Quantifying the Utility of Imperfect Reviews in Stopping Information Cascades

Tho Ngoc Le, Vijay G. Subramanian, Randall A. Berry

***Abstract*— In models of social learning where rational agents observe other agents' actions, *information cascades* are said to occur when agents ignore their private information and blindly follow the actions of other. It is well known that in some cases, incorrect cascades happen with positive probability leading to a loss in social welfare. Having agents provide reviews in addition to their actions provides one possible way to avoid such "bad cascades." In this paper, we study one such model where agents sequentially decide whether or not to purchase a good, whose true value is either "good" or "bad." If they purchase, agents also leave a review, which is imperfect. We study the impact of such reviews on the asymptotic properties of cascades. For a good underlying state, we propose an algorithm that utilizes number theory principles and Markov chain analysis to solve for the probability of a wrong cascade. We discover that the probability of a wrong cascade is a non-monotonic function of the review strength. On the other hand, for a bad underlying state, the agents always eventually reach a correct cascade; we use a martingale analysis to bound the time until this happens.**

I. Introduction

Online platforms provide an easy way for people to attempt to learn from others before making a new decision. Such "social learning" has long been studied by economists as a game among Bayesian agents. In the simplest setting, these agents sequentially make a binary decision based on their own beliefs, which depend on their own private information as well as observations of the decisions of previous agents. A key result, first shown in [2] and [3], is that such models exhibit *information cascades*. This refers to a case where some agents ignore their private information and follow the previous agents' actions. Moreover, for the models in [2] and [3], once a cascade starts, all subsequent agents also cascade, leading to *herding*. Though individually optimal, this may result in the agents making a choice that is not socially optimal, i.e., a "wrong cascade."

A wrong cascade occurs because agents observe the actions of other agents *before* the other agents receive their pay-offs, and so these actions reflect the agents' estimates of the true pay-off and not the true pay-off itself. Indeed, if agents instead were able to see the true pay-off obtained by others, then as shown in [9] there would never be an incorrect cascade in which agents buy a bad product. The use of reviews in online recommendation systems can be viewed as an attempt to provide this information. However, due for example to user errors or product variability, such reviews may be an imperfect representation of this information (instead of the true pay-off as in [9]).

The goal of this paper is to study social learning in the presence of such imperfect reviews. More precisely, we consider a variation of the models in [2], [3], where agents have the option to either buy or not buy a given item, whose true value is one of two binary states ("good" or "bad"). In addition to the actions of the previous agents, agents also see a history of reviews before making decisions. However, these reviews do not reveal the true state of the good due to two effects: first, as we have already mentioned, these reviews can be imperfect, and second, agents can only leave a review if they buy the good and so no additional information is given for agents that choose not to buy.[1]

Adding reviews changes the information structure in [2], [3]. A number of other variations of this information structure have been considered, e.g., changing the signal structure to allow for bounded signals of different qualities ([15]), or allowing signals that are arbitrarily strong ([6]). Another variation is by changing the underlying network structure: in [8], the agents randomly sample the past observations; and in [11], each agent only observes a fixed number of his predecessors. Yet another strand of related work is the literature on "word-of-mouth" learning (e.g. [4], [5], [7]) in which agents can communicate information about the payoff of past actions. However, these models consider different settings (e.g. naive rule-of-thumb decisions) whereas our paper assumes that fully-rational agents can observe all past actions and reviews. Furthermore, this class of problems is also closely related to work on sequential hypothesis testing among distributed agents (e.g. [10], [11]), though in such cases the agents decision rules may be jointly designed.

In prior work ([12], [13]), we considered another variation on the information structure, where agents observed noisy observations of the *actions* of others. This led to the following counter-intuitive result: the probability of incorrect herding is non-monotonic in the noise level, i.e., in some cases, more noise is beneficial. In this paper, agents perfectly observe the actions of previous agents and the only noise is in the reviews. Additionally, since only agents who buy the good can submit reviews, this leads to an asymmetry in the model that was not present in [12], [13].

We presented an initial analysis of this model in [14]. There it was shown that the asymmetry in reviewing leads to an asymmetry in the resulting users' behaviors depending

This work was supported by the NSF under grant IIS-1538827 and the Vietnam Education Foundation (VEF) Fellowship. T. Le and R. Berry are with the EECS Dept of Northwestern University, and V. Subramanian is with the EECS Dept of the University of Michigan, Ann Arbor. Emails: thole2012@u.northwestern.edu, vgsubram@umich.edu, rberry@eecs.northwestern.edu

[1]For example, many online platforms such as Amazon.com indicate *verified purchase reviews*; in our model only such reviews are considered.

on the underlying state of the product being either "good" or "bad." Here, we present a refined analysis of these two cases. Conditioned on the product being good, we study the probability of an incorrect cascade. We give an algorithm based on number theoretic arguments that enables characterizing this probability for a much larger set of parameter settings than in [14]. Using this, we represent the wrong cascade probability as a function of the review strength and show that this is a highly non-monotonic and discontinuous function, so that in some cases increasing the review strength leads to a higher probability of a wrong cascade.[2] Conditioned on the state being bad, as a wrong cascade never occurs, we instead focus on the expected time until a correct cascade occurs. Using martingale techniques, we give bounds on this time. We compare these bounds with simulations, and also provide an algorithm to improve the lower bound numerically.

We organize this paper as follows. In Section II we specify our model. The main results are presented in Sections III and IV for the cases where the value of product is "good" and "bad," respectively. We conclude in Section V. The technical details can be found in the archived version [16].

II. Model

We consider a model similar to [14] in which there is a countable population of agents, indexed $n = 1, 2, \dots$ with the index reflecting the time and the order in which agents act (given exogenously). There is a new product (item) with a true value (V) that can be either good (G) or bad (B); for simplicity, both possibilities are assumed to be equally likely and the value is the same for all agents. Each agent n has a one-time action choice, A_n, of saying either "Yes" (Y) or "No" (N) to this item. If an agent chooses N, his payoff is 0. On the other hand, if he chooses Y, he faces a cost of $C = 1/2$ and two possibilities depending on the true value of the item: his gain is 0 if $V = B$, and his gain is 1 if $V = G$. Assume each agent n has prior knowledge about V via a private signal $S_n \in \{1 \text{ (high)}, 0 \text{ (low)}\}$. Each agent n who chooses $A_n = Y$ submits a review $R_n \in \{G \text{ (Good)}, B \text{ (Bad)}\}$ representing his experience with the item *after* purchasing. On the other hand, if choosing $A_n = N$, agent n does not submit a review.

We consider a homogeneous population where, conditioned on V, the private signals and reviews are independent across all agents. Furthermore, assume the probability that a private signal (resp. a review) aligns with V is $p \in (0.5, 1)$ (resp. $\delta \in [0.5, 1]$), i.e., the distributions of the signals and reviews are given as:

$$\mathbb{P}[S_n = 1|V = B] = \mathbb{P}[S_n = 0|V = G] = 1 - p,$$
$$\mathbb{P}[S_n = 1|V = G] = \mathbb{P}[S_n = 0|V = B] = p, \text{ and if } A_n = Y,$$
$$\mathbb{P}[R_n = G|V = G] = \mathbb{P}[R_n = B|V = B] = \delta,$$
$$\mathbb{P}[R_n = G|V = B] = \mathbb{P}[R_n = B|V = G] = 1 - \delta.$$

Since $p \in (0.5, 1)$, the private signals are informative, but not revealing; we call p the *signal quality*. On the other hand, δ denotes the *review's strength*. The review and the private signal are assumed to be conditionally independent given V.[3] Let $\mathcal{R}_n = R_n$ when $A_n = Y$ and $\mathcal{R}_n =^*$ when $A_n = N$. The history *after* agent n decides is written as $H_n = \{A_1, \mathcal{R}_1, \dots, A_n, \mathcal{R}_n\}$; we assume that H_n is public information to subsequent agents. The agents are Bayes-rational whose decisions are based on their own private signals and public information. Each agent n updates his posterior belief about the true value V using his private signal S_n, the actions $A_1, \dots, A_{n-1}$, and the reviews $\mathcal{R}_1, \dots, \mathcal{R}_{n-1}$[4]; agents then select the action that maximizes their expected pay-off.

A. Public likelihood ratio as a Markov process

Let $q = 1 - p$. Agents' decisions are based on Bayes updates of the posterior probability of $V = B$ versus $V = G$ given the observed history H_n. However, due to the independence of signals from the public history, agent $n + 1$ can instead compare the public likelihood ratio, ℓ_n, and his private belief, β_{n+1}, of $V = B$ versus $V = G$. Since V being B or G is equally likely, $\ell_0 = 1$ and we can rewrite ℓ_n in its alternate form:

$$\ell_n = \frac{\mathbb{P}[H_n|V = B]}{\mathbb{P}[H_n|V = G]}, \text{ and } \beta_{n+1} = \frac{\mathbb{P}[S_{n+1}|V = B]}{\mathbb{P}[S_{n+1}|V = G]}. \quad (1)$$

The higher ℓ_n is, the more likely that $V = B$. Moreover, since H_n is public information, ℓ_n can be updated as:
If agent n follows his own signal, then

$$\ell_n = \begin{cases} (p/q)\ell_{n-1}, & \text{if } A_n = N, \\ (q/p)[(1-\delta)/\delta]\ell_{n-1}, & \text{if } A_n = Y, \mathcal{R}_n = G, \\ (q/p)[\delta/(1-\delta)]\ell_{n-1}, & \text{if } A_n = Y, \mathcal{R}_n = B. \end{cases} \quad (2)$$

Otherwise, if agent n cascades, then

$$\ell_n = \begin{cases} \ell_{n-1}, & \text{if } A_n = N, \\ [(1-\delta)/\delta]\ell_{n-1}, & \text{if } A_n = Y, \mathcal{R}_n = G, \\ [\delta/(1-\delta)]\ell_{n-1}, & \text{if } A_n = Y, \mathcal{R}_n = B. \end{cases} \quad (3)$$

Moreover, as shown in Lemma 1, given ℓ_n one can determine if an agent cascades or not. Thus, $\{\ell_n\}$ is a Markov process. Moreover, this is also true if in addition we condition on each value of V[5]. On the other hand, $\beta_{n+1} = q/p$ (resp. p/q) if $S_{n+1} = 1$ (resp. $S_{n+1} = 0$).

B. Agents' decision rule and cascades' condition

By (3), any cascading action provides no information about V; thus let a_n be an integer random variables denoting the difference in the number of non-cascading Y actions and N actions. In addition, let r_n be another integer random variable denoting the number of good reviews minus the number of bad reviews in the history.

[2]This is similar in spirit to the result in [12], [13], though here the "noise" is in the reviews instead of the observations of the other agents actions.

[3]The motivation being, while signal quality reflects a product's marketing efficiency, the review strength is a consequence of product reliability, e.g., due to manufacturing.

[4]For simplicity, we assume indifferent agents follow their own signals.

[5]This is an extension of results from [6].

Lemma 1. *Define $x = \log_{\delta/(1-\delta)}(p/q) \in (0, \infty)$ for $\delta \in (0.5, 1)$. Then:*

1) $\ell_n = (q/p)^{h_n}$, where the exponent $h_n = a_n + \frac{1}{x} r_n$;

2) Conditioned on V, (a_n, r_n) and h_n are 2-D and 1-D Markov chains, respectively, for $n \geq 0$; and

2) Agent $n+1$ cascades Y if $h_n > 1$, cascades N if $h_n < -1$, and follows his private signal if $h_n \in [-1, 1]$.

Proof. 1) By (2) and (3), $\ell_n = (q/p)^{a_n} ((1-\delta)/\delta)^{r_n}$, thus h_n can be written in terms of a_n and r_n as above.

2) This is a direct consequence of the fact that $\{\ell_n\}$ is a Markov process and that, from the first property, there is a 1-1 correspondence between ℓ_n and h_n. Further, since a_n and r_n are integer-valued it follows that h_n only takes on a countable number of values. Without reviews a similar Markov chain was used in [12].

3) Since agent $n+1$ makes his decision by comparing $\ell_n \beta_{n+1}$ to 1, agent $n+1$ cascades Y if $\ell_n < q/p$, cascades N if $\ell_n > p/q$, and follows his signal if $\ell_n \in [q/p, p/q]$. By 1), this is translated to the given condition on h_n. $\square$

Note that x is an indicator of how strong the reviews are with respect to the signals. That is, the lower x is, the stronger the reviews are relative to the signals. For a generic x, the dynamics of the process $\{\ell_n\}$ can be studied by investigating the 2-D Markov chain (a_n, r_n). However, for special values of x, this can be simplified. We will study two such scenarios in Section III.

C. Asymmetry by different types of cascade and item quality

This model exhibits asymmetric behaviors with respect both to the types of cascades (Y and N), and to the true value V of the item. The main reason for this is the arrival of new information (reviews) depends on the action chosen by each agent. We first highlight a key difference between Y and N cascades in the following two properties.

Property 1. *Once a Y cascade starts, there is a positive probability that it ends (unless the review are perfect).*

If agent n faces $h_{n-1} > 1$, he chooses $A_n = Y$ regardless of his signal, and thus initiates a Y cascade. Such a cascade can end if subsequent agents submit a sufficient number of bad reviews, e.g., if $\mathcal{R}_n = B$, then $h_n = h_{n-1} - \frac{1}{x}$ could be below 1, which induces agent $n+1$ to use his signal. Furthermore, if x is sufficiently small then agent n's bad review can make $h_n < -1$, so that agent $n+1$ starts a N cascade. The dynamics of a Y cascade, once it gets started, are determined solely by the reviews process (and it does not depend on the signals). Regardless of the time a Y cascade was initiated, it can be broken by a sufficiently long sequence of bad reviews. Thus, the history process $\{H_n\}$ could include sample paths where Y cascades start and stop multiple times.

On the other hand, once $h_n < -1$, an N cascade starts, it lasts forever. This is because agents who choose N do not generate reviews; thus, the likelihood ratio stays constant as soon as any agent cascades to N. Subsequent agents are left in the same state as the one who initiated the cascade, and so make the same action choice. We summarize this in the following property:

Property 2. *Once an N cascade happens, it lasts forever.*

Next we give two properties that show the differences between a good and a bad product.

Property 3. *For $V = G$, a wrong cascade happens with positive probability.*

This is a result of the existence of absorbing states for a wrong cascade. For example, if the first two agents have low signals, they both choose N and no review is collected. Therefore, all subsequent agents are drawn into an N cascade, which is irreversible. This possibility cannot be avoided by adjusting the reviews strength, δ, even to perfect quality. In case the reviews are perfect, we still need a non-cascading agent who has an H signal for his review to be submitted.

Though a wrong cascade is possible, for $V = G$, it is more likely that there will be an abundance of information, since each agent that chooses Y also creates a new review. When $V = G$, since reviews are independent of the signals, as more agents choose Y, more information accrues with the underlying Markov process having a drift towards the correct direction, but there is no absorbing state on that side since h_n is unbounded above. On the contrary, multiple absorbing states for wrong cascades might exist. For $V = G$, the quantity of interest is, therefore, the probability of wrong (N) cascade, which is a function of both p and δ. We will discuss this scenario in Section III.

On the other hand, when $V = B$, this model exhibits a different set of behaviors. As more agents purchase the item, more and more reviews are collected. Since reviews are informative, subsequent agents can track the difference in the number of reviews to learn the true value of V, eventually. In other words, in addition to there exist trapping states only for the correct cascade, the drift also leans toward this side. We summarize this result below.

Property 4. *For $V = B$ and $\delta > 0.5$, a wrong cascade can never happen.*[6]

Thus, for $V = B$, a correct cascade happens with probability 1. In this scenario, we are interested in the distribution of the time (i.e. the number of agents) until a correct cascade happens. This will be studied in section IV.

III. Probability of wrong cascade for $V = G$

In the previous section, we discussed that wrong (N) cascades could happen if the product is good. In this section, we determine the probability of this happening. For a fixed p, as x varies, the conditions on a_n and r_n when cascades happen also change. As a result, the underlying Markov chains have different structures (both in terms of the state-spaces and the transition probabilities). Despite the complexity of these dynamics for a generic x, interesting and non-intuitive insights can be drawn by looking at special values of x. In one example, for any rational x, many states of (a_n, r_n) can

[6] Note if $\delta = 0.5$, then reviews are useless, in which case wrong cascades can occur as in [2].

be mapped to one single state of (h_n); thus, it is sufficient to study the reduced 1-D Markov chain (h_n). This is generally not possible for any real value of x since there would be a 1-1 mapping between the states in (a_n, r_n) and (h_n), which prevents the simplification of the state space. However, in another example when x is real and $x < 1/3$, the state space of (a_n, r_n) can also be simplified to obtain analytical results. Summarizing, we next consider two scenarios that facilitate simplification of the underlying state space of (a_n, r_n):

1) x is a rational number in $(0, \infty)$; and
2) x is any real number in $(0, 1/3]$.

A. $x = i/j$ for positive integers i, j and $gcd(i, j) = 1$

From the discussions at the beginning of this section, it is sufficient in this case to consider the 1-D Markov chain (h_n). Let $\mathbb{P}_s$ be the asymptotic probability of wrong cascade starting from the state $h_0 = s$. We want to find $\mathbb{P}_0$. Given i, consider the finite set $\mathscr{A} \triangleq \{-1, -\frac{i-1}{i}, \dots, \frac{i-1}{i}, 1\}$. It is obvious that $\mathscr{A}$ is the set of all possible values that $h_n = a_n + \frac{j}{i} r_n$ can take in $[-1, 1]$. Depending on the value of x, the following Lemma 2 further reduces the set of accessible states for $h_n \in [-1, 1]$ to different subsets of $\mathscr{A}$.

Lemma 2. *Assume $x = i/j$ is rational, where i, j are positive integers with $gcd(i, j) = 1$:*

1) If $x \leq 1/3$ or if $x \in \{1/2, 1\}$, $h_n \in \{-1, 0, 1\}$,

2) If $1/3 < x < 1/2$, let $z = j \bmod i$ and $k = \lfloor i/z \rfloor$ Then $h_n \in \{-1, 0, 1, -\frac{z}{i}, -\frac{2z}{i}, \dots, -\frac{kz}{i}, \frac{i-z}{i}, \frac{i-2z}{i}, \dots, \frac{i-kz}{i}\}$,

3) If $x > 1/2$, h_n takes all the values in $\mathscr{A}$.

Proof Idea. The proof uses number theoretic arguments to find the accessible states in $\mathscr{A}$ (see details in [16]). □

As a consequence of Lemma 2, we can numerically solve for $\mathbb{P}_0$. The idea is based on Markov chain analysis where one can write down a system of of linear equations (LEs) with the set of variables being $\mathbb{P}_{h_n}$ for all accessible states h_n. Since there is no absorbing state for a Y cascade, h_n is not upper-bounded and the accessible state space is infinite. However, once $h_n > 1$ the state transitions dynamics are simplified to a birth-death process; thus any variable $\mathbb{P}_{h_n}$ where $h_n > 1$ can be expressed in terms of the corresponding variables where $h_n \in \mathscr{A}$. Therefore, the number of equations is finite (at most $2i+1$). We give an algorithm in Algorithm 1 to construct this system of equations and solve for $\mathbb{P}_0$.

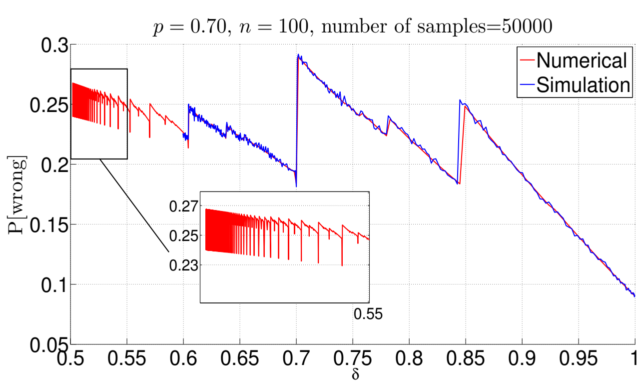

Figure 1: Wrong cascade probability versus δ for $V = G$.

Algorithm 1 Wrong cascade probability, $\mathbb{P}_0$, at rational x

---

Input: $V = 1$, p, $x = i/j$, $gcd(i, j) = 1$
Output: LEs and solution $\mathbb{P}_0$
  $\delta \leftarrow 1/\left(1 + (q/p)^{1/x}\right)$, $q \leftarrow 1 - p$, $\alpha \leftarrow (1-\delta)/\delta$
  Initialize $\mathscr{A} \leftarrow \{-1, -(i-1)/i, \dots, (i-1)/i, 1\}$
  $\mathscr{A} \supseteq \mathscr{A}' \leftarrow$ accessible states in $[-1, 1]$ (Lemma 2).
  for $h_n = s \in \mathscr{A}'$ **do**
    $s_L \leftarrow s - 1, s_{HB} \leftarrow s + 1 - j/i, s_{HG} \leftarrow s + 1 + j/i$
    $c_1 \leftarrow$ min number of steps from s_{HG} to $s_1 \in \mathscr{A}'$.
    if $s_{HB} > 1$ **then**
      $c_2 \leftarrow$ min number of steps from s_{HB} to $s_2 \in \mathscr{A}'$
    $Eq_s \leftarrow \mathbb{P}_s = q\mathbb{P}_{s_L} + p\delta\alpha^{c_1}\mathbb{P}_{s_1} + p(1-\delta)\alpha^{c_2}\mathbb{P}_{s_2}$
    Add equation Eq_s to the system of LEs
  Solve for $\mathbb{P}_0$ and return.

Note that the probability of wrong cascade is not monotonic in the review strength δ. As δ varies in $[0.5, 1]$, there are points of discontinuities resulting from the changes in the state space and the transition probabilities of the underlying Markov chains. Note that the mean number of agents needed until an N cascade occurs approaches infinity as $\delta \to 0.5$. In this regime, the time required to get an accurate simulation also blows up and so simulation results are not shown.

As a consequence of Lemma 2, for certain values of x the state space is simplified enough and we can obtain the closed-form expressions for the wrong cascade probability. In particular, Proposition 1 in [14] showed that for $x = 1$ and $x = 1/2$, $\mathbb{P}_0 = (q/p)^2$. Moreover, when there are no reviews, a result from [2] gives $\mathbb{P}_0 = (q/p)^2 / \left[(q/p)^2 + 1\right] < (q/p)^2$. Thus, having reviews with strength equal or double the signal quality strictly increases the probability of wrong cascades. For the case in Fig. 1, $\mathbb{P}_0 = 0.155$ with no reviews, and so it can be seen that for any $1/2 < x < 1$, reviews increase the probability of wrong cascades.

B. x is any real value in $(0, 1/3)$.

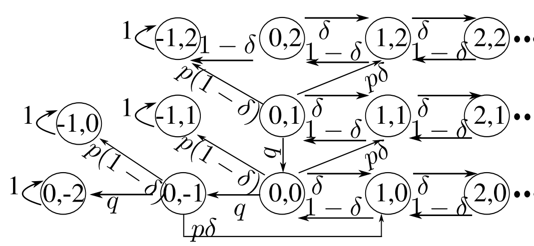

Figure 2: States transitions for $V = G$, and $x < 1/3$.

In this section, we present the second scenario when the state space (a_n, r_n) can also be simplified. In particular, we look at the cases when reviews are more than three times stronger than the private signals. For any real value of x in this region, the underlying 2-D Markov chains are shown in Fig. 2, where the first and second coordinates denote r_n and a_n, respectively. By Proposition 2 in [14], we can obtain:

$$\mathbb{P}_0 = [1 - p(2\delta - 2p\delta + 2p - p/\delta)] \,/\, [1 - 2pq(1-\delta)] \quad (4)$$

which is decreasing in δ. This is illustrated in Fig. 3 for $x = 1/5$, and $1/10$. For all values of x in this figure, the probability of wrong cascade decreases in the signal quality p. Moreover, except for reviews with perfect accuracy, there is a threshold $p^*(x)$ below which $\mathbb{P}_0$ is lower than when having no reviews.

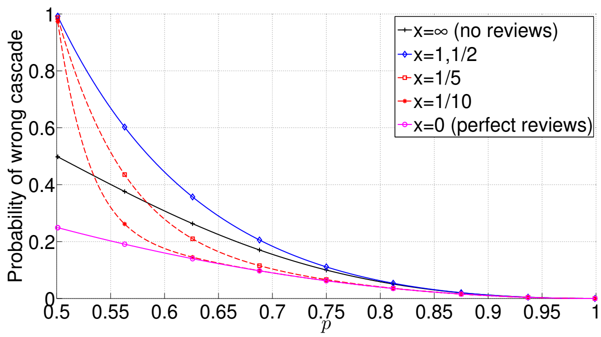

Figure 3: Wrong cascade probability for $V = G$.

The above conclusions can be explained by the discontinuity of the slopes for different curves in the above figure as $p \to 0.5$. With perfect reviews, the slope as $p \to 0.5$ is -1. With no reviews, the corresponding slope is -2. However, as long as the platform generates reviews with strength δ bounded away from 0.5, the probabilities of wrong cascade follow the set of dashed curves shown, with slopes bounded away from -2. In particular, these slopes can be studied using (4) by setting $p = 0.5 + \epsilon$ and $x = C\epsilon$ where $\epsilon \to 0, C > 0$. When C is fixed, δ is bounded away from 0.5; this yields a slope of $-\infty$ as $\epsilon \to 0$. When $C \to \infty$ and $x < 1/3$, $\delta \to 0.5$; and the slopes vary in $(-\infty, -8)$. Finally, if $C \to 0$, $\delta \to 1$; the slopes approach -1, which is exactly the slope of the perfect reviews scenario.

IV. Time until correct cascade for $V = B$

In Section II, we argued that for a bad product, only a correct (N) cascade can happen, so that it lasts forever once it occurs. In this section, we examine both upper and lower bounds on the expected time until correct cascades. In the following let $n \geq 0$. Conditioned on $V \in \{G, B\}$, let $\{\mathscr{F}_n^V\}$ be the sequence of σ-algebras generated by $\{H_n\}$. Similar to [6] and [8] where reviews do not exist, in our model the Markov process $\{\ell_n\}$ also exhibits the martingale property. In Section IV of [14] we showed that $\{1/\ell_n\}$ (resp. $\{\ell_n\}$) is a martingale process conditioned on $V = B$ (resp. $V = G$) adapted to the filtration $\{\mathscr{F}_n^B\}$ (resp. $\{\mathscr{F}_n^G\}$). Moreover, let X and Y be two random variables representing the increments $\Delta h_n = h_{n+1} - h_n$ for h_n in $[-1, 1]$ and $h_n > 1$, respectively. Let $f_1(\lambda)$ and $f_2(\lambda)$ be their corresponding moment generating functions (MGFs), where λ is a real variable. Let $\rho = \max(f_1(\lambda), f_2(\lambda))$ and define the random process $\{M_n\} = \left\{\frac{e^{\lambda h_n}}{\rho^n}\right\}$. Using techniques from [1], in [14] we showed that $\{M_n\}$ is a super-martingale adapted to $\{\mathscr{F}_n^B\}$. Let $\tau = \min\{n \geq 0 : h_n < -1\}$ be the stopping time when an N cascade happens. Now we use these results to bound the expected time until correct cascade, $\mathbb{E}[\tau]$.

A. Upper bound on $\mathbb{E}[\tau]$

Proposition 1. $\mathbb{E}[\tau] \leq e^{\lambda}/(1-\rho)$, *where* $0 < \rho < 1$, $\lambda \in (0, \ln(p/(1-p))$.

Proof. From Proposition 3 in [14], the tail distribution is upper-bounded by: $\mathbb{P}[\tau > n] \leq e^{\lambda}\rho^n$. For feasibility, we require $0 < \rho < 1$, thus $\lambda \in (0, \ln(p/(1-p))$. Now, since τ is a positive integer random variable, we can write:

$$\Rightarrow \mathbb{E}[\tau] = \sum_{n=0}^{\infty} \mathbb{P}[\tau > n] \leq e^{\lambda}/(1-\rho).$$

□

The above bound is a function of the dummy variable λ, and the two MGFs f_1, f_2. Our objective is to find λ and ρ that minimizes this bound. We solve this numerically and compare the minimum bound with the mean time obtained by Monte-Carlo simulations for different values of p and δ.

B. Lower bound on $\mathbb{E}[\tau]$

Let $\tilde{\rho} = 1/\max(f_1(\lambda), f_2(\lambda))$ for regions where $0 < \tilde{\rho} < 1$, i.e. $\lambda \in (\ln(p/(1-p)), \infty)$. The following Proposition provides a lower bound on $\mathbb{E}[\tau]$.

Proposition 2. $\mathbb{E}[\tau] \geq e^{-\lambda}\tilde{\rho}[1 - A]/[M_1(1-\tilde{\rho})]$, *where* $0 < \tilde{\rho} < 1$, $\lambda \in (\ln(p/(-p)), \infty)$,

$A = \tilde{\rho}^4 + (\tilde{\rho} - \tilde{\rho}^4)\mathcal{P}_1 + (\tilde{\rho}^2 - \tilde{\rho}^4)\mathcal{P}_2 + (\tilde{\rho}^3 - \tilde{\rho}^4)\mathcal{P}_3$, *and* $\mathcal{P}_n = \mathbb{P}[\tau = n | h_1]$ *for* $n = 1, 2, 3$.

Proof Idea. The proof uses the super-martingale property of $\{M_n\}$ and total probability theorem using the three possible values of h_1. Conditioned on each h_1, we calculate the probabilities of τ taking the first three values $n = 1, 2, 3$. We then use these probabilities to provide a lower bound on $\mathbb{E}[\tau]$. See our archived version in [16] for details. □

Note that the above lower bound is then numerically maximized over $\lambda \in (\ln(p/(1-p), \infty)$. Moreover, due to computational constraints, the bound in Proposition 2 is obtained using the closed-form expressions of $\mathcal{P}_n$ for $n = 1, 2, 3$. Next, we present an algorithm that improves this lower bound by numerically calculating $\mathcal{P}_n$ for higher values of n.

Algorithm 2 Finding $\mathbb{P}[\tau = n | h_1]$

Input: $V = 0$, p, δ, n
Output: $\mathcal{P}_n = \mathbb{P}[\tau = n | h_1]$
    idea: Build a breadth-first tree conditioned on h_1,
    tree.add(root), qualified = empty list of qualified nodes
    while tree.notempty() **do**
        Pick the first node j at lowest level i by BFS
        Check for early elimination, e.g. $h_j > 1 + (n-i)/x$
        if $i < n$ (not a leaf) **then** check for $h_j \geq -1$
            if True **then** tree.add(j's children)
                update condition on node j's children
        else(leaf) check for $h_j < -1$
            if True **then** qualified.add(j)
        tree.remove(j)
    Update $\mathcal{P}_n$ using the qualified list of leaves, return $\mathcal{P}_n$

C. Numerical and Simulation results

In Fig. 4 below, we use Algorithm 2 to show how the lower bound can be improved as n is increased. Conditioned on each h_1, computational constraints limit us to using at most $n = 17$, which generates approximately 10^5 possible realizations of the history that would lead to an N cascade. The algorithm offers greater improvement for lower values of δ. The non-monotonicities and discontinuities of the lower bound are a consequence of the same behaviors of each $\mathcal{P}_n$.

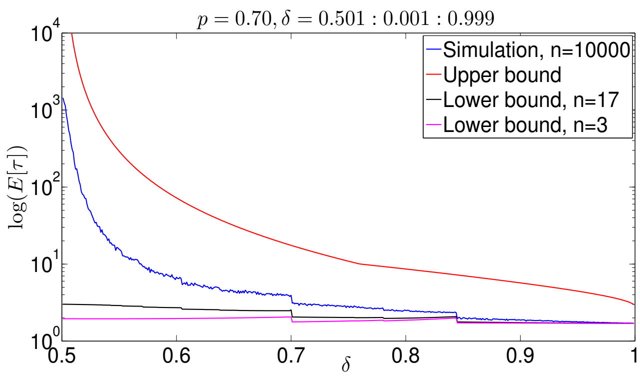

Figure 4: Bounds of $\log(\mathbb{E}[\tau])$ versus simulation.

Fig. 4 also shows the numerical bounds as compared with simulations on log-scale. Simulation results showed that $\mathbb{E}[\tau]$ is decreasing as δ increases. As $\delta \to 1$, the lower bound offers a better approximation to simulations; both converge to the same value of $1 + p$. On the other hand as $\delta \to 0.5$, both the upper bound and the result from simulations blow up, while the lower bound does not have this behavior.

In fact, we can verify that as $\delta \to 0.5$, we have $\mathbb{E}[\tau] \to \infty$. For any $\delta \in [0.5, 1]$, there is a positive probability of the underlying Markov chain hitting a state in the region where $h_n > 1$. In this region, the process becomes a simple birth-death process with transition probabilities δ, $1 - \delta$ to the left and right, respectively. By relabeling the states, assume that we start at a state $i > 0$ in this birth-death process where 0 is an absorbing state on the left and there is no absorbing state on the right. Let $\tau_i = \min\{n > 0 : h_n = 0 | h_0 = i\}$ be the stopping time when the absorbing state is 0. For this birth-death process, the recurrence equation is written as: $(1 - \delta)\mathbb{E}[\tau_{i+1}] - \mathbb{E}[\tau_i] + \delta\mathbb{E}[\tau_{i-1}] = -1$. This generates a general solution of the form $\mathbb{E}[\tau_i] = A\left(\frac{\delta}{1-\delta}\right)^i + B + \frac{i}{2\delta-1}$, where A, B are real constants. Using the boundary condition $\mathbb{E}[\tau_0] = 0$, we have $A = -B$. Moreover, since $\mathbb{E}[\tau_i] > 0$, we require $A \geq 0$. Now assume that $\delta = 0.5 + \epsilon$ where we let $\epsilon \to 0$. As a result, $2\delta - 1 = 2\epsilon \to 0$ and $\mathbb{E}[\tau_i] \to \infty$. But since $\mathbb{E}[\tau_i]$ gives a lower bound on the original $\mathbb{E}[\tau]$, we also have $\mathbb{E}[\tau] \to \infty$ as $\delta \to 0.5$.

V. Conclusions and Future Work

This paper studied a Bayesian learning model with information cascades. We assumed that subsequent agents can observe perfectly the previous actions and, in addition, feedback in the form of imperfect reviews depending on the actions. We showed that the probabilities that agents cascade toward the wrong actions are not monotonic in the review strength. In particular, imperfect reviews could increase the probability that agents misinterpret the true value of a good product. In practice, in online platforms like Yelp, Amazon, etc. customers reviews have a wide range of variability in strengths. Even though this scenario was not considered in this paper and our previous work, our results indirectly imply that a platform planner should opt to display only the reviews of high strengths while eschewing those with low strengths. In fact, this strategy is already adopted by those platforms, e.g. Amazon with verified purchase reviews, or Yelp with filtered reviews. Moreover, our results suggest that no matter how strong the reviews are, agents might not perform better if their prior knowledge is poor in quality. This implies that a platform planner should consider spending their budget on improving the product's marketing efficiency.

In the future work, we plan to study the possibility of having reviews with strengths non-homogeneously distributed across the population. In addition, we would like to study the probability of wrong cascades for more generic relationships between the signals quality and the reviews strength. Other possible directions include considering having reviews when both type of actions are taken, letting agents have the option to leave the reviews, and assuming that not all agents would exercise this option.

References

[1] B. Hajek, *Hitting and occupation time bounds implied by drift analysis with applications*, Adv. Appl. Prob., vol. 14, pp. 502-525, 1982.

[2] S. Bikhchandani, D. Hirshleifer, and I. Welch. *A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades*, J. Polit. Econ., vol. 100, No. 5, pp. 992-1026, 1992.

[3] A. Banerjee, *A Simple Model of Herd Behavior*, The Quarterly Journal of Economics, vol. 107, pp. 797-817, 1992.

[4] G. Ellison, D. Fudenberg, *Rules of thumb for social learning*, Journal of Political Economy, vol. 101, pp. 612-643, 1993.

[5] G. Ellison, D. Fudenberg, *Word-of-mouth communication and social learning*, The Quarterly Journal of Economics, pp. 93-125, 1995.

[6] L. Smith, P. Sorensen, *Pathological Outcomes of Observational Learning*, Econometrica, vol. 68, pp. 371-398, 2000.

[7] A. Banerjee, D. Fudenberg, *Word-of-mouth learning*, Games and Economic Behavior, vol. 46, pp. 1-22, 2004.

[8] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, *Bayesian learning in social networks*, Rev. Econ. Stud., vol. 78, pp. 1201-1236, 2011.

[9] H. Cao, B. Han, and D. Hirshleifer, *Taking the road less traveled by: Does conversation eradicate pernicious cascades?*, Journal of Economic Theory, vol. 146, pp. 1418-1436, 2011.

[10] Z. Zhang, E. Chong, A. Pezeshki, W. Moran, *Hypothesis Testing in Feedforward Networks With Broadcast Failures*, JSTSP, vol. 7, issue 5, pp. 797-810, 2013.

[11] K. Drakopoulos, A. Ozdaglar, J. N. Tsitsiklis, *On Learning With Finite Memory*, IEEE Transactions on Information Theory, vol. 59, issue 10, pp. 6859-6872, 2013.

[12] T. Le, V. Subramanian, R. Berry, *The Value of Noise for Informational Cascades*, IEEE ISIT Proceeding, 2014.

[13] T. Le, V. Subramanian, R. Berry, *The Impact of Observation and Action Errors on Informational Cascades*, IEEE CDC Proceeding, 2014.

[14] T. Le, V. Subramanian, R. Berry, *Are Noisy Reviews Helpful in Social Learning*,IEEE ISIT Proceeding, 2014.

[15] J. Wu, *Helpful Laymen in Informational Cascades*, Journal of Economic Behavior & Organization, 2015.

[16] T. Le, V. Subramanian, R. Berry, *Quantifying the Utility of Noisy Reviews in Stopping Information Cascades*, longer archived version of this paper, available at: `https://archive.org/details/CDC2016Archive`