Statement of Research (2021-2022)

Vartika Singh, IEOR, IIT Bombay

In the past year, I have been working on two problems under the supervision of Professor Veeraruna Kavitha. The first problem analyses the 'optimal' leader strategies (that influence the influencers and decide the dynamic vaccine supply) to eradicate an ongoing pandemic, using a layered game-theoretic framework in scenarios with partial information. The second problem analyses the optimal advertising strategies for content providers competing to win over hidden potential customers in the network under partial information. I presented this second problem in December 2021 review, hence providing only a short note on it. I am also starting a third problem in which we are trying to derive an alternate algorithm to solve the Markov Decision process with constraints.

1. Stochastic Game among influencers, leader and public ([1])

Herd immunity is often the way to curb any outbreak and can be achieved by vaccination of a large proportion of the population. For example, WHO estimates that about 95% of the population needs to be vaccinated to achieve herd immunity against Measles. However, vaccination of such a large proportion is possible only if the population is eager for the vaccine, which usually may not be the case. A recent article [2] describes the role of influencers (celebrities, medical professionals, etc.) in winning over the vaccine skeptics. We take motivation from the same and analyze the role of influencers in populations vaccination response and how the leader can harness this influence to eradicate the disease using an elaborate game-theoretic framework.

Problem formulation

We have a three-level Stackelberg game, with population dynamics modeled by an evolutionary game-theoretic framework on the bottom level; influencers' behavior captured by a stochastic game on the second level, which depends upon the equilibrium achieved by the bottom level; and leader on the top has an optimization problem based on the equilibria achieved by the lower levels.

To study the outcome of the bottom level evolutionary game, for each strategy profile of the influencers and the leader, we extend the framework of [3]; this work discusses equilibrium strategies namely evolutionary stable strategy (stable against static mutation), and the corresponding equilibrium state for a given vaccination response of the population. The extension is to include the impact of dynamic vaccine supplies and the influencer/leader strategies. The outcome, i.e., the evolutionary stable state describes the proportions of infected and vaccinated population at equilibrium. As more influencers vaccinate, the perceived cost of side-effects of vaccines reduces, and the perceived cost of insecurity upon missing the vaccine increases, for any individual in the population. If the number of vaccinated influencers is large enough, the population may vaccinate eagerly, eradicating the disease.

The leader offers incentives to some influencers for vaccination which may include monetary benefits, goodwill, etc. The leader announces the incentive scheme and provides T days to the influencers to make a vaccination decision from the day of the introduction of vaccines. This leads to a *stochastic game* among the influencers. On any day t, any susceptible influencers can either vaccinate with some probability or procrastinate the vaccination decision to the next day or choose to never vaccinate. The influencers make the vaccination decision after considering various factors, e.g., the incentives, side-effects of vaccine and eventual risk of infection based upon anticipated population dynamics at the bottom level. The influencers have only partial information about the side effects of the vaccines; every day, they receive new (random) data about the side effects (from the experiences of people getting vaccinated worldwide); this improves their estimates of cost of side-effects. Since the outcome of the stochastic game also depends upon the incentives offered by the leader, and on the corresponding response of the population, i.e., the outcome of the evolutionary game, and since the leader aims to minimize the expected value of incentives while keeping the probability of eradication above $1-\delta$ (for a desired $\delta < 1$), we have a three-layered *Stackelberg game*.

Results

- Our first result states that there indeed exists some n, such that eradication of the disease is an evolutionary stable state if and only if at least n influencers are vaccinated.
- We prove that the "wait-and-watch" is the most-preferred equilibrium among the influencers, as it results in the highest utilities. In this equilibrium, the influencers wait till the last day to make a vaccination decision which is dependent upon the then estimate (best possible estimate) of the vaccine's side effects.
- Leaders' optimization problem reduces to one-dimensional stochastic optimization; the stochastic nature results from the random evolution of the vaccines' side effects and the randomness in the number of influencers vaccinated at the Nash equilibrium of the stochastic game.
- We derive an equation whose solution provides the optimizer for the leader, i.e., optimal incentive to be provided by the leader to maintain the eradication probability above $1 - \delta$.

Discussions and observations

The minimum number of influencers who must vaccinate for eradication to be evolutionary stable depends upon various factors, such as disease severity, vaccines' side effects, etc. This number is larger if the vaccine has severe side effects and is smaller if the infection is severe. Further, we show that the probability with which influencers vaccinate in the "wait-and-watch" equilibrium increases with an increase in incentives and decreases with an increase in vaccines' side-effects. Surprisingly, this probability decreases with the number of influencers, this possibly could be attributed to free-riding behavior.

We solve the leaders' optimization problem numerically using stochastic approximation techniques and have the following observations. The optimal cost for the leader increases with an increase in the initial cost of side effects of vaccines or if the vaccines have a wider range (variance) of side effects. This is because the leader needs to provide higher incentives to encourage the influencers and eventually the population to vaccinate. Further, if the disease is severe, the optimal cost of the leader is lesser because the severity of the disease itself motivates the influencers and the public to vaccinate. Surprisingly, the leader incurs a higher cost when more number of influencers are available. One would anticipate if there were more influencers, the leader would be able to get the desired result with lesser incentives, but we observe the opposite. This is probably because the influencers start to become free riders, and the leader needs to provide higher incentives.

2. Partial information games and competitive advertising ([5])

Social networks have become a useful platform for content providers (CPs) to propagate their content to potential customers hidden in the network. CPs compete among themselves to acquire customers through advertising. The customers are limited in number, and CPs have no information if the opponents (other CPs) have already taken some or all of the customers; they are only aware of their contact status with the customers. Further, any contact is not guaranteed success; instead, the customers usually choose one among all the contacted CPs. Many times this choice is revealed to the CPs after some delay. We are interested in identifying the equilibrium advertising policies for the CPs in such a system with partial information.

Problem formulation and solution approach

The system evolves as a controlled Markov jump process, where jumps (contacts to potential customers) are controlled by the actions of the CPs. However, any CP can observe only a few jumps, to be precise the CPs can observe only their contacts. If all the jumps were known, one could use closed-loop policies, if none of the jumps were known, one would rely upon open-loop policies, but in this case, since the CPs observe some of the jumps, we propose "Open Loop Policies till Information Update (OLC-IU)". Further, the utility of any CP depends upon the system-state and hence the action of all other CPs. This leads to a partial information game with asymmetric and non-classical information with an infinite-dimensional state and action space.

We propose a general problem formulation applicable to this problem and a variety of other applications that satisfy the required assumptions, e.g., [4]. We start our analysis by deriving the best response against any strategy profile of the opponents using tools of the Markov Decision Process (MDP) and optimal control theory. Every stage of the dynamic programming equations of the MDP is solved using optimal control tools. We proved that the special type of threshold policies are the best response against any strategy profile of the opponents. We further proved that the thresholds need not depend upon the time of previous state update (e.g., previous contact instance for any CP). Hence, any optimal policy in the infinite-dimensional game can be described by a finite vector of thresholds. This implies that one can find a Nash Equilibrium (if it exists) among a reduced space of threshold strategies, hence reducing the games' complexity to a great extent.

Results and Discussion

For the advertising problem, we prove the existence of Nash Equilibrium (NE) among threshold policies and derive a finite dimensional fixed point equation, whose solution provides the NE. The thresholds defining the NE represent the time till which CPs attempts to contact any customer with full potential; the k-th component of the threshold vector represents the time till which any CP will attempt for its k-th contact. We also prove that in any best response and hence in the equilibrium, the thresholds are monotone-decreasing; when more customers are left to contact, any CP will advertise for a longer time than with fewer remaining customers. Hence, k-th contact may happen after the $(k+1)$-th threshold, in which case CP will no longer attempt for $(k+1)$-th contact. Further, the NE need not be unique.

We also provide an algorithm based on fictitious play and stochastic approximation techniques to compute the NE numerically. We derive more insights into the advertising problem with the help of numerical examples in [5]. With the remaining parameters fixed, the thresholds at NE decrease if the number of players increase and increase when the potential customers increase, in most cases. We also consider the case when the number of customers in the network are not fixed but random; the threshold policies are optimal again.

Future work - A new outlook at Constrained Markov Decision Process

It is well known that the solution of the dynamic programming equations represent the value function (optimal objective value as a function of state) of any MDP problem without constraints. To the best of our knowledge, such fixed point equations are not known when one considers MDP problem with constraints. We aim to fill this gap; the idea is to derive a fixed point equation whose solution provides value function and optimal policy for the constrained MDP problem. We further aim to design a variant of the value-iteration/Q-learning algorithm to estimate the optimal policy for the problem with constraints. Our approach is to construct a small linear programming (LP) problem at every iterate of the value-iteration/Q-learning algorithm; the solution of the LP replaces the greedy policy. We have some initial simulation results in this direction for a problem with two-states[1]. The solution derived from our algorithm matches with the solution derived using standard LP based approach to solve constrained problems. The small LP considered in every value-iteration/Q-learning step will also be instrumental in deriving the desired fixed point equation for the constrained problem.

References

[1] Singh, Vartika, and Veeraruna Kavitha. Stochastic Game among influencers, leader and public. Manuscript under preparation, to be submitted to Dynamic Games and Application (2022).

[2] Mintz, O. "How Influencers, Celebrities, and FOMO Can Win Over Vaccine Skeptics." Harvard Business School Working Knowledge (2021).

[3] Singh, Vartika, Khushboo Agarwal, Shubham and Veeraruna Kavitha. "Evolutionary Vaccination Games with premature vaccines to combat ongoing deadly pandemic." EAI International Conference on Performance Evaluation Methodologies and Tools. Springer, Cham, 2021.

[4] Singh, Vartika, and Veeraruna Kavitha. "Asymmetric Information Acquisition Games." 2020 59th IEEE Conference on Decision and Control (CDC). IEEE, 2020.

[5] Singh, Vartika, and Veeraruna Kavitha. "Partial information games and competitive advertising." Manuscript submitted to Mathematics of Operations Research (2022). Link

---

[1] joint work with a masters student