Learning from Randomly Arriving Agents

Tho Ngoc Le
Dept. of EECS
Northwestern University
Evanston, IL 60208
thole2012@u.northwestern.edu

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 who do not choose 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 incorrect cascades may occur and that the probability of such cascades is not monotonic in the arrival probability of a user. Moreover, if the agents' private signals are weak, wrong cascades are more likely to happen than correct cascades.**

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. [1], [2]). 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 [1] and [2], 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 [1] and [2], 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 [3] or allowing agents to observe only a subset of past actions [4]. In previous work ([5], [6], [7]), we considered a variation in which there were random errors in the observations of the *actions* of other agents. 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, which 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. Specifically, we consider a discrete-time model in which at each time-slot an agent either 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 an arriving agent chooses to buy, we assume that 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 always 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 [1], [2], [5], [6], [7], 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 a independent, identically distributed binary signal about this state of the world. 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.

Both this work and our prior work in [5], [6], [7] considered imperfections in the observations of the agents. In [5], [6], [7] 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

This work was supported by the NSF under grant IIS-1538827.

an empty slot decreases as the arrival probability decreases.

For this model, we show that as in [1], [2], [5], [6], [7], 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. 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.

REFERENCES

[1] 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.

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

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

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

[5] T. Le, V. Subramanian, R. Berry, "The Value of Noise for Informational Cascades," *IEEE International Symposium on Information Theory (ISIT)*, pp. 1101-1105, Honolulu, HI, July 2014.

[6] T. Le, V. Subramanian, R. Berry, "The Impact of Observation and Action Errors on Informational Cascades," *IEEE 53rd Annual Conference on Decision and Control (CDC)*, pp. 1917-1922, Los Angeles, CA, Dec. 2014.

[7] T. Le, V. Subramanian, R. Berry, "Information Cascades With Noise," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 239-251, June 2017.