Observational Learning with Negative Externalities

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

***Abstract*—Observational learning models seek to understand how distributed agents learn from observing the actions of others. In the basic model, agents seek to choose between two alternatives, where the underlying value of each alternative is the same for each agent. Agents do not know this value but only observe a noisy signal of the value and make their decision based on this signal and observations of other agents' actions. Here, instead we consider a scenario in which the choices faced by an agent exhibit a *negative externality* so that value of a choice may decrease depending on the history of other agents selecting that choice. We study the learning behavior of Bayesian agents with such an externality and show that this can lead to very different outcomes compared to models without such an externality.**

I. Introduction

There is a long history of work studying models for "observational learning," in which Bayesian agents make decisions based in part on their observations of other agents. Early papers in this area include [1]–[3] that studied models where homogeneous agents sequentially make a decision between two alternatives such as buying or not buying a given service. This service has a common quality (say "good" or "bad") that is unknown to each agent. Agents receive a noisy private signal indicating this quality and make decisions based on their signal and their observations of other agents' actions. A key results for such models is the emergence of *information cascades*, i.e., scenarios in which at some point agents ignore their private signal and simply follow the action of all previous agents. Many variations of such a model have been studied including [4]–[14].[1] In this prior work, the value an agent obtains from a given action does not depend on the action of any prior agent. Here, we instead consider a model that exhibits *negative externalities* meaning that the value of a given service may decrease depending if other agents choose the same service. We show that such externalities can lead to very different learning behavior compared to models without externalities.

In our model agents sequentially choose between two services, A and B. The value they obtain from a service decreases in the number of other agents choosing that service within a given time window, where this window can model the time an agent spends using the service. The two services differ in the degree of this externality so that in one high (H) quality service the externality is lower than in the other low (L) quality service. Again agents do not know a service's quality but only receive a noisy signal indicating this and then choose a service based on their signal and their observations of earlier agents. For example, this could model a setting where wireless devices (agents) have to choose between two frequency channels (services) accounting for the fact that a more crowded channel would yield a poorer quality of service [17], [18]. Another example might be drivers selecting between two parking garages, where the number of other vehicles in a garage could lead to a dis-utility due to more time spent finding parking [19], [20]. To see how this externality can impact learning note that even if an agent believes that service A is the better service, it may decide to choose service B because it has a lower externality. Our main result shows that, unlike in previously studied models, there exist parameters, wherein the "right" cascade does not correspond to the *optimal* action sequence had the agents known the true state of the world. We also show that the probability with which the "right" cascade occurs is not monotonically increasing in the quality of agents' private signals.[2]

II. Model

We consider a model similar to [1] in which there is a countable sequence of agents, indexed $i = 1, 2, \ldots$ where the index represents both the time and the order of actions. Each agent i takes an action $U_i \in \mathcal{U} = \{A, B\}$ of choosing to occupy one of the two available services A and B. After choosing a service, each agent i remains in that service over the next m time slots, $i+1, i+2, \ldots i+m$. Here, m represents the *service-time* for any agent, and is the same for both services. While it is common knowledge that one of the services provides a higher Quality of Service (QoS) than the other, the identity of the better (or the poorer) service is not known to the agents *a priori*. These true qualities are denoted by the pair $(V_A, V_B) \in \mathcal{V} = \{(H, L), (L, H)\}$, where $(V_A, V_B) = (H, L)$ implies that service A is better than service B; $(V_A, V_B) = (L, H)$ implies the opposite. For simplicity, both possibilities of (V_A, V_B) are assumed to be equally likely.

The difference in service qualities is reflected in the agent's cost structure, which depends on its action (choice of service), the true quality (V_A, V_B), and the number of other users of the service in the following manner. For the n^{th} agent, let the history of past actions be denoted by $H_{n-1} = \{U_1, U_2, \ldots, U_{n-1}\}$. Given this history, at time n before agent n acts, let services A

This work was supported in part by the NSF under grants CNS-1908807 and CCF-1934931/001.

[1]This line of work also has ties to work on distributed inference, e.g. [15], [16].

[2]This is similar in spirit to results in [8], [11] where noise in the observations causes this non-monotonic behavior, while here this occurs without any observation noise.

and B have m_A and m_B respective occupants. Note that due to the finite service-time m, m_A and m_B are determined by only the m most recent actions $H_{n-1}^{(m)} \triangleq \{U_{n-m}, \ldots, U_{n-1}\}$, and $m_A + m_B = m$. The cost incurred by agent n, $c_n : \mathcal{V} \times \mathcal{U} \times \mathcal{H}_{n-1} \to \mathbb{R}$ is then defined as:

$$c_n = \begin{cases} m_{U_n} + 1, & \text{if } V_{U_n} = H, \\ (m_{U_n} + 1)k, & \text{if } V_{U_n} = L, \end{cases} \tag{1}$$

where $k > 1$ denotes the *quality-factor* by which one service is poorer (costlier) than the other. Here, $\mathcal{H}_{n-1}$ denotes the set of all possible action histories H_{n-1} observed by agent n.

Every agent i receives a private signal $S_i \in \{A, B\}$. This signal reflects the agent's private beliefs and partially reveals the true service qualities through a binary symmetric channel (BSC) with crossover probability $1-p$, where $1/2 < p < 1$ (see Fig. 1). Each agent i takes a *rational* action U_i that depends on its private signal S_i and the past actions H_{n-1}. Private signals are assumed to be conditionally independent across agents, given (V_A, V_B).

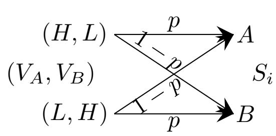

Fig. 1: The BSC through which agents receive private signals.

III. Optimal decision and cascades

For the n^{th} agent, the history of past actions H_{n-1} and its private signal S_n forms its information set $\{S_n, H_{n-1}\}$. As the first agent does not have any observation history, he always follows his private signal, i.e., he chooses service A if and only if the signal is A. For agent $n \geq 2$, the Bayes' optimal action, U_n is chosen such that it incurs the least expected cost given the information set $\{S_n, H_{n-1}\}$. Let $\gamma_n(S_n, H_{n-1}) \triangleq \mathbb{P}((H, L)|S_n, H_{n-1})$ denote the posterior probability that A is better than B, i.e., $(V_A, V_B) = (H, L)$. The expected cost E_{U_n} that agent n incurs by taking action $U_n \in \{A, B\}$, given $\{S_n, H_{n-1}\}$ is then expressed by:

$$\begin{aligned} E_{U_n} = c_n((H, L), U_n, H_{n-1})\gamma_n \\ + c_n((L, H), U_n, H_{n-1})(1 - \gamma_n). \end{aligned} \tag{2}$$

The Bayes' optimal decision rule is then:

$$U_n = \begin{cases} A, & \text{if } E_A < E_B, \\ B, & \text{if } E_A > E_B, \\ \text{follows } S_n, & \text{if } E_A = E_B. \end{cases} \tag{3}$$

When $E_A = E_B$, an agent is indifferent between the actions. Similar to [8], our decision rule in this case follows the private signal S_n, unlike [1] where random tie-breaking is used.

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

It follows from (3) that, agent n cascades to service A (B) if and only if $E_A < E_B$ ($E_A > E_B$) for all $S_n \in \{A, B\}$. The other case being $E_A \leq E_B$ for $S_n = A$ and $E_A \geq E_B$ for $S_n = B$; in which case, agent n follows S_n. To better understand the above cascade conditions, we encapsulate the information contained in the history H_{n-1} observed by agent n in the term $g_{n-1}(H_{n-1}) \triangleq (\Delta_{(L,H)}/\Delta_{(H,L)})l_{n-1}(H_{n-1})$, where $l_{n-1}(\cdot) \triangleq \mathbb{P}(\cdot|(L, H))/\mathbb{P}(\cdot|(H, L))$ is the likelihood ratio function of the action history H_{n-1}. Further $\Delta_{(V_A, V_B)}$ denotes the difference between the costs incurred for actions A and B given the true qualities (V_A, V_B) and is given by:

$$\Delta_{(L,H)}(H_{n-1}) = c_n((L, H), B, H_{n-1}) - c_n((L, H), A, H_{n-1}),$$
$$\Delta_{(H,L)}(H_{n-1}) = c_n((H, L), A, H_{n-1}) - c_n((H, L), B, H_{n-1}).$$

For a history H_{n-1}, where services A and B have m_A and m_B occupants, $\Delta_{(L,H)}(H_{n-1}) = (m_B + 1) - k(m_A + 1)$ and $\Delta_{(H,L)}(H_{n-1}) = (m_A + 1) - k(m_B + 1)$. A more intuitive way to characterize an information cascade is given next.

Lemma 1: Given a history H_{n-1} such that $\Delta_{(H,L)}(H_{n-1}) > 0\ (< 0)$; agent n cascades to service A if $g_{n-1} > \frac{p}{1-p}$ $\left(g_{n-1} < \frac{1-p}{p}\right)$, cascades to service B if $g_{n-1} < \frac{1-p}{p}$ $\left(g_{n-1} > \frac{p}{1-p}\right)$, and otherwise follows its private signal S_n.

To prove Lemma 1, define $\beta(\cdot) \triangleq \mathbb{P}(\cdot|(L, H))/\mathbb{P}(\cdot|(H, L))$ as the likelihood ratio of the private signal S_n, where $\beta(A) = (1-p)/p$ and $\beta(B) = p/(1-p)$. Applying Bayes' formula gives $\gamma_n = \frac{1}{1+\beta_n l_{n-1}}$. Now, using the expressions for E_A and E_B from (2) given $\{S_n, H_{n-1}\}$, the inequality $E_A > E_B$ can be simplified to the form $g_{n-1}(H_{n-1}) >$ (or $<$) $1/\beta(S_n)$. Agent n cascades to a service only if the above inequality holds for both $S_n = A$ and B. This gives bounds on $g_{n-1}(H_{n-1})$ for a cascade to occur which completes the proof. Here, recall that the term g_{n-1} has a denominator $\Delta_{(H,L)}$, which if negative will flip the sign of the inequality. This in turn swaps the conditions for agent n to cascade to services A and B as stated in Lemma 1. Next, the evolution of the likelihood ratio process $\{l_n\}$ is characterized by Lemma 2, and is a common property of cascade models studied so far, such as [1]–[4], [8].

Lemma 2: Given a history H_n, with $I \subseteq \{1, 2, \ldots, n\}$ denoting the set of past agents who have followed their private signals, the likelihood ratio $l_n(H_n)$ depends only on the difference between the number of A's (denoted by n_A) and B's (denoted by n_B) in the set $F_n = \{U_i : i \in I\}$. Specifically, $l_n = \left(\frac{1-p}{p}\right)^{h_n}$, where $h_n = n_A - n_B$.

The proof follows by noting that for any agent $i \notin I$, the action U_i does not provide any additional information about the true qualities (V_A, V_B) to the successors beyond what is contained in H_{i-1}. As a result, $l_i = l_{i-1}$. On the other hand, if agent i does not cascade ($i \in I$), then Lemma 1 implies that it follows its private signals S_i, which means $U_i = S_i$. Now, as S_i is conditionally independent of the history H_{i-1} given (V_A, V_B), $l_i = (\frac{1-p}{p})l_{i-1}$ if $U_i = A$, else $l_i = (\frac{p}{1-p})l_{i-1}$ if $U_i = B$. Thus, $l_n = \left(\frac{1-p}{p}\right)^{h_n}$ for h_n as defined in the lemma.

Note that in models such as [1]–[4], [8], [11], the condition for agent n to cascade, solely depends on the likelihood ratio: $l_{n-1}(H_{n-1})$. Thus, if agent n cascades, Lemma 2 implies that $l_j = l_{n-1}$ for all $j \geq n$. Therefore once a cascade occurs, it lasts forever with subsequent agents herding to the cascading action. In our model, however, a cascade at time n does not necessarily imply $g_n = g_{n-1}$. So a cascade may not cause subsequent agents to cascade or herd to a fixed action.

A sufficient statistic of an agent's observation history

In prior models such as [1]–[3], [8], [11], the likelihood ratio l_{n-1} is a sufficient statistic for agent n's past observations, which facilitates analyzing these models. Here, the analogous quantity is g_{n-1}. However, it follows from Lemma 1 that this is not a sufficient statistic of the history H_{n-1}. To see this, consider the following example with $m = 2$ and $k \in (1, 3)$. At time n, let H_{n-1} and H'_{n-1} be two different histories such that $H_{n-1}^{(m)} = \{A, A\}$, $H_{n-1}^{'(m)} = \{B, B\}$ and $h_{n-1} = 2$, $h'_{n-1} = -2$. Now, there exists a value of k, given by $k_0 = (3\alpha^2 + 1)/(3 + \alpha^2)$, where $\alpha := \frac{p}{1-p}$, for which $g_{n-1} = g'_{n-1}$. However, it follows from Lemma 1 that given H_{n-1}, agent n cascades to B, whereas given H'_{n-1} it cascades to A. This is because the two histories have dissimilar service occupancies such that while $\Delta_{(H,L)}(H_{n-1}) > 0$, $\Delta_{(H,L)}(H'_{n-1}) < 0$. Note that $\Delta_{(H,L)}$ and $\Delta_{(L,H)}$ only depend on $H_{n-1}^{(m)}$; whereas l_{n-1} is fully represented by h_{n-1} (Lemma 2). Hence, it follows that $\big(H_{n-1}^{(m)}, h_{n-1}\big)$ is a sufficient statistic of the history H_{n-1}.

Next, we explore how $H_n^{(m)}$ and h_n influence the $(n+1)^{\text{th}}$ agent's choice. First, consider the following scenario. At time $n+1$, let service A have more occupants than service B, i.e., $m_A > m_B$ and let $h_n \in \{0, 1, 2, \ldots\}$, i.e., given the history of past actions H_n, it is more (or equally) likely that A is of a higher quality than B. Here, the question arises: should the agent choose B as it is less congested? Or should it choose A as it has a higher probability of being the better service? Or should it use its private signal S_{n+1} to resolve this issue? It turns out that the answer to these questions depends on the value of k, as per the following lemma.

Lemma 3: Consider the Markov state $\big(H_n^{(m)}, h_n\big)$ such that $h_n \in \{0, 1, 2, \ldots\}$ and $m_A > m_B$. Define $\hat{m} = (m_A + 1)/(m_B + 1)$ and $\alpha = p/(1-p)$. For $r = 0, 1, \ldots, (|h_n| + 1)$ define the increasing sequence of thresholds $\{k_r\}$, given as:

$$k_r = \begin{cases} \dfrac{\hat{m}\alpha^{(|h_n|+1-r)} - 1}{\alpha^{(|h_n|+1-r)} - \hat{m}}, & \text{for } p > \Big(\hat{m}^{\frac{-1}{(|h_n|+1-r)}} + 1\Big)^{-1} \\ \infty, & \text{o.w.} \end{cases} \tag{4}$$

Agent $n+1$ cascades to B if the quality factor $k < k_0$, follows its private signal S_{n+1} if $k \in [k_0, k_2]$ and cascades to A if $k > k_2$. The same thresholds $\{k_r\}$ exist for the opposite Markov state $\big(H_n^{(m)}, h_n\big)$ where $h_n \in \{0, -1, -2, \ldots\}$, $m_A < m_B$ and $\hat{m} = (m_B + 1)/(m_A + 1)$.

Remark 1: For any private signal quality p, the thresholds $\{k_r\}$ defined in Lemma 3 satisfy $\hat{m} < k_r$ for all r. Further, k_r for any r decreases with p.

Next, consider the scenario where service A has more occupants than B, i.e., $m_A > m_B$ but $h_n \in \{-1, -2, \ldots\}$ i.e., given the history H_n, it is more likely that service B is of a higher quality than A. In this case, both $H_n^{(m)}$ and h_n favor the choice of B, hence the agent will choose B regardless the value of k. We state this in the next lemma.

Lemma 4: Consider the Markov state $\big(H_n^{(m)}, h_n\big)$ such that $h_n \in \{-1, -2, \ldots\}$ and $m_A > m_B$. For all values of k, agent $n+1$ always cascades to B. Likewise, for $h_n \in \{1, 2, \ldots\}$ and $m_A < m_B$, agent $n+1$ always cascades to A.

Lemma 5: Consider the Markov state $\big(H_n^{(m)}, h_n\big)$ such that $m_A = m_B$. Then, the $(n+1)^{\text{th}}$ agent's action solely depends on the value of $h_n \in \mathbb{Z}$. It cascades to service A (B) if $h_n > 1$ (< -1) and otherwise follows its private signal S_{n+1}.

Next we consider: what would the agents' optimal actions be if everyone knew the the true qualities?

Definition 2: The *optimal action sequence* refers to the repeating pattern of agents' actions if all agents know the true quality (V_A, V_B) a priori.
For example, consider $(V_A, V_B) = (H, L)$ and $m = 2$. The optimal action sequence for $k = 1.5$ is found to be AAB, whereas for any $k > 3$, it is AAA. Similarly, for any general m, the following holds:

Remark 2: For any quality factor $k > m + 1$, the optimal action sequence given the qualities (H, L) is AAA, and for (L, H) is BBB.

IV. Markovian Analysis of Cascades for $m = 2$

In this section, we consider the service time $m = 2$. Let $p_f \triangleq \mathbb{P}(S_n = A | (V_A, V_B))$ denote the probability that an agent observes private signal A. Depending on the true qualities, $p_f = p$ for $(V_A, V_B) = (H, L)$ whereas $p_f = 1 - p$ for $(V_A, V_B) = (L, H)$. It follows from the previous section that the process $\big\{(H_n^{(m)}, h_n)\big\}$ is a discrete-time 2-D Markov process (m.p.) where h_n takes integer values. For $m = 2$, $H_0^{(m)} = \{\}$, $H_1^{(m)} \in \big\{\{A\}, \{B\}\big\}$ and at any time $n \geq 2$, $H_n^{(m)} \in \big\{\{A, A\}, \{A, B\}, \{B, A\}, \{B, B\}\big\}$. Figure 2 depicts this 2-D Markov chain for quality factor $k < \min(t_0, k_0)$, where t_0 and k_0 are thresholds defined shortly. To begin, $H_0^{(m)} = \{\}$ and $h_0 = 0$ since the first agent has no observation history. Hence the m.p. starts at state $(\{\}, 0)$. In this starting state, $g_0 = 0$ and so the 1st agent follows its private signal S_1. Therefore, w.p. p_f the next state is $(\{A\}, 1)$, and otherwise is $(\{B\}, -1)$. At time $n = 1$, $H_1^{(m)} \in \big\{\{A\}, \{B\}\big\}$. Suppose that at $n = 1$, the m.p. is in state $(\{A\}, 1)$. At this point, the transition to the next state at time $n = 2$ depends on the quality factor k. Applying Lemma 3 to state $(\{A\}, 1)$ gives thresholds $\{t_r\}_{r=0}^2$ where $t_2 = \infty$. It follows that for $k \in (1, t_0)$, the m.p. transitions from $(\{A\}, 1)$ to $(\{A, B\}, 1)$ w.p. 1 (shown in Figure 2). On the other hand for $k \in [t_0, \infty)$, the m.p. transitions from $(\{A\}, 1)$ to $(\{A, A\}, 2)$ w.p. p_f and otherwise to $(\{A, B\}, 0)$. Next, for all states where $m_A = m_B$, namely: $\big(\{A, B\}, h_n\big)$ and $\big(\{B, A\}, h_n\big)$ for any h_n; the behaviour of agents is dictated solely by h_n, as stated in Lemma 5. Lastly, for the remaining states, namely: $\big(\{A, A\}, h_n\big)$ and $\big(\{B, B\}, h_n\big)$, once again the value of k determines whether an agent chooses to cascade to A, or B or follow its private signal (as per Lemma 3). Let thresholds $\{k_r\}_{r=0}^3$ be defined as per (4) for the state $\big(\{A, A\}, 2\big)$. Then specifically for $k < k_0$, it follows from Lemma 3 that agents in states $\big(\{A, A\}, h_n\big)$, $h_n \in \{0, 1, 2\}$ cascade to B, whereas agents in states $\big(\{B, B\}, h_n\big)$, $h_n \in \{0, -1, -2\}$ cascade to A (see Fig. 2).

We highlight the following salient features of the 2-D m.p. in Fig. 2. First, any state transition that corresponds to a cascade occurs w.p. 1 and only translates along the vertical axis. This is due to Lemma 2,

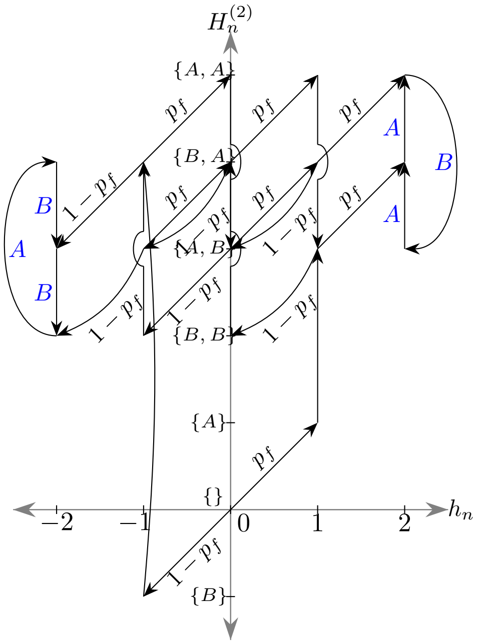

Fig. 2: Transition diagram of the 2-D m.p. for service time $m = 2$. Here, quality factor $k < \min(t_0, k_0)$.

which implies that $h_n = h_{n-1}$ if a cascade occurs at time n. Secondly, the m.p. has two classes of recurrent states, namely: $C_1 = \{(\{A,A\},2),(\{B,A\},2),(\{A,B\},2)\}$ and $C_2 = \{(\{B,B\},-2),(\{B,A\},-2),(\{A,B\},-2)\}$. It can be shown that the m.p. eventually gets absorbed into one of these classes w.p. 1. Once absorbed into class C_1 (C_2), all subsequent agents cascade resulting in an infinite sequence of actions $AABAAB\ldots$ ($BBABBA\ldots$). Later, we show that for other values of k, herding to a fixed action could also occur.

Property 1: The m.p. $\{(H_n^{(m)}, h_n)\}$, for any $k > 1$, gets absorbed almost surely in either C_1 or C_2. Once absorbed, a cascade lasts forever, but may not result in agents herding to a fixed action.

Now, if $(V_A, V_B) = (H, L)$, then the average cost (per agent) for agents absorbed into C_1 is $(2+2+k)/3$ and for agents absorbed into C_2 is $(2k+2k+1)/3$. Thus, clearly absorption to C_1 causes average costs to be lower than for C_2 when $(V_A, V_B) = (H, L)$.

Property 2: Consider $(V_A, V_B) = (H, L)$ and any $k > 1$, then the cascade to C_1 incurs a lower average cost per agent compared to a C_2. We hence refer to C_1 (C_2) as the *Right* (*Wrong*) cascade when $(V_A, V_B) = (H, L)$. Vice-versa if $(V_A, V_B) = (L, H)$.

Lastly, note that a cascade to C_1 (C_2) occurs if and only if $h_n = 2$ (-2). Now, let T_i, $i \in \mathbb{Z}$ denote the group of all states which have $h_n = i$. Then, any transition from T_i to T_j, $i \neq j$, occurs with the same probability for all $s \in T_i$. The 2-D m.p. is thereby said to be *lumpable* with respect to the partition $\{T_i\}_{i\in\mathbb{Z}}$, and the lumped chain is a 1-D m.p. as in Fig. 3, with $M = 2$ (see Thm. 6.3.2 in [21]). As a result, the probability of absorption of the 2-D m.p. to a C_1 (or C_2) cascade starting from state $(\{\},0)$ is equal to the probability of absorption to M (or $-M$) starting from state 0 of the 1-D m.p.

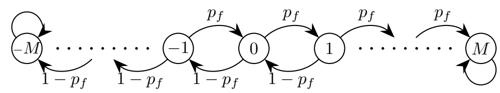

Fig. 3: Transition diagram of a 1-D r.w. with absorption states: $\pm M$.

with $M = 2$. Now, for the 2-D m.p., let $P_i^{(V_A,V_B)}$ denote the absorption probability to C_i, given (V_A, V_B). It follows from symmetry of the 2-D m.p. that $P_{\text{R-cas}} = P_1^{(H,L)} = P_2^{(L,H)} = p^M/[p^M + (1-p)^M]$. Here, $P_{\text{R-cas}}$ refers to the unconditional probability of absorption to the Right cascade.

A. Changes in the absorption states as k increases

Interestingly, at $k = k_0$ an agent in state $(\{A,A\},2)$ instead of cascading to B follows its private signal. As a result, more informative states, i.e., states with $|h_n| = 3,4,5$ are also possible. In particular, when an agent enters the state $(\{A,A\},5)$, it chooses to cascade to A, resulting in absorption to the state itself. Similarly, $(\{B,B\},-5)$ also becomes an absorption state. Therefore, in contrast to the earlier case of $k < k_0$, here the m.p. gets absorbed w.p. 1 into one of the two recurrent classes: $C_1 = \{(\{A,A\},5)\}$ and $C_2 = \{(\{B,B\},-5)\}$. Absorption into C_1 (C_2) corresponds to the infinite sequence of cascading actions $AAA\ldots$ ($BBB\ldots$). Further, the cascade probabilities can be obtained from the 1-D m.p. in Fig. 3, with $M = 5$. Similarly, it can be shown that for values of k in $(k_0, k_1]$, $(k_1, k_2]$ and (k_2, ∞), the corresponding absorption states $\pm M$ for the equivalent 1-D m.p. are ± 4, ± 3 and ± 2, respectively (see Table I).

The above discussion outlines how for a fixed p, the absorption states $\pm M$ vary with k. Alternatively, consider a fixed k, say $\hat{k}$, and varying p. Note that for any p, thresholds $\{k_r\}_{r=0}^2$, which are defined for state $(\{A,A\},2)$ are always greater than 3 (Remark 1). Thus, if $\hat{k} \leq 3$, then the absorption states $\pm M$ for the equivalent 1-D m.p. are ± 2 (see Table I) for all values of p. On the other hand, if $\hat{k} > 3$, then there exist an increasing sequence of thresholds $\{p_r\}_{r=0}^2$ such that when $p = p_r$, $k_r = \hat{k}$. Here $p_r := \left(\delta^{\frac{-1}{(3-r)}} + 1\right)^{-1}$, where $\delta = (3\hat{k}-1)/(\hat{k}-3)$ and $r \in \{0,1,2\}$. This implies that for values of p in $(0.5, p_0)$, p_0, $(p_0, p_1]$, $(p_1, p_2]$, $(p_2, 1)$, the values for $\pm M$ are ± 2, ± 5, ± 4, ± 3 and ± 2, respectively. This is depicted in Fig. 4, where $P_{\text{R-cas}}$ is plotted as a function of p for two cases: $k \leq 3$ and for a fixed value of $\hat{k} = 10$. Note the discontinuities in $P_{\text{R-cas}}$ at $\{p_r\}_{r=0}^2$ corresponding to abrupt changes in $\pm M$.

B. Learning the cost-optimal sequence of actions

Consider the realization $(V_A, V_B) = (H, L)$. If all agents knew the true quality a priori, then for $k \leq 3$ it can be shown that their optimal actions would result in the sequence $AABAAB\ldots$. Whereas, for $k > 3$, the optimal action sequence would be $AAA\ldots$ (Remark 2). When agents sequentially learn, if absorption to class C_1 gives the above mentioned optimal action sequence, then in this respect, class C_1 is not only the "Right" cascade but more importantly, it is the *Optimal* cascade. This is indeed the case for all k except $k \in (3, k_0)$, as shown in Table I. For $k \in (3, k_0)$, the sequentially arriving agents never *learn* the optimal action sequence. Whereas, for all $k \notin (3, k_0)$, agents learn the optimal sequence w.p. $P_{\text{R-cas}}$. Alternatively, for a fixed $k = \hat{k}$ and varying p, if $\hat{k} \leq 3$ then learning happens (w.p. $P_{\text{R-cas}}$) for any p. However, if $\hat{k} > 3$, then for $p \in (0.5, p_0)$ agents never learn the optimal sequence. This is indicated in red in Fig. 4 for the case $\hat{k} = 10$.

		Cascade pattern		Optimal sequence if (V_A, V_B) known apriori	
k	$\pm M$	Class C_1	Class C_2	$(V_A, V_B) = (H, L)$	$(V_A, V_B) = (L, H)$
$[1, 3]$	± 2	AAB	BBA	AAB	BBA
$(3, k_0)$	± 2	AAB	BBA	AAA	BBB
k_0	± 5	AAA	BBB	AAA	BBB
$(k_0, k_1]$	± 4	AAA	BBB	AAA	BBB
$(k_1, k_2]$	± 3	AAA	BBB	AAA	BBB
(k_2, ∞)	± 2	AAA	BBB	AAA	BBB

TABLE I: Comparing cascade pattern with optimal action sequence for different values of quality factor k. Thresholds $[k_0, k_1, k_2]$ defined for state $(\{A, A\}, 2)$ as per (4). Here, $3 < k_0 < k_1 < k_2$ (Remark 1).

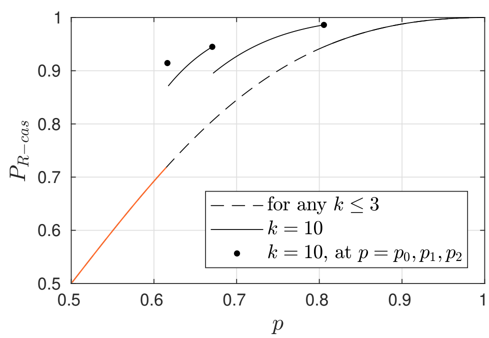

Fig. 4: Probability of Right cascade versus private signal quality for the indicated values of k. For any $k > 3$, when $p \in (0.5, p_0)$ (indicated in red for $k = 10$), the Right cascade is not optimal.

V. Generalization to even values of m

In this section, we consider the service times m to be an even integer value. Recall from Section III the 2-D m.p. $\{(H_n^{(m)}, h_n)\}$, where h_n is defined as per Lemma 2 and $H_n^{(m)}$ is the history of the m most recent actions till time n. Let $\mathcal{S} = \{A, B\}^m$ denote the space of all possible m-length histories $H_n^{(m)}$. Histories of length less than m occur only over the first $m-1$ time slots and are transient states. As in the case of $m = 2$, these transient states do not affect the probability of absorption and hence will be ignored for the following discussion. Now, for any $s \in \mathcal{S}$, let $m_A(s)$ and $m_B(s)$ denote the prior occupants in service A and B, respectively. Let $\mathcal{S}$ be partitioned into sets $\mathcal{S}_A$, $\mathcal{S}_B$ and $\mathcal{S}_{AB}$ defined as: $\mathcal{S}_A := \{s : m_A(s) > m_B(s)\}$, $\mathcal{S}_B := \{s : m_A(s) < m_B(s)\}$ and $\mathcal{S}_{AB} := \{s : m_A(s) = m_B(s)\}$. Consider the m.p. $\{(s_n, h_n)\}$ with the initial state $(s_0, 0)$ such that $s_0 \in \mathcal{S}_A$ ($s_0 \in \mathcal{S}_B$). Then, it can be shown using Lemma 3 that if the m.p. does not exit the $h = 0$ axis, then the successive sequence of states $\{(s_n, 0)\}$ will have an increasing number of B's (A's). Thus, given that the m.p. starts in state $(\{\}, 0)$, at some point, it either transitions from $h = 0$ axis to $h = 1$ ($h = -1$) axis w.p. p_f $(1 - p_f)$ or remains on the $h = 0$ axis and eventually enters a state $(s, 0)$ where $s \in \mathcal{S}_{AB}$. Now in this state, the agent always follows its private signal (Lemma 5) based on which the m.p. shifts to either the $h = 1$ or $h = -1$ axis. Thus, on the whole, the m.p.transitions from the $h = 0$ axis to the axes $h = 1$ and $h = -1$ w.p. p_f and $1 - p_f$ respectively. Similar arguments show that this is in fact true for all axes $h \leq |M - 1|$, for a certain M defined in Property 4. However, this is not true if m is odd valued, as the set $\mathcal{S}_{AB}$, in this case, is empty. For odd values of m, recurrent classes may exist (absorption may occur) along the $h = 0, 1, -1$ axes. As a result, the effect of service congestion dominates and thereby hinders the process of agents' learning of (V_A, V_B). For even m, the m.p. exhibits Properties 1 and 2 where the two recurrent classes C_1 and C_2 exist along the $h = M$ and $h = -M$ axis respectively. From above arguments, it follows that for even m, the 2-D m.p. is lumpable with respect to the partiton $\{T_i\}_{i \in \mathbb{Z}}$ as defined in Section IV.

Property 3: For any even-valued service time m, the probability of absorption of the 2-D m.p. to a C_1 (or C_2) cascade starting from state $(\{\}, 0)$ is equal to the probability of absorption to M (or $-M$) starting from state 0 of the 1-D m.p. (Fig. 3), where $M \geq 2$.

The absorption states $\pm M$ are characterized by following property, whose proof is omitted due to space considerations.

Property 4: The value of M referring to the absorption states $\pm M$ equals the smallest positive integer h such that for all states (s, h), $s \in \mathcal{S}_A$, the agent cascades (to A or B).

Now, recall that for $k > m + 1$, the optimal action sequence when $(V_A, V_B) = (H, L)$ is $AAA \ldots$ (Remark 2). Then, for the 2-D m.p. considered here, it follows that the Right cascade (learnt w.p. $P_{\text{R-cas}}$) is optimal only if for every state (s, M), $s \in \mathcal{S}_A$, the agent cascades to A (and never to B). From simulations we observe that there exists a threshold k_{th} such that for $k \in (m + 1, k_{th})$, the above condition is not satisfied, i.e., agents never learn the optimal sequence. Whereas, for $k \geq k_{th}$ agents learn the optimal sequence w.p. $P_{\text{R-cas}}$. Figure 5 depicts this for $m = 10$, where we find $k_{th} = 30.7$.

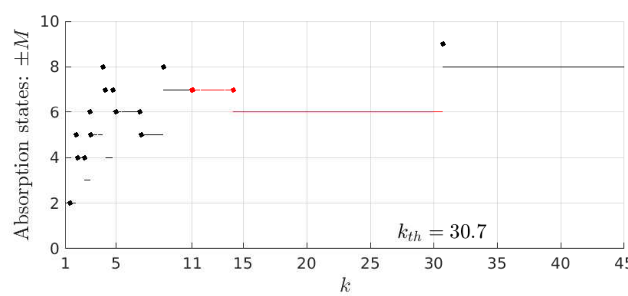

Fig. 5: Absorption state index M versus quality factor k for $m = 10$ and $p = 0.6$. For $k \in (11, k_{th})$ (indicated in red), the Right cascade is not optimal.

VI. Conclusions

We considered a simple model for observational learning with negative externalities and showed that such externalities can lead to very different learning behavior compared to models without externalities. In particular, there are cases in which the optimal sequence of actions is never learned and better signal quality may result in worse learning outcomes.

REFERENCES

[1] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.

[2] A. V. Banerjee, "A simple model of herd behavior," *The quarterly journal of economics*, vol. 107, no. 3, pp. 797–817, 1992.

[3] I. Welch, "Sequential sales, learning, and cascades," *The Journal of finance*, vol. 47, no. 2, pp. 695–732, 1992.

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

[5] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

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

[7] T. Le, V. Subramanian, and R. Berry, "Quantifying the utility of noisy reviews in stopping information cascades," in *IEEE CDC Proceeding*, 2016.

[8] 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.

[9] D. Sgroi, "Optimizing information in the herd: Guinea pigs, profits, and welfare," *Games and Economic Behavior*, vol. 39, no. 1, pp. 137–166, 2002.

[10] G. Schoenebeck, S. Su, and V. G. Subramanian, "Social learning with questions," *CoRR*, vol. abs/1811.00226, 2018. [Online]. Available: http://arxiv.org/abs/1811.00226

[11] P. Poojary and R. Berry, "Observational learning with fake agents," in *2020 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2020, pp. 1373–1378.

[12] Y. Peres, M. Z. Rácz, A. Sly, and I. Stuhl, "How fragile are information cascades?" *arXiv preprint arXiv:1711.04024*, 2017.

[13] T. N. Le, V. G. Subramanian, and R. A. Berry, "Bayesian learning with random arrivals," in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 926–930.

[14] D. Seo, R. K. Raman, and L. R. Varshney, "Social learning with beliefs in a parallel network," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1265–1270.

[15] T. M. Cover, "Hypothesis testing with finite statistics," *The Annals of Mathematical Statistics*, vol. 40, no. 3, pp. 828–835, 1969.

[16] M. E. Hellman and T. M. Cover, "Learning with finite memory," *The Annals of Mathematical Statistics*, pp. 765–782, 1970.

[17] Y.-H. Yang, Y. Chen, C. Jiang, C.-Y. Wang, and K. R. Liu, "Wireless access network selection game with negative network externality," *IEEE Transactions on Wireless Communications*, vol. 12, no. 10, pp. 5048–5060, 2013.

[18] T. Nguyen, H. Zhou, R. A. Berry, M. L. Honig, and R. Vohra, "The cost of free spectrum," *Operations Research*, vol. 64, no. 6, pp. 1217–1229, 2016.

[19] B. Badia, R. A. Berry, and E. Wei, "Investment in EV charging spots for parking," *IEEE Transactions on Network Science and Engineering*, vol. 7, no. 2, pp. 650–661, 2019.

[20] R. Arnott and E. Inci, "An integrated model of downtown parking and traffic congestion," *Journal of Urban Economics*, vol. 60, no. 3, pp. 418–442, 2006.

[21] J. G. Kemeny and J. L. Snell, *Finite Markov chains: with a new appendix" Generalization of a fundamental matrix"*. Springer, 1983.