Bayesian Learning with Random Arrivals

Tho Ngoc Le
Data Scientist
JD.com
Santa Clara, CA 95054
thonle2012@gmail.com

Vijay G. Subramanian
Dept. of EECS
University of Michigan
Ann Arbor, MI 48109
vgsubram@umich.edu

Randall A. Berry
Dept. of EECS
Northwestern University
Evanston, IL 60208
rberry@eecs.northwestern.edu

***Abstract*—We add to a line of work considering the impact of observation imperfections in models of Bayesian observational learning. In particular, we study a discrete-time model in which in each time-slot, an agent may randomly arrive. Agents who arrive have the opportunity to buy a given item. If an agent chooses to buy, this action is recorded for subsequent agents. However, the decisions of agents that choose not to buy are not recorded. Hence, if no one buys in a given slot, agents are unaware if this was due to no agent arriving or an agent choosing not to buy. We study the impact of this uncertainty on the emergence of information cascades. Using a Markov chain based analysis, we show that the probability of incorrect cascades and the expected time until a cascade happens are not monotonic in the arrival probability of a user. We find that adding a small uncertainty in the arrival information from the perfect information setting will make a buy cascade happen with higher probability than a not-buy cascade. However, if the agents' private signals are weak, then a not-buy cascade is more likely to occur for most arrival rates, resulting in wrong cascades dominating when the item is good and vice-versa when the item is bad.**

I. Introduction

A key feature of many on-line platforms is that they provide users with data about the actions of the other users such as how many users bought a given item or a summary of the users' experiences via reviews. A subsequent user can then attempt to learn from this data about the action she should take, e.g., whether to buy an item or not. Such settings can be studied under the framework of Bayesian observational learning, which has its roots in the economics literature (e.g. [2], [3]). In this framework, the agents are viewed as players in a dynamic game with incomplete information who form beliefs about the value of different actions based on observations of the actions of other players as well as their own private signals. In the simplest setting, these agents sequentially make a binary decision while perfectly observing the actions of all prior agents. A key result, first shown in [2] and [3], is that such models may exhibit *information cascades*, meaning that at some point an agent ignores their own private information about the item and chooses an action solely dependent on their type, which for homogenous agents results in blindly following the actions of the preceding agents. 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" where an item is bought even when it is not beneficial to do so or *vice versa*.

Many variations of these types of models have been studied, such as differences in the types of signals received by each agent [4] or allowing agents to observe only a subset of past actions [6]. In [5] and [7], the agents can see only a random sample of past observations, which are of unknown order. In this paper, we instead assume that agents may randomly arrive in discrete time even though the order of arrivals is public information. In previous work [9], we considered another variation in which the arrivals were deterministic but there were random errors in the observations of the *actions* of other agents, hence imperfect observations. This led to the following counter-intuitive result: the probability of a wrong cascade is non-monotonic in the error level, i.e., in some cases, a higher error rate is beneficial. In this work, we consider a different way in which the observations of the agents may be imperfect.

Our work is motivated by on-line platforms that report information on how many users have bought a given item, but do not specify how many users have considered buying the item and did not buy. For example, e-commerce websites do not publicize their products' landing pages conversion rates. Specifically, we consider a discrete-time model in which at each time-slot an agent randomly arrives with a given probability (and doesn't arrive with the complementary probability). If an agent arrives it has the opportunity to either buy an item or not. If she chooses to buy, this action is recorded for all subsequent agents to see. On the other hand, if either no agent arrives in a certain epoch or one arrives and chooses not to buy, the subsequent agents will observe an "empty slot" in such an epoch. This introduces uncertainty in the observation history: if there is an empty slot, the agents are unsure as to whether it was because an agent chose not to buy or simply because no agent arrived.

In our model, as in [2], [3], [9], all agents receive the same value from buying the given item, which is determined by an unknown binary state of the world that takes values from *{Good, Bad}*. If the state is *Good*, agents benefit from buying the item, while if it is *Bad*, agents are better off not buying. Each agent receives an independent, identically distributed binary signal about this state of the world. Our Bayes-rational agents then determine the *a posteriori* probability distribution of the state of the world given their private signal and observations, and choose to buy the item if their expected pay-off is greater from that action than from not buying.

In our prior work [9], random errors occurred in the observations, where the error rate is assumed to be the same for each action choice ("buy" or "not buy"). This led to a symmetry in the model: an observation of an agent buying and not buying conveys the same amount of information if agents are not herding. In this paper, instead the uncertainty in the observations is asymmetric. A slot in which an agent arrives and buys conveys more information than an "empty" slot. Further, the information in an empty slot decreases as the arrival probability decreases.

We show that as in [2], [3], [9], an information cascade will eventually occur with probability one, and once started, such a cascade will persist forever, i.e., the agents will herd. We then study the probability of a wrong cascade. To do this we utilize a Markov chain based analysis in which herding corresponds to absorbing states. Using this, we show that the probability of wrong/incorrect herding is not monotonic in the probability that an agent arrives in any time epoch. In some cases having a higher arrival probability (and thus more information conveyed in empty slots) increases the probability of a wrong cascade. We also discover that if the private signals are weak, the probability of a wrong cascade can be higher than that of a correct cascade. Moreover, we find that adding a small uncertainty will make a buy cascade happen with higher probability than a not-buy cascade due to a bias toward buy actions in the history.

We organize the paper as follows. In Section II we specify our model. The model's properties and an information theoretic explanation are presented in Section III. The main results are presented in Section IV. We conclude in Section V. The technical details can be found in the archived version [10].

II. Model

We consider a variation of the model in [9] in which time is divided into discrete slots indexed $n = 1, 2, \ldots$. Different from [9], we assume that in each time slot n, an agent is present with a probability $a \in [0, 1]$.[1] If an agent is present at the time slot n, he chooses an action A_n of either buying (Y) or not buying (N) a new item. The true value (V) of the item can be either good (G) or bad (B) and is the same for all agents. For simplicity, both possibilities are assumed to be equally likely.

The agents are Bayes-rational utility maximizers[2] whose payoff structure is based on the agent's action and the true value V. If an agent chooses N, his pay-off is 0. On the other hand, if he chooses Y, he faces a cost of $C = 1/2$ and gains one of two amounts depending on the true value of the item: his gain is 0 if $V = B$ and 1 if $V = G$. The total pay-off of an agent choosing Y is then the gain minus the cost. Thus, the *ex ante* expected pay-off of each agent is 0.

To reflect the agents' prior knowledge about V, if an agent arrives at time n, he receives a private signal $S_n \in \{1 \text{ (high)}, 0 \text{ (low)}\}$ through a binary symmetric channel (BSC) with crossover probability $1 - p$, where $0.5 < p < 1$. In other words, we have

$$\begin{aligned} &\mathbb{P}(S_n = 1|V = G) = \mathbb{P}(S_n = 0|V = B) = p, \text{ and} \\ &\mathbb{P}(S_n = 0|V = G) = \mathbb{P}(S_n = 1|V = B) = 1 - p. \end{aligned} \quad (1)$$

Thus, the private signals are informative, but not revealing.

Denote an observation at the time slot n as O_n, we have either $O_n = Y$ when an agent is present at time n and chooses to buy the item, or $O_n = E$ denoting an empty slot if either there is no agent arriving at time n or there is one who chooses not to buy the item. Denote the observation history after time n as $H_n = \{O_1, \ldots, O_n\}$. Similar to [9], we assume that H_n is recorded via a common database that is available to all subsequent agents but without any further imperfections.

III. Model properties and Bayesian updates

In this section, we present the properties of our model. In particular, the observation history H_n can be summarized by a sufficient statistic that is a weighted difference of the number of Ys and Es in the history, which we will elaborate on shortly. As in [9], both Y and N cascades are permanent. However, different from [9], there is an asymmetry in the type of observations. As the probability that an agent is present in any time slot, a, decreases, an empty slot $O_n = E$ conveys less information even though an observation $O_n = Y$ conveys the same amount of information. As a result, it requires a greater number of empty slots to create an N cascade as a decreases.

A. *Public likelihood ratio as a Markov process*

Let $q = 1 - p$. Similar to [9], 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 conditional independence of signals from the public history given the value of V, if present agent $n+1$ can instead compare the public likelihood ratio, ℓ_n, and his private belief, β_{n+1}, of $V = B$ versus $V = G$. There is one subtle difference, however. In [8] and [9], ℓ_n denotes the ratio after an agent n has decided. In this paper, ℓ_n is the public likelihood ratio after the discrete time slot n, as an agent may not be present at time n. Since V being B or G is equally likely, $\ell_0 = 1$ and we can rewrite ℓ_n and β_{n+1} as follows:

$$\ell_n = \frac{\mathbb{P}[H_n|V = B]}{\mathbb{P}[H_n|V = G]}, \text{ and } \beta_{n+1} = \frac{\mathbb{P}[S_{n+1}|V = B]}{\mathbb{P}[S_{n+1}|V = G]}. \quad (2)$$

The higher ℓ_n is, the more likely that $V = B$. Moreover, since H_n is public information, if after slot $n - 1$ a cascade does not happen, ℓ_n can be updated as:

$$\ell_n = \begin{cases} \frac{1-aq}{1-ap}\ell_{n-1}, & \text{if } O_n = E, \\ \frac{q}{p}\ell_{n-1}, & \text{if } O_n = Y \end{cases} \quad (3)$$

Otherwise, if agent $n - 1$ cascades, then $\ell_n = \ell_{n-1}$ with probability 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 the value of V. In

[1] Two special cases need highlighting: i) $a = 0$, where the model is degenerate with no incoming agents; and ii) $a = 1$, where the agents are always present and yields the model in [2] with symmetry in both states of the world.

[2] This assumption leads to a Perfect Bayesian Equilibrium that uses the common information based belief to determine strategies, as proved in [6].

addition, we can show that $\{\ell_n\}$ (resp. $\{1/\ell_n\}$) is a martingale conditioned on $V = G$ (resp. $V = B$). 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, and a cascade is permanent. Thus, conditioned on a cascade not occurring before time $n-1$, let a_n be a random variable denoting the number of Y actions in the history until n. In addition, let $e_n = n - a_n$ denote the number of empty slots at time n in the history. We have the following lemma:

Lemma 1. *Let $y = \log_{p/q} \frac{1-aq}{1-ap} \in (0,1)$ for $a \in (0,1)$. Then:*

1) $\ell_n = (q/p)^{h_n}$, where the exponent $h_n = a_n - ye_n$;

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

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

Proof. 1) By (3), $\ell_n = (q/p)^{a_n} \left[(1-ap)/(1-aq)\right]^{e_n}$, thus h_n can be written in terms of a_n and e_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 e_n are integer-valued it follows that h_n only takes on a countable number of values.

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

Note that y is an indicator of how weak an empty slot is with respect to the signals. That is, the lower y is, the weaker the empty slots are relative to the signals. In addition, y is an increasing function of a, and $y = 0$ (resp. 1) if and only if $a = 0$ (resp. 1). For a generic y, the dynamics of the process $\{\ell_n\}$ can be studied by investigating the 2-D Markov chain (e_n, a_n). However, for special values of y, this can be simplified. We will study a few of such scenarios in Section IV.

Note that both Y and N cascades are permanent and the probabilities of reaching each cascade is positive. In addition, in contrast to [9] where Y and N observations in succession cancel each other out, here Y and E observations in succession leads to a bias towards a Y cascade. Hence, a sufficiently long sequence of Y and E pairs can also create a Y cascade. In Section IV, we present results on the probability of each cascade conditioned on the state of the world.

C. An information theoretic view on the parameter y

Recall that given a cascade has not occurred, the sufficient statistic is $h_n = a_n - ye_n$, where we can rewrite y as:

$$y = \left[\log(1-aq) - \log(1-ap)\right] / \left[\log(p) - \log(q)\right]. \quad (4)$$

This is the exponent of the public likelihood ratio, ℓ_n. More precisely, let $L_n = \log \ell_n$, be the public log-likelihood ratio, then we have

$$L_n = h_n \log(q/p) = -h_n(\log(p) - \log(q)). \quad (5)$$

Next we give an information theoretic interpretation of these results. Following [1], we can express the public log-likelihood ratio given an observation sequence x_n in terms of relative entropies (K-L divergenges) as follows:

$$L_n(x^n) = n\left[D(P_{x^n}|P_B) - D(P_{x^n}|P_G)\right] \quad (6)$$

where P_{x^n} is the empirical distribution of the observations, and P_G and P_B are the probability of an observation conditioned on V being equal to G and B, respectively.

Given a public observation history at time n consisting of a_n Y actions and e_n empty slots, we have that $n = a_n + e_n$, $P_{x^n}(y) = \frac{a_n}{a_n+e_n}$ and $P_{x^n}(e) = \frac{e_n}{a_n+e_n}$. Moreover, we have that $P_B(y) = aq$ and $P_B(e) = 1 - aq$. Likewise, $P_G(y) = ap$ and $P_B(e) = 1 - ap$. Using these in (6) gives

$$\begin{aligned} L_n(x^n) &= a_n \log\left(\frac{a_n/(a_n+e_n)}{aq}\right) + e_n \log\left(\frac{e_n/(y_n+e_n)}{1-aq}\right) \\ &- \left(a_n \log\left(\frac{a_n/(a_n+e_n)}{ap}\right) + e_n \log\left(\frac{e_n/(a_n+e_n)}{1-ap}\right)\right) \\ &= a_n \log(p/q) - e_n \log((1-aq)/(1-ap)). \end{aligned}$$

Note the parameter y in (4) is exactly the ratio between the weights given to the a_n and e_n terms here and using this we have the relationship in (5), as expected.

Using this we can further interpret this parameter as follows. Let P_e denote the empirical distribution given a single observation of an empty slot (i.e., this is simply a single atom on e). Then we have

$$D(P_e|P_B) - D(P_e|P_G) = \log\left(\frac{1}{1-aq}\right) - \log\left(\frac{1}{1-ap}\right) \quad (7)$$

$$= -\log((1-aq)/(1-ap)). \quad (8)$$

Likewise, let P_y denote the empirical distribution given a single observation of a Y action, so that

$$D(P_y|P_B) - D(P_y|P_G) = \log(1/aq) - log(1/ap) = \log(p/q).$$

Comparing with the above we have that

$$-y = \left[D(P_e|P_B) - D(P_e|P_G)\right] / \left[D(P_y|P_B) - D(P_y|P_G)\right]$$

i.e., y gives the ratio of the value of observing a single empty slot to that from observing a single Y action.

IV. Analysis and Results

In general, the underlying Markov chain is a 2-D chain with the states being the pairs (e_n, a_n), which denote the number of empty slots and the number of Y actions in the history. For rational values of y, the state space of the chain can be enumerated and this facilitates the analysis of the chain's asymptotic properties. The following lemma further simplifies the 2-D chain to a 1-D chain with the states being $h_n = a_n - ye_n$, and then enumerates these states.

Lemma 2. *Let $y = \frac{k}{l}$, where k, l are integers satisfying $gcd(k,l) = 1$. Then the non-cascading state space of the underlying 1-D Markov chain is the finite set $\mathscr{A} \triangleq \{-1, -\frac{l-1}{l}, \ldots, \frac{l-1}{l}, 1\}$.*

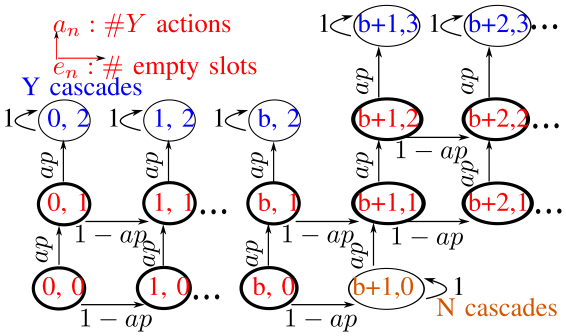

Figure 1: The 2-D MC for rational values $y = k/l$ and $V = G$, where $l = kb + c, c = 0, \dots, k-1, b = 1, 2, \dots$

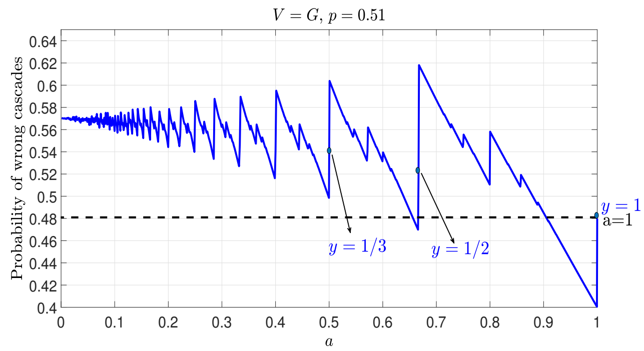

Figure 2: Wrong cascade probability versus a, the probability of an agent being present in each time slot, for $V = G$, $p = 0.51$.

Given that $h_n \in [-1, 1]$, its values must belong to a subset of $\mathscr{A}$. Lemma 2 further states that the state space before a cascade happens is exactly $\mathscr{A}$ for all rational values of y.

Proof. See Appendix A in our archived version [10]. □

In Fig. 1, we show the state space and the transition dynamics of the underlying Markov chain for rational values of $y = k/l$. A related state space simplification was given for the model in [8]; however, a different approach is required here. In [8], the state transitions in a Y-cascade region follow a simple birth-death chain; this allows the recurrence equations to convert any states $h_n > 1$ to a corresponding state in $[-1, 1]$. In this paper, there is no birth-death chain. However, for any rational y the states will eventually repeat themselves in a structured way as shown in Fig. 1. We rely on this observation to simplify those states to a finite subset, based on which a finite set of linear equations can be written down and solved to get $\mathbb{P}$ and $\mathbb{E}\tau$. In the remainder of this section, we provide closed-form expressions for $\mathbb{P}$ and $\mathbb{E}\tau$ for a few values of y when $V = G$. Note that for $V = B$, similar expressions can be derived and these are given in the appendices of our archived version [10].

A. Probability of wrong cascades

For a rational value of y, we can use Lemma 2 to numerically solve for and plot the wrong cascade probabilities, $\mathbb{P}$. As a result, we can show that the probability of wrong/incorrect cascade is not monotonic in the probability that an agent arrives in any time epoch, as demonstrated in Fig. 2 for a low value of the signal quality (i.e., $p = 0.51$). The dashed lines in these figures show the corresponding probability when $a = 1$. There are points of discontinuity which happen when the y value changes. In addition, for a sufficiently low value of $p = 0.51$, as shown in Fig. 2, the probability of a wrong cascade can be higher than that of a correct cascade. Moreover, adding a small uncertainty will make a Y cascade happen with higher probability than a N cascade due to the bias toward Y actions in the history. Results for different values of p can be found in our archived version [10].

For special values of y, we can derive closed-form expressions of this probability, which give additional insights on the discontinuities seen in the numerical results. In particular when $V = G$, let $\alpha = ap$ denote the probability of having one more Y action in the history before a cascade happens. For any positive integer l, the following two propositions provide closed-form formulae for the probability of wrong cascade, $\mathbb{P}$, when $y = 1/l$ and $y = (l-1)/l$, respectively.

Proposition 1. *Let $y = 1/l$, where l is any positive integer. When $V = G$, the probability of wrong cascades, $\mathbb{P}$, satisfies:*

$$\mathbb{P} = (1-\alpha)^{l+1} / \left[1 - \alpha(1-\alpha)^l(1+l)\right]. \tag{9}$$

Proof. See Appendix B in our archived version [10]. □

Proposition 2. *Let $y = (l-1)/l$, where the positive integer $l \geq 3$.*[3] *Let $x = \alpha(1-\alpha)$. When $V = G$, the probability of wrong cascades, $\mathbb{P}$, satisfies:*

$$\mathbb{P} = \frac{(1-\alpha)^2\left[2\frac{1-x^{l-2}}{1-x} + 2\alpha^{l-2}(1-\alpha)^{l-1} - 1\right]}{1 - 4\alpha^{l-1}(1-\alpha)^l}. \tag{10}$$

Proof. See Appendix C in our archived version [10]. □

For a fixed p, both Propositions 1 and 2 show that $\mathbb{P}$ decreases as a increases such that we always satisfy $y = 1/l, (l-1)/l$, respectively. In other words, $\mathbb{P}$ is increasing in l in Proposition 1 and is decreasing in l in Proposition 2. More insights into the discontinuities at $a = 0$ and $a = 1$ can be gained by studying the limits of $\mathbb{P}$ when $y = 1/l, (l-1)/l$ as $l \to \infty$ for a fixed p. In particular, in the first scenario, let $l \to \infty$ when $y = 1/l$; this yields $a \downarrow 0^+$. Using Proposition 1, $\lim_{a\downarrow 0^+} \mathbb{P} = \frac{1}{\frac{p(\log q - \log p)}{p-q} + \left(\frac{p}{q}\right)^{\frac{p}{p-q}}}$, which is higher than the value of $\mathbb{P} = q$ when $a = 0$. In the second scenario, let $l \to \infty$ when $y = (l-1)/l$; this yields $a \uparrow 1^-$. Using Proposition 2, we have $\lim_{a\uparrow 1^-} \mathbb{P} = \frac{q^2(1+pq)}{1-pq}$, which is lower than the value of $\mathbb{P} = \frac{q^2}{q^2+p^2}$ when $a = 1$.[4] Therefore both scenarios again demonstrate discontinuities that happen at $a = 0$ and $a = 1$, as shown in Fig. 2. More importantly, both limits also show that a lower wrong cascade probability can be achieved if a is near one of these discontinuities by

[3] If $l = 1$ then $y = 0$, $a = 0$ and this reduces to a trivial case of no observation history. If $l = 2$, then $y = 1/2$ and one can refer to Proposition 1 for an expression of $\mathbb{P}$.

[4] This expression is a consequence of the wrong cascade probability in [2], [3], and in [9] for the special case where there are no observations errors (i.e., $\epsilon = 0$.)

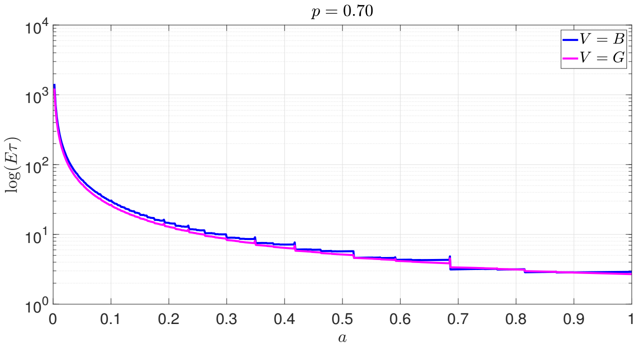

Figure 3: Expected time until a cascade happens versus a, for $p = 0.70$.

reducing it a small amount so that $\mathbb{P}$ moves to the left hand side of the discontinuity. In other words, agents can do better if a controlled amount of uncertainty is introduced through reducing the probability of arrival a. This observation applies for the case $V = G$. However, in contrast, when $V = B$, one can show that reducing uncertainty through increasing a leads to better outcomes.

B. Expected time until a cascade happens, $\mathbb{E}\tau$

Using similar approach, one can numerically solve for the expected time until a cascade occurs, $\mathbb{E}\tau$. The results are shown in Fig. 3 for $p = 0.70$. Depending on the range of a, different values of $V \in \{G, B\}$ can lead to higher $\mathbb{E}\tau$. There are points of discontinuity which happen as different values of a cause changes in the underlying state space. One interesting observation is that unlike the case $V = G$, for $V = B$ the expected time $\mathbb{E}\tau$ is higher as p increases, given a fixed a. (See Figs. 22, 24, 26 in Appendix G in [10].) The reason being a higher signal quality leads to more agents having low signals; therefore a higher number of empty slots induces a longer time to N cascades. Furthermore for values of $y \in \{1/l, (l-1)/l\}$, one again can write down the closed-form expressions for $\mathbb{E}\tau$. These expressions are given in the next two propositions.

Proposition 3. *Let $y = 1/l$, where l is any positive integer. When $V = G$, the expected time until a cascade, $\mathbb{E}\tau$, satisfies:*

$$\mathbb{E}\tau = \frac{2 - (l+1)(1-\alpha)^l + (l-1)(1-\alpha)^{l+1}}{\alpha\,[1 - (l+1)\alpha(1-\alpha)^l]}. \quad (11)$$

Proof. See Appendix D in our archived version [10]. □

Proposition 4. *Let $y = (l-1)/l$, where the positive integer $l \geq 3$. Let $x = \alpha(1-\alpha)$. When $V = G$, the expected time until a cascade, $\mathbb{E}\tau$, satisfies:*

$$\mathbb{E}\tau = \frac{2\frac{1-x^l}{1-x} + 2(1-\alpha)\frac{x-x^{l-1}}{1-x} + 2\alpha^{l-2}(1-\alpha)^l}{1 - 4\alpha^{l-1}(1-\alpha)^l}. \quad (12)$$

Proof. See Appendix E in our archived version [10]. □

For a fixed p, Proposition 3 shows that $\mathbb{E}\tau$ is increasing in l (i.e., decreasing in a). Similarly, Proposition 4 shows that $\mathbb{E}\tau$ is decreasing in l (i.e., decreasing in a). More insights into the discontinuities at $a = 1$ can be gained by studying the limits of $\mathbb{E}\tau$ when $y = (l-1)/l$ as $l \to \infty$ for a fixed p.[5] In particular, let $l \to \infty$ when $y = (l-1)/l$; this yields $a \uparrow 1^-$. Using Proposition 4 gives $\lim_{a\uparrow 1^-} \mathbb{E}\tau = 2 + 2p\frac{pq}{1-pq}$, which is lower than the value of $\mathbb{E}\tau = \frac{2}{1-2pq}$ when $a = 1$. Therefore, adding a little uncertainty leads to a faster expected time toward either type of cascades. Moreover, we can also show that $\lim_{a\uparrow 1^-} \mathbb{E}\tau$ is not monotonically decreasing in p. This means when a little uncertainty is present, a higher signal quality does not necessarily lead to a faster time to a cascade. (See Fig. 17 in Appendix G in [10].)

V. Conclusions and Future Work

This paper considers Bayesian observational learning in which agents arrive in discrete time-slot and only one action ("buying") is recorded. This introduces an asymmetry in the observation history as a slot in which an agent arrives and buys conveys more information than an "empty" slot. Further, the information in an empty slot decreases as the arrival probability decreases. We show that an information cascade will eventually occur with probability one, and once started such a cascade will persist forever, i.e., the agents will herd. We then utilize a Markov chain based analysis to study the probability of a wrong cascade and the expected time until a cascade happens. We show that both quantities are not monotonic in the probability that an agent arrives in any time epoch. In some cases, having a higher arrival probability (and thus more information conveyed in empty slots) leads to worse outcomes. Moreover, we also discover that if the private signals are weak, the probability of a wrong cascade can be higher than that of a correct cascade. In the future, we are interested in generalizations where agents also arrive randomly and leave reviews (as in [8]) that are observable (but not their actions).

References

[1] T. Cover, J. Thomas, *Elements of Information Theory*, New York: Wiley, 1991.

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

[5] A. Banerjee, D. Fudenberg, *Word-of-mouth learning*, GEB, vol. 46, pp. 1-22, 2004.

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

[7] L. Smith, P. Sorensen, *Rational social learning with random sampling*, working paper. Available at SSRN 1138095, 2013.

[8] T. Le, V. Subramanian, R. Berry, *Quantifying the Utility of Noisy Reviews in Stopping Information Cascades*, IEEE CDC Proceeding, 2016.

[9] T. Le, V. Subramanian, R. Berry, "Information Cascades With Noise," *IEEE TSIPN* , vol. 3, no. 2, pp. 239-251, June 2017.

[10] T. Le, V. Subramanian, R. Berry, *Learning From Randomly Arriving Agents*, longer archived version of this paper, available at: https://archive.org/details/randomly_arriving_agents.

[5] Note that there is little interest in the limiting value of $\mathbb{E}\tau$ as $a \to 0$, since if $y = 1/l$ and we let $l \to \infty$, then $\mathbb{E}\tau \to \infty$, which can be verified from both Fig. 3 and equation (11).