 1) = l_{j-1} + t + 2$, for $t = 1, 2, \dots, r_j$ and $j = 1, 2, \dots$. This yields the values for $\{u_n\}$ as follows:

$$u_n = \begin{cases} t p_f^{t-1} (1-p_f)^2 \prod_{i=1}^{j-1} r_i (1-p_f) p_f^{r_i}, & \text{if } n = l_{j-1} + t + 2, \\ 0, & \text{o.w.} \end{cases} \quad (23)$$

Now, as $n \rightarrow \infty$, the asymptotic value of the quantity v_n^{*V} (u_n^{*V}) refers to the probability that a Y (N) cascade occurs eventually. Thus,

$$\lim_{n \rightarrow \infty} v_n^{*V} = \mathbb{P}_{Y\text{-cas}}^V \quad \text{and} \quad \lim_{n \rightarrow \infty} u_n^{*V} = \mathbb{P}_{N\text{-cas}}^V \stackrel{(a)}{=} 1 - \mathbb{P}_{Y\text{-cas}}^V, \quad (24)$$

where equality (a) holds since it can be shown that $\{h_n\}$ is absorbed into either of the cascades w.p. 1.

V. WELFARE FOR ORDINARY AGENTS

In this section, we analyze the expected pay-off or welfare of agent n , π_n if it is ordinary (rational), *i.e.*, if it takes a pay-off optimal action. Recall from Section II that $\pi_n = 0$ if $A_n = N$ whereas $\pi_n = 1/2$ or $-1/2$ if $A_n = Y$ depending on whether $V = G$ or B , respectively. Now, Figure 2 shows that it takes at least two steps to begin a cascade. This implies that the first two agents always follow their private signals, and hence have the same welfare given by

$$\mathbb{E}[\pi_n] = \mathbb{P}(A_n = Y|V = G) \frac{1}{4} - \mathbb{P}(A_n = Y|V = B) \frac{1}{4}, \quad (25)$$

$$= (2p - 1)/4 \triangleq F, \quad \text{for } n \in \{1, 2\}. \quad (26)$$

In fact, F defined in (26) refers to the welfare for any agent n , if A_n always follows S_n disregarding the optimal decision rule in (2), *i.e.*, $\mathbb{E}[\pi_n|A_n \text{ always follows } S_n] = F$, for all n . However, for agents $n > 2$, the unconditional welfare must also account for the possibilities that the history \mathcal{H}_{n-1} could cause agent n to cascade to a Y or a N , which can be expressed as follows.

$$\mathbb{E}[\pi_n] = F + \frac{1}{4} [(1-p)v_{n-1}^{*G} - p v_{n-1}^{*B} + (1-p)u_{n-1}^{*B} - p u_{n-1}^{*G}]. \quad (27)$$

The above equation explicitly relates agent n 's welfare to the probabilities of Y and N cascades resulting from its history \mathcal{H}_{n-1} under both $V = G$ and $V = B$.

Now, assuming that agent n can choose to observe only a certain subset of the available observations $I_n = \{O_1, \dots, O_{n-1}, S_n\}$, the next theorem shows that its welfare cannot get worse if more elements are added to this subset. This implies the non-redundancy of the observations in I_n as agent n can never achieve a welfare higher than in (27) by ignoring any of the observations in I_n .

Theorem 3. *Let \mathcal{J}_n be the collection of all subsets of the observations $\{O_1, \dots, O_{n-1}, S_n\}$ that are available to agent n . Let $\mathbb{E}[\pi_n|J]$ refer to the optimal welfare achieved by agent n by only observing the set $J \in \mathcal{J}_n$. Then, for any two sets $J, K \in \mathcal{J}_n$ such that $K \subset J$, we have $\mathbb{E}[\pi_n|J] \geq \mathbb{E}[\pi_n|K]$.*

Theorem 3 can be proved by applying the celebrated Blackwell's Theorem [11] which implies that it is sufficient to show that the signals from observing the smaller set K are obtained as a stochastic mapping (garbling) of the signals from the larger set J . Let \bar{J} and \bar{K} be the j and k -length random vectors ($j > k$) corresponding to the observations sets J and K , respectively, such that the two vectors share the first k elements. Then, the desired mapping is given by $\bar{K} = G\bar{J}$, where G is a $k \times j$ diagonal matrix. Then, Blackwell's result for the corresponding optimal welfares states that $\mathbb{E}[\pi_n|J] \geq \mathbb{E}[\pi_n|K]$.

As a corollary of Theorem 3, the next property shows that the welfare under complete observation, *i.e.*, $\mathbb{E}[\pi_n]$ is monotonic in the agents' indices.

Property 5. *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 .*

To see this property, consider two consecutive agents, $n-1$ and n . Under the informational equivalence of their private signals: S_{n-1} and S_n , we have $I_{n-1}, I_n \in \mathcal{J}_n$ and $I_{n-1} \subset I_n$. By applying Theorem 3, this implies that $\mathbb{E}[\pi_n] \geq \mathbb{E}[\pi_{n-1}]$; and $\mathbb{E}[\pi_n] \geq F$ for all n follows from (26).

Next, by taking the limit $n \rightarrow \infty$ of (27) and using (24), the asymptotic welfare $\lim_{n \rightarrow \infty} \mathbb{E}[\pi_n(\epsilon)]$ denoted by $\Pi(\epsilon)$ is given by

$$\Pi(\epsilon) = \frac{1}{4} [\mathbb{P}_{Y\text{-cas}}^G(\epsilon) - \mathbb{P}_{Y\text{-cas}}^B(\epsilon)] \stackrel{(a)}{=} \frac{1}{4} (1 - 2\mathbb{P}_{\text{wrong-cas}}(\epsilon)). \quad (28)$$

In Step (a), $\mathbb{P}_{\text{wrong-cas}} := [\mathbb{P}_{N\text{-cas}}^G + \mathbb{P}_{Y\text{-cas}}^B]/2$ refers to the unconditional probability of a wrong cascade, i.e., the probability that a N cascade occurs and $V = G$ or a Y cascade occurs and $V = B$. Equation (28) implies that an *improved learning*, i.e., a lower probability of wrong cascade results in a higher asymptotic welfare. The probability $\mathbb{P}_{Y\text{-cas}}^V(\epsilon)$ for $V \in \{G, B\}$ can be computed using the recursive method described in Section IV, which is outlined by equations (12), (13) and (14). Then, substituting these obtained values in (28) yields the value for $\Pi(\epsilon)$. Figure 6 shows a plot of $\Pi(\epsilon)$ with respect to $\epsilon \in (0, 1)$, for $p = 0.7$ and compares it with the constant level of $\Pi(0)$ which refers to the asymptotic welfare in the absence of fake agents. Substituting (18) in (10) gives

$$\Pi(0) = (1/4)(2p - 1)/[p^2 + (1 - p)^2]. \quad (29)$$

It can be observed in Figure 6 that for $p = 0.7$, $\Pi(\epsilon) < \Pi(0)$ for all $\epsilon \in (0, 1)$. Further, note that the relatively larger jumps in $\Pi(\epsilon)$ observed in Figure 6 (marked by \times and \circ) occur exactly at the threshold points $\{\epsilon_r\}_{r=2}^\infty$. Here, counter to expectation, a slight increase in ϵ beyond ϵ_r causes an abrupt and significant increase in the asymptotic welfare. This abrupt increase does not simply follow from Property 3 because a drop in the Y cascade probability improves learning when $V = B$, whereas it worsens learning when $V = G$. Thus, when averaged across V , Property 3 is not sufficient to imply that learning improves in turn causing the asymptotic welfare to increase at each ϵ_r -threshold. This is unlike the unbiased noise model in [17], where learning at noise thresholds improves for both V 's, thereby implying an increase in asymptotic welfare. We propose the following theorem, which shows that such an abrupt increase in welfare at ϵ_r occurs not only as $n \rightarrow \infty$ but also occurs for every agent n . Refer to Appendix E in [25] for a detailed proof.

Theorem 4. *Given a fixed private signal quality p , for every agent $n \in \{1, 2, \dots\}$, at any r^{th} ϵ -threshold ϵ_r , $r \geq 2$,*

$$\mathbb{E}[\pi_n(\epsilon_r^+)] - \mathbb{E}[\pi_n(\epsilon_r^-)] > 0. \quad (30)$$

It follows from Theorem 4 that for any p , there occurs an abrupt increase of Δ_r in the asymptotic welfare at each of the threshold points $\{\epsilon_r\}_{r=2}^\infty$, where

$$\Delta_r := \Pi(\epsilon_r^+) - \Pi(\epsilon_r^-) > 0. \quad (31)$$

Through the relation in (28), this improvement of Δ_r in the asymptotic welfare corresponds to an abrupt reduction of $(4\Delta_r - 1)/2$ in the wrong cascade probability, $\mathbb{P}_{\text{wrong-cas}}$. The expression for Δ_r can be obtained by first using equations (15) and (16) to compute the values: $\mathbb{P}_{Y\text{-cas}}^V(\epsilon)$ for $V \in \{G, B\}$ and

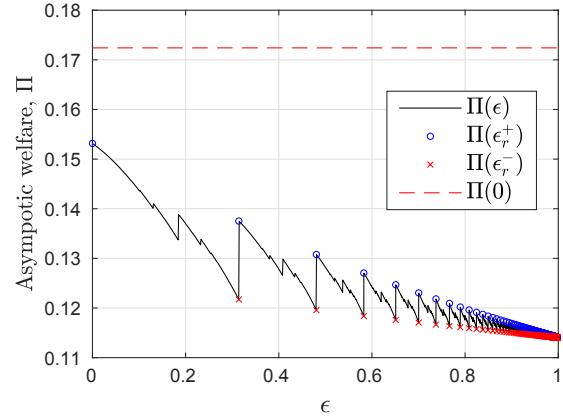


Fig. 6: Asymptotic welfare as a function of ϵ for $p = 0.7$.

$\epsilon = \epsilon_r^+$ and ϵ_r^- . Next, substituting these values in (28) to obtain $\Pi(\epsilon_r^+)$ and $\Pi(\epsilon_r^-)$ and then using (17) yields

$$\Delta_r = \frac{1}{4} [\delta_r^B \mathbb{P}_{Y\text{-cas}}^B(\epsilon_r^-) - \delta_r^G \mathbb{P}_{Y\text{-cas}}^G(\epsilon_r^-)]. \quad (32)$$

Property 6. *If the possibility of fake agents in the history equals the r^{th} ϵ -threshold, $r = 2, 3, \dots$, then a further marginal increase in fake agents improves the welfare at every agent index $n \in \{1, 2, \dots\}$. Moreover, the asymptotic welfare of the agents improves by Δ_r .*

Therefore, increasing ϵ over the r^{th} ϵ -threshold is not only counter-productive for the fake agents (due to Property 3), but it also leads to a higher social welfare for every ordinary agent. The intuition underlying Property 6 is that when ϵ increases from ϵ_r^- to ϵ_r^+ , the drop in the Y cascade probability for $V = B$ (better learning) is more pronounced than the corresponding drop for $V = G$ (worse learning). This is because this drop is a decreasing function of p_f , and that p_f for $V = G$ is greater than p_f for $V = B$ (i.e., $a > 1 - b$). Therefore, this increase in ϵ implies a net decrease in the wrong cascade probability when averaged over V , which as per (28) leads to a higher welfare for rational agents.

VI. LEARNING IN THE LIMIT

In this section, we evaluate the asymptotic welfare of agents under two limiting regimes of ϵ , namely, $\epsilon \rightarrow 0$ and $\epsilon \rightarrow 1$. In the first case, taking the limit $\epsilon \rightarrow 0$ in (10) and then substituting the expressions obtained in (20) yields

$$\lim_{\epsilon \rightarrow 0} \Pi = \frac{1}{4}(2p - 1) \frac{1 + p(1 - p)}{1 - p(1 - p)}. \quad (33)$$

Comparing (33) with the asymptotic welfare at $\epsilon = 0$ given in (29), it can be shown that $\lim_{\epsilon \rightarrow 0} \Pi(\epsilon) < \Pi(0)$ for any p . This implies that even an infinitesimal presence of fake agents causes an abrupt deterioration in welfare. Figure 6 shows this drop in welfare at $\epsilon = 0$, where the value $\Pi(0)$ drops to the value $\lim_{\epsilon \rightarrow 0} \Pi(\epsilon)$, which is marked by \circ . Next, we define the fractional reduction (f.r.) in welfare relative to $\Pi(0)$, as $\lim_{\epsilon \rightarrow 0} [\Pi(0) - \Pi(\epsilon)]/\Pi(0)$. We then plot it against p in Figure 7a in order to understand the effects of varying the signal quality p . It can be shown analytically from the expressions in (29) and (33) that the f.r. in welfare is monotonically decreasing

with p . The greatest f.r. thus occurs as $p \rightarrow 0$ and is found to equal $1/6$, whereas the f.r. is 0 in the limit $p \rightarrow 1$.

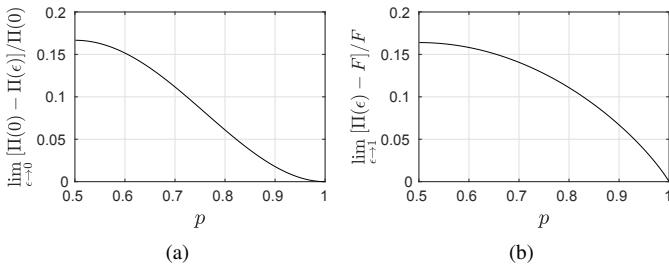


Fig. 7: (a) Fractional reduction in welfare relative to $\Pi(0)$ for varying p , as $\epsilon \rightarrow 0$. (b) Fractional increase in welfare relative to F for varying p , as $\epsilon \rightarrow 1$.

In the case of $\epsilon \rightarrow 1$, the information contained in a Y observation becomes negligible. As a result, an agent would need to observe infinitely many consecutive Y 's in his history for him to be convinced of starting a Y cascade. Hence, one would expect that if $V = G$, learning would never occur, whereas if $V = B$, then learning would always occur. However, recall that as $\epsilon \rightarrow 1$, the occurrence of Y 's becomes increasingly frequent, *i.e.* $p_f \rightarrow 1$; for both $V = B$ and G . This motivates studying $\mathbb{P}_{Y\text{-cas}}$ in this limiting scenario. First, recall that in the process of enumerating all sequences leading to a Y cascade, for $\epsilon \in \mathcal{I}_r$, in each stage $i \geq 2$, r_i is either r or $r+1$. However, $\epsilon \rightarrow 1$ implies $r \rightarrow \infty$, in which case $r \approx r+1$. As a result, the expressions obtained in (15) and (16) yield the same limiting value as $r \rightarrow \infty$, which also equals $\mathbb{P}_{Y\text{-cas}}$ as $\epsilon \rightarrow 1$. In particular,

$$\lim_{\epsilon \rightarrow 1} \mathbb{P}_{Y\text{-cas}}(\epsilon) = \lim_{r \rightarrow \infty} \mathbb{P}_{Y\text{-cas}}(\epsilon_r^+) \stackrel{(b)}{=} \lim_{r \rightarrow \infty} \mathbb{P}_{Y\text{-cas}}(\epsilon_r^-), \quad (34)$$

where Step (b) can also be proved by recalling that $\delta_r \rightarrow 0$ as $r \rightarrow \infty$. By using (16) in (34), the limiting probability of a Y cascade, in terms of $\alpha = p/(1-p)$ can be obtained as:

$$\lim_{\epsilon \rightarrow 1} \mathbb{P}_{Y\text{-cas}}(\epsilon) = \lim_{r \rightarrow \infty} p_r^r \frac{1}{1 - r(1-p_f)p_r^r} = \frac{1}{e^t - t}; \quad (35)$$

where $t = \frac{1}{\alpha-1} \log \alpha$ for $V = G$, and $t = \frac{\alpha}{\alpha-1} \log \alpha$ for $V = B$. A detailed proof of (35) is provided in Appendix F of [25].

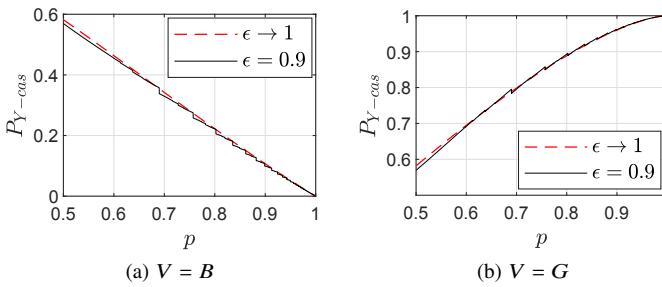


Fig. 8: Probability of Y cascade versus private signal quality for the indicated values of ϵ under (a) $V = B$ and (b) $V = G$.

Figure 8 illustrates (35) and the corresponding probability when $\epsilon = 0.9$ as a function of the signal quality. For both $V = B$ and $V = G$, a better signal quality leads to improved learning even when fake agents have overwhelmed

the ordinary agents. Also note that for $V = B$, for a weak signal quality, the incorrect cascade is more likely than the correct one (while for $V = G$, this is never true).

The asymptotic welfare in this scenario can be obtained by taking the limit $\epsilon \rightarrow 1$ in (10), then substituting (35). In Figure 6, observe that the sequences $\{\Pi(\epsilon_r^-)\}$ and $\{\Pi(\epsilon_r^+)\}$, marked by \times and \circ respectively, converge to this limiting value as $r \rightarrow \infty$. Further, it can be proved that for any p , $\lim_{\epsilon \rightarrow 1} \Pi(\epsilon) > F$, where we recall that F (defined in (26)) is the welfare when an agent acts based on its private signal alone. Therefore, even when fake agents have corrupted almost all the actions, it is still better for any ordinary agent to observe its past in addition to its private signal. Figure 7b plots the fractional improvement in asymptotic welfare relative to F against varying values of p and shows that it decays to 0 with increasing p .

VII. EFFECTS OF A PLATFORM CO-ORDINATOR

In this section, we introduce an additional entity called the *Platform Co-ordinator*, which at each time i , randomly modifies (manipulates or filters) the observation O_i before presenting it to future agents. We assume that the Co-ordinator has the same information about private signal quality p and fraction of fake agents ϵ as the rational buyers, while the item's underlying true value (V) is unknown to both and assumed to be equiprobable. Moreover, the type of modification and the corresponding parameter(s) used by the Co-ordinator are common knowledge. By modifying the observations, the Co-ordinator who acts as a Bayesian *persuader* aims to improve the asymptotic welfare. A related work is [27], where a similar entity instead designs the information structure for the agents' private signals so as to maximize its utility which depends on the eventual learning outcome. We investigate the effects on agents' asymptotic welfare under three scenarios, namely wherein the Co-ordinator (a) modifies a N to a Y w.p. κ or (b) discards a Y w.p. e or (c) modifies a Y to a N w.p. β . Otherwise, in all scenarios, observations that are neither discarded nor modified are retained. Let ${}^{(i)}\Pi$ refer to the asymptotic welfare, where $i \in \{a, b, c\}$ indicates the particular scenario being considered. The superscript (i) is dropped when referring to the asymptotic welfare without the Co-ordinator.

Scenario (a): Manipulating the N 's

Here, the Co-ordinator modifies only the N 's in the observations, *i.e.*, at any time i , if $O_i = N$ then the modified observation $O'_i = Y$ w.p. $\kappa \in [0, 1]$ and $O'_i = N$ w.p. $1 - \kappa$. Whereas, if $O_i = Y$ then $O'_i = Y$ w.p. 1. Here, $\mathcal{H}'_{n-1} := \{O'_1, \dots, O'_{n-1}\}$ denotes the modified history observed by agent n . This has the effect that agents now perceive their history as having an effectively increased probability, $\epsilon_{\text{eff}} := \epsilon + (1 - \epsilon)\kappa$ of past agents being fake. As a result, the asymptotic welfare in this scenario denoted by ${}^{(a)}\Pi(\epsilon, \kappa)$ is given by ${}^{(a)}\Pi(\epsilon, \kappa) = \Pi(\epsilon_{\text{eff}})$, where $\Pi(\cdot)$ is defined in (28). Note that by varying κ , ϵ_{eff} can attain any desired value in $[\epsilon, 1]$. Specifically, to get the best improvement in asymptotic welfare, the Co-ordinator can set κ to the optimal value:

$$\kappa^* = \arg \max_{\kappa \in [0, 1]} \Pi(\epsilon + (1 - \epsilon)\kappa), \quad (36)$$

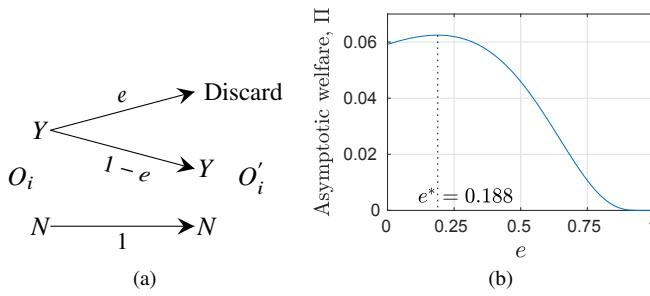


Fig. 9: (a) The channel through which observations are filtered by the Co-ordinator at each time i . (b) Asymptotic welfare $\Pi(\epsilon, e)$ as a function of the filtration parameter e , for $p = 0.6$ and $\epsilon = 0.9$. Maximal value obtained at $e = e^*$.

which yields the optimal asymptotic welfare,

$$^{(a)}\Pi(\epsilon, \kappa^*) = \Pi(\epsilon + (1 - \epsilon)\kappa^*). \quad (37)$$

Scenario (b): Filtering out the Y 's

In this scenario, the Co-ordinator randomly filters out the Y 's from the observations instead of manipulating the observations as was done in Scenario (a). Specifically, at every time n , it discards a Y observation w.p. $e \in [0, 1]$ and retains it otherwise. Whereas, a N observation is always retained. This bias is due to the fact that, while a N observation always represents the action of a rational agent, a Y observation could also represent the action of a fake agent. So, filtering out some of the Y observations while retaining all the N observations seems reasonable. The channel in Figure 9a depicts the filtering of observation O_i at each time i , where $O'_i \in \{Y, N, \text{Discard}\}$ denotes the channel output. Here, $O'_i = \text{Discard}$ implies that a Y at time i has been discarded by the Co-ordinator.

Note that the arrival of an agent into the platform is recorded for (and visible to) subsequent agents only if its corresponding observation is retained by the Co-ordinator. If an agent's observation is discarded, future agents are unaware of its arrival into the buying platform. Therefore, we define a new index set $\{i_1, i_2, \dots\} \subseteq \{1, 2, \dots\}$ which indexes the arrivals of only those agents whose observations are retained. Here, i_n is the index in the unfiltered arrival sequence of the agent whose observation is the n^{th} undiscarded observation.

Now, for the i_n^{th} agent which observes a history $\mathcal{H}_{i_{n-1}} := \{O_{i_1}, \dots, O_{i_{n-1}}\}$, the Bayes' optimal decision rule is still given by (2), except that the posterior probability for the item being good, $\gamma_{i_n} := \mathbb{P}(G|S_{i_n}, \mathcal{H}_{i_{n-1}})$. Further, with the likelihood ratio for its history $\mathcal{H}_{i_{n-1}}$ being redefined as $l_{i_{n-1}} \triangleq \mathbb{P}(\mathcal{H}_{i_{n-1}}|B)/\mathbb{P}(\mathcal{H}_{i_{n-1}}|G)$, it can be shown that both Lemma 1 and Property 1 remain true. Now, assuming the i_n^{th} agent does not cascade, Property 1 and Lemma 1 imply that for every observation O_{i_k} , $k \leq n$, A_{i_k} follows S_{i_k} . Then, for every such observation O_{i_k} , the probability that it follows the true value V is implicitly conditioned on the fact that it has not been discarded, and hence is given by

$$a' = \mathbb{P}(O_{i_k} = Y|G, O_{i_k} \neq \text{Discard}) = \frac{a(1 - e)}{a(1 - e) + (1 - a)}, \quad (38)$$

$$b' = \mathbb{P}(O_{i_k} = N|B, O_{i_k} \neq \text{Discard}) = \frac{b}{b + (1 - b)(1 - e)}. \quad (39)$$

Further, for every such O_{i_k} , as S_{i_k} is conditionally independent of the history $\mathcal{H}_{i_{k-1}}$ given V , the likelihood ratio updates are $l_{i_k} = \left(\frac{1-b}{a}\right)l_{i_{k-1}}$ if $O_{i_k} = Y$, and $l_{i_k} = \left(\frac{b}{1-a}\right)l_{i_{k-1}}$ if $O_{i_k} = N$. These updates are identical to the likelihood updates in (4) for the original model without the Co-ordinator. This is because, such an observation O_{i_k} , which could be an undiscarded Y or an N , still possesses the same information about V as in the original model. Now, as a result of the updates, l_{i_n} can be expressed as $l_{i_n} = \left(\frac{1-p}{p}\right)^{h_{i_n}}$ where h_{i_n} is a sufficient statistic of the information contained in history \mathcal{H}_{i_n} , given that agent i_n is not in a cascade. Similar to eq. (7) with η as per (8),

$$h_{i_n} := \eta n_Y - n_N, \quad (40)$$

which is the difference between the number of Y 's (n_Y) weighted by η and the number of N 's (n_N) present in \mathcal{H}_{i_n} .

Now, Lemma 1 implies that until a cascade occurs, $-1 \leq h_{i_n} \leq 1$ for all such times n , and (40) implies that the update rule for h_{i_n} is given by

$$h_{i_n} = \begin{cases} h_{i_{n-1}} + \eta & \text{if } O_{i_n} = Y, \\ h_{i_{n-1}} - 1 & \text{if } O_{i_n} = N. \end{cases} \quad (41)$$

Once $h_{i_n} > 1$ (< -1), a Y (N) cascade begins and h_{i_n} stops updating (Property 1). Observe that the updates of $\{h_{i_n}\}$ stated in (41) are identical to the updates of $\{h_n\}$ in (7). However, $\{h_{i_n}\}$ has two notable differences. Firstly, $\{h_{i_n}\}$ is indexed by the new set $\{i_1, i_2, \dots\}$ and secondly, conditioned on V , $\{h_{i_n}\}$ moves to the right by η w.p. $\mathbb{P}(O_{i_n} = Y|V)$ or to the left by 1 w.p. $\mathbb{P}(O_{i_n} = N|V)$ until a cascade occurs, where these probabilities are defined in terms of a' and b' in (38) and (39). This is unlike in the case of $\{h_n\}$ where a and b in (6) define the transition probabilities.

Next, let the Y cascade probability for the process $\{h_{i_n}\}$ be denoted by $^{(b)}\mathbb{P}_{Y\text{-cas}}^V(\epsilon, e)$ to highlight that it is a function of the fraction of fake agents ϵ and the filtration parameter e set by the Co-ordinator. Similarly, let $^{(b)}\Pi(\epsilon, e)$ denote the asymptotic welfare associated with this scenario, which relates to the cascade probabilities as per (10) as

$$^{(b)}\Pi(\epsilon, e) = (1/4) \left[^{(b)}\mathbb{P}_{Y\text{-cas}}^G(\epsilon, e) - ^{(b)}\mathbb{P}_{Y\text{-cas}}^B(\epsilon, e) \right]. \quad (42)$$

Now, as both processes: $\{h_{i_n}\}$ and $\{h_n\}$ share the same update rule, $^{(b)}\mathbb{P}_{Y\text{-cas}}^V(\epsilon, e)$ can be computed using the recursive method described in Section IV by equations (12) and (13), except that $p_f = a'$ if $V = G$, else $p_f = 1 - b'$ if $V = B$. Then, applying these probabilities to (42) yields $^{(b)}\Pi(\epsilon, e)$. For the best improvement in asymptotic welfare, the Co-ordinator needs to set e to the optimal value:

$$e^* = \arg \max_{e \in [0, 1]} ^{(b)}\Pi(\epsilon, e), \quad (43)$$

which yields the best welfare, $^{(b)}\Pi(\epsilon, e^*)$. As an example, Figure 9b plots this welfare against e for $p = 0.6$ and $\epsilon = 0.9$, where the best welfare is obtained at $e^* = 0.188$.

Scenario (c): Manipulating the Y 's

Here, as opposed to Scenario (a), the Co-ordinator modifies only the Y 's in the observations, i.e., at any time i , if $O_i = Y$ then the modified observation $O'_i = N$ w.p. $\beta \in (0, 1)$ and

$O'_i = Y$ w.p. $1 - \beta$. Whereas, if $O_i = N$ then $O'_i = N$ w.p. 1. The net effect of these random modifications is that, at any time i , the channel between action A_i and O'_i becomes a binary asymmetric channel with cross-over probabilities:

$$\begin{aligned} \epsilon_Y &:= \mathbb{P}(O'_i = Y | A_i = N) = \epsilon(1 - \beta) \\ \text{and } \epsilon_N &:= \mathbb{P}(O'_i = N | A_i = Y) = \beta. \end{aligned} \quad (44)$$

For this scenario, let the Y -cascade probability be denoted by ${}^{(c)}\mathbb{P}_{Y\text{-cas}}^V(\epsilon, \beta)$ to highlight that it is a function of the fraction of fake agents ϵ and the modification parameter β . Similarly, let ${}^{(c)}\Pi(\epsilon, \beta)$ denote the asymptotic welfare associated with this scenario. It is difficult to compute the Y cascade probabilities for general values of channel parameters: $\epsilon_Y \in (0, \epsilon)$ and $\epsilon_N = (0, 1)$, except when $\epsilon_Y = \epsilon_N$. This occurs if the Co-ordinator sets $\beta = \beta_{\text{sym}} := \frac{\epsilon}{1+\epsilon}$ at which the channel between A_i and O'_i effectively becomes a BSC with cross-over probability β_{sym} . In Appendix G of [25], we show that for this special case, a sufficient statistic of history \mathcal{H}'_n , at each time n is a random walk $\{s_n\}$ that occupies a finite state-space (unlike $\{h_n\}$ in (7) that typically has a countably infinite state-space). Thus, closed-form expressions for cascade probabilities exist and are derived in Appendix G. The resultant asymptotic welfare is

$${}^{(c)}\Pi(\epsilon, \beta_{\text{sym}}) = (1/4) \frac{a_{\text{sym}}^k - (1 - a_{\text{sym}})^k}{a_{\text{sym}}^k + (1 - a_{\text{sym}})^k}. \quad (45)$$

where $a_{\text{sym}} := p(1 - \beta_{\text{sym}}) + (1 - p)\beta_{\text{sym}}$ and $k := \left\lfloor \log_{(1 - a_{\text{sym}})/a_{\text{sym}}} \left(\frac{1-p}{p} \right) \right\rfloor + 1$.

Asymptotic welfare comparisons

We now compare the three asymptotic welfares, namely ${}^{(a)}\Pi(\epsilon, \kappa^*)$, ${}^{(b)}\Pi(\epsilon, e^*)$ and ${}^{(c)}\Pi(\epsilon, \beta_{\text{sym}})$ for fixed values of the private signal quality p . We also contrast these welfares with the default welfare in the absence of the Co-ordinator, $\Pi(\epsilon)$ and the baseline welfare $\Pi(0)$ (given by (28) and (29) resp.). Figure 10 plots the different welfares against ϵ for $p = 0.7$. It can be seen that for low values of ϵ , Scenario (c) provides the best improvement in welfare. We also observe that as $\epsilon \rightarrow 0$, Scenario (c) entirely mitigates the reduction in welfare caused by the presence of fake agents. For high values of ϵ , Scenario (b) provides the best improvement in welfare. Here, Scenario (c) in fact significantly worsens the welfare as compared to $\Pi(\epsilon)$. There also exist several intervals in $[0, 1]$ with moderate values of ϵ where Scenario (a) performs the best. We find the above characteristics of the three scenarios to be consistent for all values of p . We demonstrate this in Appendix H of [25] with plots for two more values, $p = 0.55$ and $p = 0.9$.

VIII. GENERAL CASE OF ZERO EX ANTE PAY-OFF

In this section, we extend our model to consider a general (possibly non-uniform) prior for the true value V of the item and a general pay-off structure for the agents, while still retaining the condition of zero *ex-ante* pay-off. Recall that this condition implies agents' indifference to the two actions a priori. We begin by assuming a non-revealing general prior for the true value, $\mathbb{P}(V = G) = q \in (0, 1)$. Recall from Section II that in general, the amount a buyer gains (loses) is x (y)

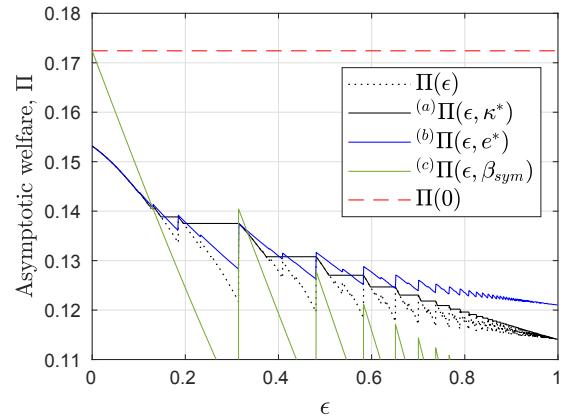


Fig. 10: Asymptotic welfare for the indicated scenarios versus the fraction of fake agents for $p = 0.7$.

if $V = G$ ($V = B$), where $x, y \geq 0$ and C denotes the cost of buying the item. Restricting the *ex-ante* pay-off to 0 implies that $C = qx - (1 - q)y$. Then, the general pay-off for any agent i is given by

$$\pi_i = \begin{cases} x - C = (1 - q)(x + y), & \text{if } A_i = Y \text{ and } V = G, \\ -y - C = -q(x + y), & \text{if } A_i = Y \text{ and } V = B, \\ 0, & \text{if } A_i = N. \end{cases} \quad (46)$$

In all previous sections, we have considered agents' pay-offs under the specific case of $x = 1$, $y = 0$ and $C = 1/2$.

In this more general scenario, the decision rule is similar to (2) except $\gamma_n = q$ is the new threshold at which agent n is indifferent to the two actions (and thus follows S_n). Next, using Bayes' rule, we express γ_n in terms of l_{n-1} and β_n as $\gamma_n = 1/(1 + \beta_n l_{n-1} \frac{1-q}{q})$. Then, the condition on γ_n for a $Y(N)$ cascade is $\gamma_n > q$ ($< q$) for all S_n , which translates to $\beta_n l_{n-1} < 1$ (> 1) for all S_n . This cascade condition is the same as before (see proof of Lemma 1) and hence is still defined in terms of l_{n-1} as per Lemma 1. As a result, the sequence of agents' actions once again satisfies Properties 1 and 2 where it is governed by $\{h_n\}$ which starts in the state $h_0 = 0$,⁶ evolves as per the update rule in (9), and is depicted in Figure 2. Now, conditioned on the true value V , $\{h_n\}$ does not change with the prior q . Thus, all conditional probabilities derived in this paper remain unaltered in the general scenario. On the contrary, agents' welfares being their pay-offs averaged over V with a general prior q will now change as follows. Evaluating (25) given that agent follows its private signal yields a new $F := q(1 - q)(x + y)(2p - 1)$. Next, the constant term $1/4$ is replaced by the term $q(1 - q)(x + y)$ in all equations pertaining to agents' welfares, namely, equations (27), (28), (29), (32), (33), (42), (45). However, in Eq. (28), Step (a) is no longer true and $\mathbb{P}_{\text{wrong-cas}} := q\mathbb{P}_{N\text{-cas}}^G + (1-q)\mathbb{P}_{Y\text{-cas}}^B$. With the above changes in place, all results and discussions presented in this paper extend to any case with zero ex ante pay-off. Our analytic techniques can also be readily modified for non-zero ex-ante pay-off, as done in [26].

⁶If the ex-ante pay-off is not restricted to 0, then for an arbitrary cost $C \in [-y, x]$, the resulting r.w. $\{h_n\}$ still updates as per (9), except that the starting state $h_0 = \log \left(\frac{q(x-C)}{(1-q)(y+C)} \right) / \log \left(\frac{p}{1-p} \right)$.

IX. CONCLUSIONS AND FUTURE WORK

We studied the effect of randomly arriving fake agents, who by taking a fixed action seek to influence the outcome of an information cascade. We focussed on the impact of varying the fraction of fake agents on the probability of their preferred cascade. To study this impact, we developed a Markov chain model which typically has a countably infinite state-space and does not readily allow for a closed-form solution to the cascade probabilities. Instead, we presented an iterative method that can compute the cascade probabilities with arbitrary precision. This process also yields exact values for any given agent's chances of herding and its welfare, which is the expected payoff it receives.

Our main result identified scenarios where surprisingly, an increase in the fraction of fake agents not only reduces the chances of their preferred cascade but also effects a significant improvement in the welfare of every rational agent. Further, we analysed three approaches to modify the observation database such that learning (asymptotic welfare) can be improved, namely, (1) increasing, (2) filtering out and (3) modifying the possibly fake actions. Interestingly, we observed that the third approach provides the best improvement in learning when the original fraction of fake agents is low, with a complete mitigation of their welfare-reducing effects as their fraction tends to zero. Whereas, the second approach performs the best for high values of this fraction. Lastly, we showed that our analysis, results and discussions readily extend to general priors and agent pay-offs, as long as the ex-ante pay-off is zero. For non-zero ex-ante pay-off, our analytic techniques can also be easily modified to study the platform's behaviour.

Potential future directions for this work include studying the effects of time-varying fractions and/or multiple types of fake agents, non-Bayesian rationality, random and asymmetric tie-breaking rules, and imperfect observations such as allowing each agent to only observe the total number for each action-type in its history, instead of perfectly observing the sequence of past actions. Moreover, as our results rely on the probability of fake agents being common knowledge, a natural extension would be to relax this assumption.

REFERENCES

- [1] P. Poojary and R. Berry, "Observational learning with fake agents," in *2020 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2020, pp. 1373–1378.
- [2] 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.
- [3] A. V. Banerjee, "A simple model of herd behavior," *The quarterly journal of economics*, vol. 107, no. 3, pp. 797–817, 1992.
- [4] I. Welch, "Sequential sales, learning, and cascades," *The Journal of finance*, vol. 47, no. 2, pp. 695–732, 1992.
- [5] M. Förster, A. Mauleon, and V. J. Vannetelbosch, "Trust and manipulation in social networks," *Network Science*, vol. 4, pp. 216–243, 2016.
- [6] N. Hu, I. Bose, N. S. Koh, and L. Liu, "Manipulation of online reviews: An analysis of ratings, readability, and sentiments," *Decision support systems*, vol. 52, no. 3, pp. 674–684, 2012.
- [7] K. K. Aköz, C. E. Arbatli, and L. Çelik, "Manipulation through biased product reviews," *Available at SSRN 3068345*, 2017.
- [8] D. Proserpio, B. Hollenbeck, and S. He, "How fake customer reviews do and don't work," <https://hbr.org/2020/11/how-fake-customer-reviews-do-and-dont-work>, Nov. 2020.

- [9] R. Kumar, S. Mukherjee, and N. P. Rana, "Exploring latent characteristics of fake reviews and their intermediary role in persuading buying decisions," *Information Systems Frontiers*, vol. 26, pp. 1091–1108, 2024.
- [10] Y. Huang, J. M. Villas-Boas, and M. Zhao, "Unmasking the deception: The interplay between fake reviews, rating dispersion, and consumer demand," *Rating Dispersion, and Consumer Demand*, 2023.
- [11] M. Leshno and Y. Spector, "An elementary proof of blackwell's theorem," *Mathematical Social Sciences*, vol. 25, pp. 95–98, 1992.
- [12] D. Blackwell, "Equivalent comparisons of experiments," *The annals of mathematical statistics*, pp. 265–272, 1953.
- [13] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.
- [14] 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.
- [15] Y. Song, "Social learning with endogenous observation," *Journal of Economic Theory*, vol. 166, pp. 324–333, 2016.
- [16] D. Sgroi, "Optimizing information in the herd: Guinea pigs, profits, and welfare," *Games and Economic Behavior*, vol. 39, 2002.
- [17] 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.
- [18] D. Acemoğlu, G. Como, F. Fagnani, and A. Ozdaglar, "Opinion fluctuations and disagreement in social networks," *Mathematics of Operations Research*, vol. 38, no. 1, pp. 1–27, 2013.
- [19] M. Mobilia, "Does a single zealot affect an infinite group of voters?" *Physical review letters*, vol. 91, no. 2, p. 028701, 2003.
- [20] Y. Peres, M. Z. Rácz, A. Sly, and I. Stuhl, "How fragile are information cascades?" *The Annals of Applied Probability*, vol. 30, no. 6, pp. 2796–2814, 2020.
- [21] J. Wu, "Helpful laymen in informational cascades," *Journal of Economic Behavior and Organization*, vol. 116, pp. 407–415, 2015.
- [22] 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.
- [23] T. M. Cover, "Hypothesis testing with finite statistics," *The Annals of Mathematical Statistics*, vol. 40, no. 3, pp. 828–835, 1969.
- [24] M. E. Hellman and T. M. Cover, "Learning with finite memory," *The Annals of Mathematical Statistics*, pp. 765–782, 1970.
- [25] P. Poojary and R. Berry, "Impact of fake agents on information cascades," *arXiv preprint*, vol. arXiv:2005.05518, 2025. [Online]. Available: <https://arxiv.org/abs/2005.05518>
- [26] P. Poojary and R. Berry, "Welfare effects of ex-ante bias and tie-breaking rules on observational learning with fake agents," in *2023 21st International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, 2023, pp. 334–341.
- [27] I. Arieli, R. Gradwohl, and R. Smorodinsky, "Herd design," *American Economic Review: Insights*, vol. 5, no. 4, pp. 460–76, December 2023.



Pawan Poojary received the M.S. degree in electrical engineering from the Indian Institute of Technology Madras, India, in 2018. He is currently a doctoral student at Northwestern University and a member of the Communications and Networking (CommNet) Laboratory. His research focuses on analyzing multi-agent learning and decision-making problems that arise from agent-based network interactions (eg. recommendation-based markets, ad platforms, online retail, social networks), using tools from dynamic estimation, Bayesian learning, stochastic modeling, game theory, network science and optimization.



Randall Berry joined Northwestern University in 2000, where he is currently the Chair and John A. Dyer Professor in the Department of Electrical and Computer Engineering. His research interests span topics in wireless communications, computer networking, network economics, and information theory. Dr. Berry received his PhD degrees in Electrical Engineering and Computer Science from MIT in 2000. He has served as an Editor for the IEEE Transactions on Wireless Communications, the IEEE Transactions on Information Theory, and an Area Editor for the IEEE Open Journal of the Communication Society. He has also served as the co-chair of the 2012 IEEE Communication Theory Workshop, a technical co-chair of 2010 IEEE ICC Wireless Networking Symposium, and steering committee chair for the IEEE DySPAN conference.