23rd International Symposium on Mathematical Theory of Networks and Systems
Hong Kong University of Science and Technology, Hong Kong, July 16-20, 2018

On Influencing Opinion Dynamics over Finite Time Horizons

Bhumesh Kumar[1], Neeraja Sahasrabudhe[2], and Sharayu Moharir[1]

[1]Indian Institute of Technology Bombay
[2]Indian Institute of Science Education and Research Mohali

***Abstract*— In this work, we focus on strategies to influence the opinion dynamics of a well-connected society. We propose a generalization of the popular voter model. This variant of the voter model can model a wide range of individuals including strong-willed individuals whose opinion evolution is independent of their neighbors as well as conformist individuals who tend to adopt the opinion of the majority.**

Motivated by political campaigns which aim to influence opinion dynamics by the end of a fixed deadline, we focus on influencing strategies for finite time horizons. We characterize the nature of the optimal influencing strategies as a function of the nature of individuals forming the society. Using this, we show that for a society consisting of strong-willed individuals, the optimal strategy is to influence towards the end of the finite time horizon, whereas, for a society consisting of individuals who are affected by their peers, it could be optimal to influence in the initial phase of the finite time horizon.

I. Introduction

Opinion dynamics have been a subject of study in various fields including sociology, philosophy, mathematics, and physics for a very long time [1]. In this work, we focus on a variant of a widely studied binary opinion dynamics model known as the voter model [2], [3]. In the voter model, society is modeled using a graph where each individual is a node and edges represent links between these individuals. Each individual holds one of two possible opinions, e.g., pro-government and anti-government. The opinions of individuals evolve over time. Assuming time is slotted, one individual is chosen uniformly at random at the beginning of each time-slot. This individual then adopts the opinion of one of its neighbors, chosen uniformly at random. The voter model is a useful framework to study opinion dynamics and the spread of competing epidemics. Variants and generalizations of the voter models have also been studied [4], [5].

Our model differs from the voter model in two key ways. Firstly, in each time-slot, the opinion of the selected individual evolves according to a general function of the opinion of its neighbors. This modification to the voter model allows us to model a variety of natures of individuals in society. For example, we can model stubborn individuals by making the opinion evolution of the selected individual independent of the opinions of its neighbors and/or any other external influence. Similarly, we can model conformist individuals by forcing the selected individual to adopt the opinion of the majority. Secondly, we focus on the setting where the graph between the individuals is a complete graph. This is justified in the presence of social media platforms like Twitter and the abundance of publicly available poll results on most important issues.

Use of social networks and other media outlets for political campaigning and project marketing is on the rise. While opinions of individuals evolve organically over time, this evolution can be influenced by effective campaigning. Resource limitations like a fixed budget or limited manpower restrict the set of feasible influencing strategies and motivate the need to use the available resources efficiently.

In political campaigning, the goal is to influence as many individuals as possible by the end of a fixed deadline. Motivated by this, we focus on designing influencing strategies that maximize the number of individuals with a positive opinion at the end of a known and finite time horizon [6]. The optimal influencing strategy is one that maximizes the number of individuals with a favorable opinion at the end of this time horizon. In this work, our goal is to study how the nature of the optimal influencing strategy varies with the nature of individuals in society.

The key takeaway from this work can be summarized as follows. If the society consists of strong-willed individuals who are unaffected by the opinion of their peers but are susceptible to external influence, the optimal influencing strategy is to influence towards the end of the finite time horizon. Contrary to this, if individuals are heavily influenced by their peers in addition to being susceptible to external influence, in some cases, it is optimal to influence at the beginning of the time horizon. Intuitively, this is because increasing the fraction of individuals with a favorable opinion at the beginning of the finite time horizon has a cascading effect on the opinions of the society as a whole.

A. Related Work

Closest to our setting, [5] focuses on the voter model and generalizes it to include external influences. The key takeaway is that the effect of external influences overpowers node-to-node interactions in driving the network to consensus in the long term. In [4], the focus is on studying the effect of stubborn agents, i.e., agents who influence others but do not change their opinion, on the opinion dynamics of the network. The authors also study the problem of optimal placement of these stubborn agents to maximize the effect on the network.

Designing optimal influencing strategies has been the subject of study in many works including [6]–[9]. Refer to [9] for a detailed survey of various works in this domain.

Unlike our work, most of these works focus on the infinite time horizon setting. In [6], the focus is on characterizing the optimal influence strategy to maximize the spread of an epidemic in a network. In [7], the focus is on minimizing the cost incurred by the influencer to reach a fixed fraction of nodes in the network. In [8], the authors propose a general model of influence propagation called the decreasing cascade model and analyze its performance with respect to maximizing the spread of an idea. In [9], the focus is on designing optimal advertising strategies in the presence of multiple advertising channels.

A related body of work focuses on preventing the spread of disease/viruses in networks (refer to [10]–[12] and the references therein). Our work differs from this body of work since we focus on strategies to increase the spread of the favorable opinion in the network.

In [13]–[15], the focus is on analyzing the performance of various rumor spreading strategies. These works do not focus on finding the optimal strategies for information spread.

B. Organization

The rest of this paper is organized as follows. In Section II, we formally define our opinion dynamics model. In Section III, we discuss some preliminary results on stochastic approximation which are used in the subsequent analysis. In Section IV, we discuss our results for the setting where individuals are unaffected by the opinion of their peers but are susceptible to external influence. In Section V, we present our results for the setting where individuals are influenced by their peers in addition to being susceptible to external influence. We conclude the paper in Section VI. Some additional results are presented in the appendix.

II. Setting

A. System Evolution

Consider a fixed population of N people, where each individual has either a positive or a negative opinion on the topic of interest. We model this as an urn with balls of two colors (representing positive/negative opinions): Green and Red. Time is slotted, and in each time-slot, the campaigner decides whether to influence the opinion dynamics or not. One ball is chosen uniformly at random in each time-slot. In each time-slot, the system evolves as follows:

a. If the ball chosen in time-slot t is a red ball, and the campaigner decides not to influence in this time-slot, the red ball turns green with probability p_t.
b. If the ball chosen in time-slot t is a green ball, and the campaigner decides not to influence in this time-slot, the green ball turns red with probability q_t.
c. If the ball chosen in time-slot t is a red ball, and the campaigner decides to influence in this time-slot, the red ball turns green with probability $\tilde{p}_t$.
d. If the ball chosen in time-slot t is a green ball, and the campaigner decides not to influence in this time-slot, the green ball turns red with probability $\tilde{q}_t$.

Formally, let $r(t)$ denote the number of red balls in the urn at time t, $\rho_r(t) = r(t)/N$ denote the fraction of red balls at time t for all $t \in [0, T]$, and $\chi(t)$ denote the change in the number of red balls in the urn at time $t+1$. Then, we have that,

$$r(t+1) = r(t) + \chi(t+1), \tag{1}$$

$$\text{where, } \chi(t+1) = \begin{cases} -1 & \text{w.p. } \rho_r(t)\widehat{p}_t, \\ 1 & \text{w.p. } (1-\rho_r(t))\widehat{q}_t, \\ 0 & \text{otherwise,} \end{cases}$$

$$\widehat{p}_t = \begin{cases} p_t & \text{if campaigner does not influence in time-slot } t, \\ \tilde{p}_t & \text{if campaigner influences in time-slot } t, \end{cases}$$

$$\widehat{q}_t = \begin{cases} q_t & \text{if campaigner does not influence in time-slot } t, \\ \tilde{q}_t & \text{if campaigner influences in time-slot } t. \end{cases}$$

B. Goal

Given a finite influencing budget of bT time-slots, where $0 \leq b \leq 1$, the goal is to design effecient influencing strategies which maximize the number/fraction of green balls at the end of time-slot T.

III. Preliminaries

Let $\mathcal{F}_t$ denote the σ-field generated by random variables $\{\chi(s)\}_{1 \leq s \leq t}$. The dynamics of the urn are governed by the random variables $\{\chi(t)\}_{t \geq 1}$ according to the random reinforcement scheme defined in (1). The dynamics in (1) can be re-written as follows:

$$\begin{aligned} r(t+1) = r(t) &+ \mathbb{E}\big[\chi(t+1)|\mathcal{F}_t\big] \\ &+ \big(\chi(t+1) - \mathbb{E}\big[\chi(t+1)|\mathcal{F}_t\big]\big), \end{aligned} \tag{2}$$

where, $\big(\chi(t+1) - \mathbb{E}\big[\chi(t+1)|\mathcal{F}_t\big]\big)$ is zero-mean Martingale difference. We analyze the difference equation above using stochastic approximation. The theory of stochastic approximation states that the solution of a difference equation of the form $x_{n+1} = x_n + a(n)[h(x_n) + M_{n+1}], n \geq 0$, where h is Lipschitz, $a(n)$s are decreasing step-sizes and M_n is a square-integrable Martingale difference sequence, are close to the solutions of the ODE $\dot{x}(n) = h(x(n))$ as $n \to \infty$. Further, for sufficiently large n, the solutions of the difference equation remain "close" to that of the ODE with high probability. Explicit bounds on the probability are provided in Corollary 14 of Chapter 4 in [16]. These results were further extended in [17] and [18]. These results allow us to analyze the solution of a suitable ODE to obtain an optimal strategy for the influencing agents in the opinion dynamics defined in (1).

As mentioned above, in (2), $\chi(t) - \mathbb{E}\big[\chi(t)|\mathcal{F}_{t-1}\big]$ is a zero-mean Martingale difference and $\mathbb{E}\big[\chi(t)|\mathcal{F}_{t-1}\big]$ is a Lipschitz function. The corresponding ODE with constant step size (See [19] for extensions of stochastic approximation results to constant step-size) is given by:

$$\dot{r}(t) = \mathbb{E}\big[\chi(t)|\mathcal{F}_{t-1}\big] = -r(t)\left[\frac{(\widehat{p}_t + \widehat{q}_t)}{N}\right] + \widehat{q}_t, \tag{3}$$

where N is the total number of balls in the urn. In other words,

$$\dot{\rho_r}(t) = \frac{\widehat{q}_t}{N} - \rho_r(t)\frac{\widehat{p}_t + \widehat{q}_t}{N}. \tag{4}$$

It is worth noting that a larger step-size, scaled by some $m > 1$, results in a faster convergence of the ODE, whereas a smaller step-size would mean that the ODE trajectory will track the difference equation (1) better.

IV. Strong-Willed Population

In this section, we focus on a population of strong-willed individuals whose opinion evolves independent of the opinion of the remaining population but is susceptible to external influence.

A. Time-invariant Adamancy

We first focus on the case where the probability of an individual changing their opinion is independent of the time-slot index and only depends on whether the campaigner decides to influence in that time-slot or not. The next assumption formally characterizes the setting.

Assumption 1 (Strong-Willed Population with Time-invariant Adamancy)**.** *The system evolves according to* (1) *with* $p_t = p$, $q_t = q$, $\tilde{p}_t = \tilde{p}$, $\tilde{q}_t = \tilde{q}$, $\forall t$, *with* $\tilde{p} \geq p$ *and* $\tilde{q} \leq q$.

Remark 1. *We limit our interest to the case where* $\tilde{p} \geq p$ *and* $\tilde{q} \leq q$ *since the goal of influencing is to maximize the number of green balls. Given this, compared to the case without influence, influencing should make the event that a red ball turns green more likely and the event that a green ball turns red less likely.*

Definition 1 (Optimal Strategy)**.** *We call a strategy optimal if the influence according to that strategy results in a larger expected fraction of green balls at the end of time* T *than the expected fraction of green balls at the end of time* T *using any other influence strategy.*

Throughout the paper $\rho_r(t)$ denotes the solution of (4) at time t and $\rho_r(0)$ denotes the fraction of red balls at time 0.

Recall that the campaigner has a fixed influencing budget of bT time-slots for some $0 < b < 1$. We start by discussing two influence strategies: influence in the first bT time-slots and influence in the last bT time-slots. Recall that the influencing agent is interested in maximizing the number of green balls at the end of time-slot T.

Lemma 1. *Let* $\rho_r(T)|_{first}$ *and* $\rho_r(T)|_{last}$ *denote the expected fraction of red balls at the end of time-slot* T *under the influence in the first* bT *time-slots and influence in the last* bT *time-slots policies respectively. Under Assumption 1,*

$$\rho_r(T)|_{first} = \frac{1}{p+q}\left(q - e^{\frac{-(p+q)(1-b)T}{N}}\left(q - \frac{p+q}{\tilde{p}+\tilde{q}} \times \left(\tilde{q} - e^{\frac{-(\tilde{p}+\tilde{q})bT}{N}}\left(\tilde{q} - r(0)\frac{\tilde{p}+\tilde{q}}{N}\right)\right)\right)\right), \tag{5}$$

$$\rho_r(T)|_{last} = \frac{1}{\tilde{p}+\tilde{q}}\left(\tilde{q} - e^{\frac{-(\tilde{p}+\tilde{q})bT}{N}}\left(\tilde{q} - \frac{\tilde{p}+\tilde{q}}{p+q} \times \left(q - e^{\frac{-(p+q)(1-b)T}{N}}\left(q - r(0)\frac{p+q}{N}\right)\right)\right)\right). \tag{6}$$

Proof. Recall that the ODE corresponding to our system is given by:

$$\frac{d\rho_r(t)}{dt} = \frac{\widehat{q}_t}{N} - \frac{\widehat{p}_t + \widehat{q}_t}{N}\rho_r(t),$$

where $\widehat{p}_t$ and $\widehat{q}_t$ are as defined in Section II. Under Assumption 1 and the strategy to influence in the first bT time-slots, we get that for the influence period $[0, bT]$:

$$\rho_r(bT) = \frac{\tilde{q} - e^{\frac{-(\tilde{p}+\tilde{q})bT}{N}}\big(\tilde{q} - \rho_r(0)(\tilde{p}+\tilde{q})\big)}{\tilde{p}+\tilde{q}}.$$

Similarly, for the no-influence period $[bT, T]$, we get:

$$\rho_r(T) = \frac{q - e^{\frac{-(p+q)(1-b)T}{N}}\big(q - \rho_r(bT)(p+q)\big)}{p+q}.$$

Combining the two solutions, we get an expression for total number of red balls at the end of time horizon T when the influence strategy is to influence in the first bT time-slots, thus completing the proof.

The proof for the strategy to influence in the last bT time-slots follows on similar lines. □

Next, we use Lemma 1 to show that, under Assumption 1, the strategy to influence in the last bT time-slots outperforms the strategy to influence in the first bT time-slots.

Lemma 2. *Let* $\rho_r(T)|_{first}$ *and* $\rho_r(T)|_{last}$ *denote the fraction of red balls at the end of time* T *when the influence strategy is to influence in first* bT *slots and the last* bT *slots respectively. Under Assumption 1, we have that* $\rho_r(T)|_{first} \geq \rho_r(T)|_{last}$.

Proof. From the ODE solutions computed for both the strategies in (5) and (6), we get:

$$\rho_r(T)|_{\text{first}} - \rho_r(T)|_{\text{last}} = \left(\frac{q}{p+q} - \frac{\tilde{q}}{\tilde{p}+\tilde{q}}\right) \times \big(1 - e^{\frac{-(p+q)(1-b)T}{N}} - e^{\frac{-(\tilde{p}+\tilde{q})(b)T}{N}} + e^{\frac{-(\tilde{p}+\tilde{q})bT-(p+q)(1-b)T}{N}}\big). \tag{7}$$

Since $\tilde{p} \geq p$ and $\tilde{q} \leq q$, we have $\frac{\tilde{p}}{\tilde{q}} \geq \frac{p}{q}$. Hence the first term in (7) is always positive. The second term is of the form $f(x, y) = 1+e^{-x-y}-e^{-x}-e^{-y}$. Note that, $f(0, 0) = 0$ and, $f_x(x, y) = e^{-x}(1-e^{-y}) \geq 0$, $f_y(x, y) = e^{-y}(1-e^{-x}) \geq 0$. Hence, $f(x, y)$ is non-negative for $x, y \geq 0$. This completes the proof of the lemma. □

In Figures 1 and 2, we compare the performance of the two policies (influence in the first and last bT time-slots) for different values of N, T and b. We plot the values obtained via the ODE solutions (given by (5) and (6)) as well as results obtained by simulating the system described by Assumption 1. We observe that the ODE solutions accurately track the results obtained via simulations. In addition, the strategy of influencing the last bT outperforms the strategy to influence in the first bT time-slots.

Theorem 1 (Optimality of last bT influence strategy)**.** *Under Assumption 1 and the ODE dynamics described in* (4)*, the strategy to influence in the last* bT *time-slots, i.e., at the end of the finite time-horizon, is optimal.*

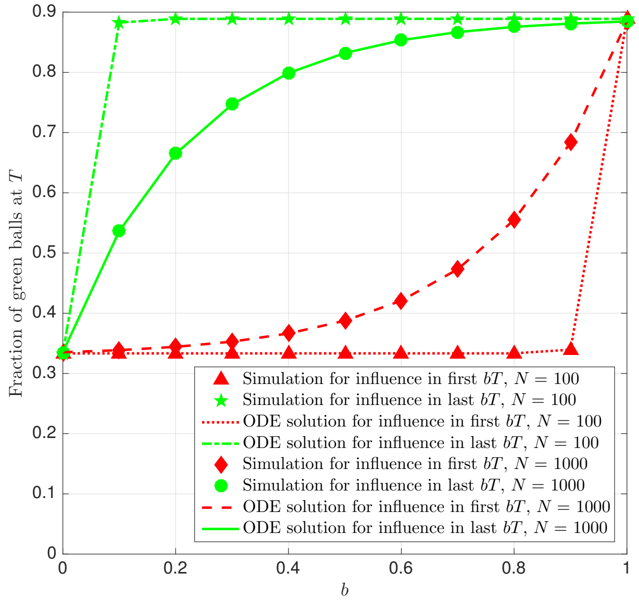

Fig. 1. Expected fraction of green balls at the end of time-slot T obtained by solving the ODE and simulating the system as a function of the influence budget b for two policies, namely, influence in the first bT timeslots and influence in the last bT time-slots. The system under consideration satisfies Assumption 1 with parameters $p = 0.3$, $q = 0.6$, $\tilde{p} = 0.8$, $\tilde{q} = 0.1$, $\rho(0) = 0.5$ and $T = 5000$. The ODE solutions accurately track the results obtained via simulations.

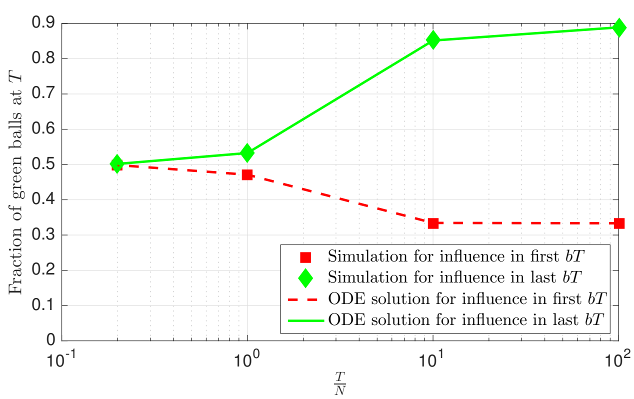

Fig. 2. Expected fraction of green balls at the end of time-slot T obtained by solving the ODE and simulating the system as a function of the influence budget b for two policies, namely, influence in the first bT timeslots and influence in the last bT time-slots. The system under consideration satisfies Assumption 1 with parameters $p = 0.3$, $q = 0.6$, $\tilde{p} = 0.8$, $\tilde{q} = 0.1$, $\rho(0) = 0.5$ and $T = 1000$. The ODE solutions accurately track the results obtained via simulations.

Proof. Assign values 1 and 0 to each time slot depending on whether the dynamics in that time slot is influenced (i.e., $(\tilde{p}, \tilde{q})$-dynamics) or not (i.e., (p, q)-dynamics). A binary string of length L will correspond to an influence strategy pattern over a time duration with L slots.

Suppose that there exists a ‘10’ sub-sequence (otherwise the configuration of the string is of last bT type). Then by Lemma 2, if we just consider the time window consisting of this ‘10’ sub-sequence, we could strictly improve the target ball count at the end of this time window by doing a local swap and getting ‘01’. In other words, Lemma 2 implies that in this time window influence in the last slot is better.

The ODE solution is a monotonically increasing function of the configuration count at the start of any time window over which the ODE is written. If we look at the composition of these increasing functions along the time horizon after the small window, we get an increasing function. Thus, the local swap also results in a global improvement i.e. larger number of green balls at the end of time T.

Finally, this ‘10’ to ‘01’ swapping operation is clearly a terminating (finite number of permutations) and non-oscillating (always decrease the decimal value of the string) algorithm. If we assume that some strategy $S(\neq$ *last bT* influence strategy) is better than last bT, it has to have a 10 sub-sequence and can therefore be improved by swapping operation and hence is not optimal. Therefore, last bT is the optimal strategy to achieve a larger fraction of green balls at the end of time T. $\square$

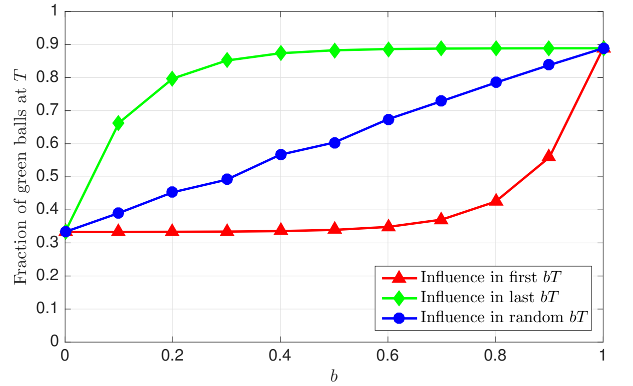

Fig. 3. Expected fraction of green balls at the end of time-slot T as a function of the influence budget b for three policies, namely, influence in the first bT timeslots, influence in the last bT time-slots, and influence in each time-slot with probability b. The system under consideration satisfies Assumption 1 with parameters $p = 0.3$, $q = 0.6$, $\tilde{p} = 0.8$, $\tilde{q} = 0.1$, $N = 500$, $\rho(0) = 0.5$ and $T = 5000$. The strategy of influencing the last bT outperforms the other two strategies for all values of b.

Note that the fraction of green balls is a Markov chain on a finite state space that mixes fast. It is clear from figures 1 and 2 that the ODE solution at time T is close to the solution of the simulated trajectory. In fact, the solutions of the difference equation or the stochastic approximation scheme track the solution trajectory of the corresponding differential equation (See figure 6 in Appendix). An explicit probabilistic bound for this in terms of T is obtained in the Appendix using concentration inequalities.

In Figure 3, we compare the performance of three policies (influence in the first bT time-slots, influence in the last bT time-slots and influencing in each time-slot with probability b) as a function of the influence budget b via simulations. We observe that the strategy of influencing the last bT outperforms the other two strategies.

B. Time-varying Adamancy

Next, we study the case where the probability of an individual changing their opinion changes with time. This model incorporates the adamant nature of masses. The general wisdom is that humans have the tendency to get attached to the preferences they make early in their life. So, it gets difficult to recast their opinion as time passes. We present a two-phase opinion dynamics model. In the first phase called the flexible phase, individuals are more likely to change their opinion as compared to the second phase called the adamant phase. More specifically, compared to the flexible phase, in the adamant phase, the probability of an individual changing their mind (both with and without influence) goes down by a constant factor of $\eta > 1$. Our next assumption formally characterizes this setting.

Assumption 2 (Strong-Willed Population with Time-varying Adamancy)**.** *Given positive constants $\eta > 1$ and $\xi < 1$, the system evolves according to* (1) *with*

- $p_t = p$, $q_t = q$, $\tilde{p}_t = \tilde{p}$, $\tilde{q}_t = \tilde{q}$ *for* $1 \le t \le (1-\xi)T$,
- $p_t = \dfrac{p}{\eta}$, $q_t = \dfrac{q}{\eta}$, $\tilde{p}_t = \dfrac{\tilde{p}}{\eta}$, $\tilde{q}_t = \dfrac{\tilde{q}}{\eta}$ *for* $t > (1-\xi)T$,

such that $\tilde{p} \ge p$ and $\tilde{q} \le q$.

Our next result characterizes the structure of the optimal policy for this setting. As discussed in Remark 1, we limit our interest to the case where $\tilde{p} \ge p$ and $\tilde{q} \le q$.

Lemma 3. *Under Assumption 2, $\exists\ t_1, t_2$, where $0 \le t_1 \le (1-\xi)T$ and $(1-\xi)T < t_2 \le T$, such that the optimal influencing strategy is to influence in time-slots $t_1+1, t_1+2, \ldots, (1-\xi)T$ and $t_2, t_2+1, \ldots, T$.*

Proof. This is direct consequence of applying Theorem 1 for each phase of opinion dynamics. ☐

Definition 2 (Optimal Split)**.** *Under the optimal policy structure, a split or distribution of the slots to each phase of influence is said to be optimal if no other split gives a higher expected fraction of green balls at the end of time horizon T. Optimal split is characterized by a parameter $\kappa \in [0,1]$ such that the system is influenced in the last $(1-\kappa)bT$ time-slots of the fliexible phase and the last κbT time-slots of the adamant phase.*

Theorem 2. *Under Assumption 2, given a constant $\eta > 1$, let $\kappa^* = \dfrac{1}{1+\eta^{-1}}\left(1 + \dfrac{\xi}{b\eta} - \dfrac{2N}{bT}\ln\eta\right)$. If $p+q = \tilde{p}+\tilde{q} = c$, for some constant $c > 0$ and $\xi \ge b$, the optimal split is given by:*

$$\kappa = \begin{cases} 0 & \text{if } \kappa^* < 0, \\ 1 & \text{if } \kappa^* > 1, \\ \kappa^* & \text{otherwise}. \end{cases}$$

Proof. Without loss of generality, assume $p+q = \tilde{p}+\tilde{q} = 1$. Since from Theorem 3, we know that advertising towards the end of each phase is optimal, we define $t_1 = (1-b-\xi+\kappa b)T$, $t_2 = (1-\xi)T$, and $t_3 = (1-\kappa b)T$, where the interval $[0, t_1]$ is the period of zero adamance and no influence, $[t_1, t_2]$ is the period of zero adamance and advertising, $[t_2, t_3]$ is the period of adamant opinions and no influence and $[t_3, T]$ is the period of adamant opinions and advertising influence. Our assumption on b and ξ ensures that the influence regime can be completely contained in the adamant phase. Now, from the ODE analysis, we can obtain explicit expressions for $\rho_r(t_1), \rho_r(t_2)$ and $\rho_r(t_3)$. This gives us:

$$\begin{aligned}\rho_r(T) = \tilde{q}\Bigg(&\frac{1}{\eta} - \frac{1}{\eta}e^{\frac{-\kappa bT}{\eta N}} + e^{\frac{-\xi T}{\eta N}}\left(1 - e^{\frac{-(1-\kappa)bT}{N}}\right)\Bigg) \\ + q\Bigg(&\frac{1}{\eta}e^{\frac{-\kappa bT}{\eta N}} - \frac{1}{\eta}e^{\frac{-\xi T}{\eta N}} + e^{\frac{-\xi T}{\eta N}}e^{\frac{-(1-\kappa)bT}{N}} \\ &- e^{\frac{-\xi T}{\eta N}}e^{\frac{-(1-\xi)T}{N}}\Bigg) + \rho_r(0)e^{\frac{-(1-\xi)T}{N}}e^{\frac{-\xi T}{\eta N}}.\end{aligned}$$

Optimizing for the optimal split parameter κ, we get:

$$\kappa^* = \frac{1+\frac{\xi}{b\eta}}{\left(1+\frac{1}{\eta}\right)} - \frac{2N\ln(\eta)}{bT\left(1+\frac{1}{\eta}\right)}.$$

Since $\tilde{q} \le q$, this is a point of minima for $\rho_r(T)$ and since $\kappa \in [0,1]$, we get the expression in the theorem. ☐

Remark 2. *According to our setup $\kappa \le \xi/b$. As long as, $\xi/b \ge 1$, the above theorem gives the optimal split. However, when $\kappa^*, \xi/b < 1$ and $\kappa^* > \xi/b$, the optimal split is given by $\kappa = \xi/b$. This means that advertising is done in the last bT slots. This seems non-intuitive at first but notice that in this case ξ is small and $\kappa^* > \xi/b$ means that η is also quite small.*

Notice that κ^* is a decreasing function of η. That immediately gives us the threshold values of η for transition of the optimal κ.

Corollary 1. *Under Assumption 2 with $\xi \ge b$, let η_1 and η_2 be such that $\ln(\eta_1) = \frac{T}{2N}\left(b + \frac{\xi}{\eta_1}\right)$ and $\eta_2 \log \eta_2 = \frac{bT}{2N}\left(\frac{\xi}{b} - 1\right)$. Then, we have that,*

- *if $\eta \ge \eta_1$, the optimal strategy is to influence in the final bT time-slots of the flexible phase,*
- *if $\eta_1 > \eta > \eta_2$, the optimal strategy is to influence in the final $(1-\kappa)bT$ time-slots of the flexible phase and the final κbT time-slots of the adamant phase, where κ is as defined in Theorem 2,*
- *if $\eta \ge \eta_2$, the optimal strategy is to influence in the final bT time-slots of the adamant phase.*

In Figure 4, we compare the performance of three policies, namely, influence in the last bT time-slots of the flexible phase (Phase 1), influence in the last bT time-slots of the adamant phase (Phase 2) and influence according to the split characterized in Theorem 2) as a function of the parameter η defined in Assumption 2. We observe that the strategy characterized in Theorem 2 outperforms the other two strategies for all values of η considered. For small values of η, the optimal policy is very close to the policy to influence in the last bT time-slots of the adamant phase and for large values of η, the optimal policy is very close to the policy to influence in the last bT time-slots of the flexible phase. For intermediate values of η, the optimal policy outperforms both policies.

The key takeaway from this section is that if individuals are unaffected by the opinion of their peers and the effect

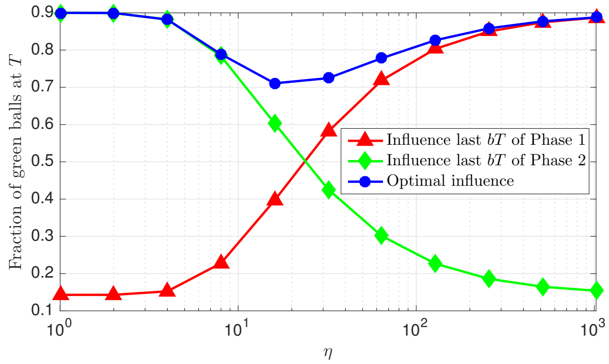

Fig. 4. Expected fraction of green balls at the end of time-slot T as a function of the influence budget b for three policies, namely, influence in the last bT time-slots of the flexible phase (Phase 1), influence in the last bT time-slots of the adamant phase (Phase 2) and influence according to the split characterized in Theorem 2) as a function of the parameter η defined in Assumption 2. The system satisfies Assumption 2 with parameters $p = 0.1$, $q = 0.6$, $\tilde{p} = 0.9$, $\tilde{q} = 0.1$, $b = 0.3$, $N = 1000$, $\rho(0) = 0.5$ and $T = 5000$. The strategy characterized in Theorem 2 outperforms the other two strategies for all values of η considered.

of external influence is time-invariant, the optimal influence strategy is to influence at the end of the finite time-horizon. In the case where the external influence becomes ineffective over time, influencing at the end of the finite time-horizon can be strictly sub-optimal.

V. Population Affected by Peer Pressure

We now focus the task of influencing a population where, in the absence of external influence, individuals are affected by the opinions of their peers. More specifically, in the absence of external influence, the probability of an individual changing their opinion increases with the fraction of the population holding the contrary opinion. The next assumption formally characterizes this setting.

Assumption 3 (Population Affected by Peer Pressure)**.** *The system evolves according to* (1) *with*

- $p_t = p_0(1-\rho_r(t))^\gamma$, $q_t = q_0(\rho_r(t))^\gamma$, *with constants* $\gamma > 0$ *and* $0 \le p_0, q_0 \le 1$,
- $\tilde{p}_t = 1$, $\tilde{q}_t = 0$.

Remark 3. *The assumption* $\tilde{p}_t = 1$, $\tilde{q}_t = 0$ *corresponds to an extremely strong external influence where if influenced, a red ball turns green with probability 1 and a green ball never changes its color. This assumption is made for mathematical tractability.*

In the analytical results in this section, we consider two special cases, namely, $\gamma = 1$ and $\gamma = 2$. We use simulations to explore the performance of influencing strategies for other values of γ. We first present our results for $\gamma = 1$.

Throughout this section, $\rho_r(T)|_{\text{first}}$ and $\rho_r(T)|_{\text{last}}$ denote the fraction of red balls at the end of time T when the influence happens in the first bT and the last bT slots respectively.

Theorem 3. *Under Assumption 3 with* $\gamma = 1$,

- *if* $p_0 > q_0$, *the optimal strategy is to influence in the first* bT *time-slots of the finite time-horizon,*
- *if* $p_0 < q_0$, *the optimal strategy is to influence in the last* bT *time-slots of the finite time-horizon,*
- *if* $p_0 = q_0$, *the performance of all policies which influence in* bT *out of the* T *time-slots in the finite horizon perform equally well.*

Proof. Let $\beta = p_0 - q_0$. We have the following ODEs corresponding to the stochastic approximation schemes.

- Without influence:

$$\frac{d\rho_r(t)}{dt} = \frac{-(p_0 - q_0)}{N}\rho_r(t)(1-\rho_r(t)). \tag{8}$$

- With Influence:

$$\frac{d\rho_r(t)}{dt} = \frac{-\rho_r(t)}{N}. \tag{9}$$

From (8) and (9), for the first bT influence strategy, we have:

$$\rho_r(bT) = \rho_r(0)e^{\frac{-bT}{N}} \;\&\; \rho_r(T) = \frac{1}{1 + \left(\frac{e^{\frac{bT}{N}}}{\rho_r(0)} - 1\right)e^{\frac{\beta(1-b)T}{N}}}. \tag{10}$$

Similarly, for the last bT strategy, we get:

$$\rho_r(T) = \frac{1}{e^{\frac{bT}{N}} + \left(\frac{1}{\rho_r(0)} - 1\right)e^{\frac{\beta(1-b)T+bT}{N}}}. \tag{11}$$

Notice that for $\beta = 0$, final fraction of red balls $\rho_r(T)$ is same for both the strategies. For $\beta \neq 0$, the result follows from (10) and (11) the monotonicity argument. $\square$

Our next result shows that for $\gamma = 2$, for a specific set of initial states, influencing in the first bT slots is a better strategy than influencing the last bT slots.

Proposition 1. *Under Assumption 3 with* $\gamma = 2$, *if the fraction of red balls at the beginning of time-slot 1 is less than half, the strategy of influencing in the first* bT *time-slots strictly outperforms the strategy of influencing in the last* bT *time-slots.*

Proof. The ODEs for influence and no-influence regimes are given by:

- Without influence:

$$\frac{d\rho_r(t)}{dt} = \frac{-\rho_r(t)(1-\rho_r(t))^2}{N}. \tag{12}$$

- With influence:

$$\frac{d\rho_r(t)}{dt} = \frac{-\rho_r(t)}{N}. \tag{13}$$

Since $\rho_r(0) < \frac{1}{2}$, for $\rho_r(t) \in [0, 1/2]$, $\exists c_1 > 0, c_2 > 0$ such that $\frac{3\rho_r(t)^2 - 2\rho_r(t)}{c_1 N} < \frac{-\rho_r(t)(1-\rho_r(t))^2}{N} < \frac{3\rho_r(t)^2 - 2\rho_r(t)}{c_2 N}$. In particular, the bounds hold for $c_1 = 2$ and $c_2 = 2.5$.

Now, for the strategy to influence in the first bT time-slots, we have: $\rho_r(bT) = \rho_r(0)e^{\frac{-bT}{N}}$ for the first bT time slots and $\frac{d\rho_r(t)}{dt} < \frac{3\rho_r(t)^2 - 2\rho_r(t)}{c_2 N}$ for $t \in ((1-b)T, T]$, and therefore, $1 - \frac{\frac{2}{3}}{\rho_r(T)} < \left(1 - \frac{2}{3}\right)\rho_r(bT)e^{\frac{2(1-b)T}{c_2 N}}$. Substituting $b = 1/2$

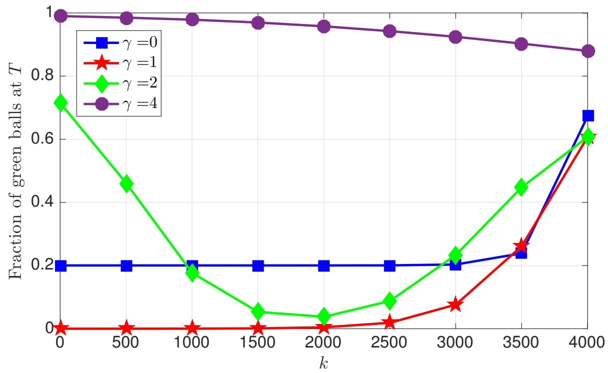

Fig. 5. Expected fraction of green balls at the end of time-slot T as a function of the influence budget b for a set of policies indexed by k of various values of γ. Policy k corresponds to the strategy of influencing in a block of bT time-slots from k to $k + bT - 1$. The system satisfies Assumption 3 with parameters $p_0 = 0.1$, $q_0 = 0.4$, $\tilde{p} = 0.1$, $\tilde{q} = 0.01$, $b = 0.2$, $N = 100$, $\rho(0) = 0.35$ and $T = 5000$. We observe that for the set of parameters considered, influencing towards the end of the time horizon outperforms the other policies for $\gamma = 0,\ 1$. For higher values of γ, influencing towrds the end of the time horizon is not necessarily optimal.

and denoting $x := \frac{T}{N}$, we get: $\rho_r(T)_{\text{first}} < \frac{2}{3-(3-4e^{x/2})e^{\frac{x}{c_2}}}$. For the policy with influence in last bT slots, we have: $\rho_r(T) = \rho_r((1-b)T)e^{\frac{-bT}{N}}$ for the last bT time slots and $\frac{d\rho_r(t)}{dt} > \frac{3\rho_r(t)^2-2\rho_r(t)}{c_1 N}$ for $t \in (0,(1-b)T]$.
For $b = 1/2$ and x, this implies $\rho_r(T)_{\text{last}} > \frac{2e^{-x/2}}{3+e^{x/c_1}}$. Then, for $c_1 = 2,\ c_2 = 2.5$, standard arguments imply $\frac{2e^{-x/2}}{3+e^{x/c_1}} > \frac{2}{3-(3-4e^{x/2})e^{\frac{x}{c_2}}}$. Hence, we have: $\rho_r(T)|_{\text{last}} > \rho_r(T)|_{\text{first}}$. □

In Figure 5, we compare the performance of multiple policies (indexed by k) for different values of γ. Policy k corresponds to the strategy of influencing in a block of bT time-slots from k to $k + bT - 1$. We observe that for the set of parameters considered, influencing towards the end of the time horizon outperforms the other policies for $\gamma = 0,\ 1$. For higher values of γ, influencing towrds the end of the time horizon is not necessarily optimal.

VI. Conclusions and Future Work

In this work, we proposed a variant of the voter model which can be used to model variation in the nature of the individuals in society. We evaluate the performance of campaigning strategies as a function of the nature of individuals when the goal is to maximize the fraction of individuals with a favorable opinion at the end of a known finite time-horizon.

We conclude that if individuals are unaffected by the opinion of their peers and the effect of external influence is time-invariant, influencing at the end of the finite time-horizon is optimal. In the case where individuals are affected by the opinion of their peers and/or external influence becomes ineffective over time, influencing at the end of the finite time-horizon can be strictly sub-optimal.

Possible extensions of this work include modeling the connections between individuals in the society using a graph such that individuals susceptible to being influenced by others are only influenced by their neighbors in this graph. Another direction worth exploring is allowing for heterogeneity in the nature of individuals in the same society.

VII. Acknowledgements

This work was supported in part by an Indo-French grant on “Machine Learning for Network Analytics". The work of Neeraja Sahasrabudhe was also supported in part by the DST-INSPIRE Faculty Fellowship from the Govt. of India. The work of Sharayu Moharir was supported in part by a seed grant from IIT Bombay.

References

[1] H. Xia, H. Wang, and Z. Xuan, “Opinion dynamics: A multidisciplinary review and perspective on future research,” *International Journal of Knowledge and Systems Science (IJKSS)*, vol. 2, no. 4, pp. 72–91, 2011.
[2] R. A. Holley and T. M. Liggett, “Ergodic theorems for weakly interacting infinite systems and the voter model,” *The annals of probability*, pp. 643–663, 1975.
[3] P. Clifford and A. Sudbury, “A model for spatial conflict,” *Biometrika*, vol. 60, no. 3, pp. 581–588, 1973.
[4] E. Yildiz, D. Acemoglu, A. E. Ozdaglar, A. Saberi, and A. Scaglione, “Discrete opinion dynamics with stubborn agents,” 2011.
[5] J. Majmudar, S. M. Krone, B. O. Baumgaertner, and R. C. Tyson, “The voter model and jump diffusion.” *arXiv preprint arXiv:1511.04160*, 2015.
[6] K. Kandhway and J. Kuri, “How to run a campaign: Optimal control of sis and sir information epidemics,” *Applied Mathematics and Computation*, vol. 231, pp. 79–92, 2014.
[7] B. Kotnis, A. Sunny, and J. Kuri, “Incentivized campaigning in social networks,” *IEEE/ACM Transactions on Networking*, 2017.
[8] D. Kempe, J. M. Kleinberg, and É. Tardos, “Influential nodes in a diffusion model for social networks.” in *ICALP*, vol. 5. Springer, 2005, pp. 1127–1138.
[9] S. Eshghi, V. M. Preciado, S. Sarkar, S. S. Venkatesh, Q. Zhao, R. D'Souza, and A. Swami, “Spread, then target, and advertise in waves: Optimal capital allocation across advertising channels,” *arXiv preprint arXiv:1702.03432*, 2017.
[10] E. Asano, L. J. Gross, S. Lenhart, and L. A. Real, “Optimal control of vaccine distribution in a rabies metapopulation model.” *Mathematical biosciences and engineering: MBE*, vol. 5, no. 2, pp. 219–238, 2008.
[11] U. Ledzewicz and H. Schättler, “On optimal singular controls for a general sir-model with vaccination and treatment,” *Discrete and Continuous Dynamical Systems*, vol. 2, pp. 981–990, 2011.
[12] A. A. Lashari and G. Zaman, “Optimal control of a vector borne disease with horizontal transmission,” *Nonlinear Analysis: Real World Applications*, vol. 13, no. 1, pp. 203–212, 2012.
[13] F. Chierichetti, S. Lattanzi, and A. Panconesi, “Rumor spreading in social networks,” in *International Colloquium on Automata, Languages, and Programming*. Springer, 2009, pp. 375–386.
[14] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Gossip algorithms: Design, analysis and applications,” in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 3. IEEE, 2005, pp. 1653–1664.
[15] B. Pittel, “On spreading a rumor,” *SIAM Journal on Applied Mathematics*, vol. 47, no. 1, pp. 213–223, 1987.
[16] V. S. Borkar, “Stochastic approximation,” *Cambridge Books*, 2008.
[17] S. Kamal, “On the convergence, lock-in probability, and sample complexity of stochastic approximation,” *SIAM J. Control Optim.*, vol. 48, no. 8, pp. 5178–5192, 2010.
[18] G. Thoppe and V. S. Borkar, “A concentration bound for stochastic approximation via alekseev's formula,” *arXiv preprint arXiv:1506.08657*, 2015.
[19] B. Kumar, V. S. Borkar, and A. Shetty, “Bounds for tracking error in constant stepsize stochastic approximation,” *arXiv preprint arXiv:1802.07759*, 2018.

VIII. Appendix

In this section, we use Concentration inequalities for Martingales to obtain a concentration results for fraction of red balls at the end of time horizon for the first bT strategy. The result for the strategy to influence in the last bT time-slots follows via similar arguments. Throughout this section, $x \overset{y}{\approx} z$ means $|x-z| \le y$. We use the following notation: $\alpha = \frac{p+q}{N}$, $\tilde{\alpha} = \frac{\tilde{p}+\tilde{q}}{N}$ and $\widehat{\alpha}_t = \frac{\widehat{p}_t+\widehat{q}_t}{N}$. Note that for the first bT influence strategy:

$$\widehat{\alpha}_t = \begin{cases} \tilde{\alpha} & \text{for } t \in [0, bT] \\ \alpha & \text{for } t \in (bT, T]. \end{cases} \tag{14}$$

Proposition 2. *Given $\epsilon > 0$, there exists $\delta = \delta(\epsilon, T)$ such that:*

$$P\left(\left|\frac{r(T)}{N} - M(T)\right| < \epsilon\right) \ge 1 - \delta,$$

where, $M(T) = \frac{q}{p+q}\left[1-(1-\alpha)^{(1-b)T}\right] + (1-\alpha)^{(1-b)T}\left[\frac{\tilde{q}}{\tilde{q}+\tilde{p}}\left(1-(1-\tilde{\alpha})^{bT}\right) + \frac{r(0)}{N}(1-\tilde{\alpha})^{bT}\right]$.

Proof. We know that $\mathbb{E}\left[r(t+1)|\mathcal{F}_t\right] = r(t)\left[1-\widehat{\alpha}_t\right] + \widehat{q}_t$. Let $P_r = \prod_{k=0}^{r}(1-\widehat{\alpha}_k)$, $P_{-1} = 1$, and $X(t) = \frac{r(t)}{P_{t-1}} - \sum_{r=0}^{t-1}\frac{\widehat{q}_r}{P_r}$. $X(t)$ is a Martingale with respect to $\mathcal{F}_t$. Define $Y(t) := P_T X(t)$. Clearly, $Y(t)$ is also a martingale. In fact, $Y(t)$ is a bounded martingale with $-1-\widehat{q}_t \le Y(t+1)-Y(t) \le \widehat{p}_t+1$. Assume that $\widehat{p}_s + \widehat{q}_s = \widehat{p} + \widehat{q}$ for all $s \in [0,t]$. Then, by Azuma-Hoeffding inequality, given $\epsilon > 0$, we have:

$$P\left(|Y(t) - Y(0)| > \epsilon\right) \le 2\exp\left(\frac{-\epsilon^2}{2t(\widehat{p}+\widehat{q}+2)^2}\right).$$

Henceforth, for sake of simplicity, we write that with probability at least $1 - 2\exp\left(\frac{-\epsilon^2}{2t(\widehat{p}+\widehat{q}+2)^2}\right)$, $Y(t) \overset{\epsilon}{\approx} Y(0)$. From this, we get that with probability at least $1 - 2\exp\left\{\frac{-\epsilon'^2\left(1-\widehat{\alpha}_t\right)^{2t}}{2t(\widehat{p}+\widehat{q}+2)^2}\right\}$,

$$r(t) \overset{\epsilon'}{\approx} \sum_{r=0}^{t-1} P_{t-1}\frac{\widehat{q}_r}{P_r} + P_{t-1}r(0), \tag{15}$$

where $\epsilon' := \epsilon/P_t$. Using (14) and (15) for the first bT influence strategy, given $\epsilon_1, \epsilon_2 > 0$, we get:

- With probability at least $1 - 2\exp\left\{\frac{-\tilde{\epsilon}_1^2 N^2\left(1-\tilde{\alpha}\right)^{2bT}}{2bT(\tilde{p}+\tilde{q}+2)^2}\right\}$,

$$\frac{r(bT)}{N} \overset{\tilde{\epsilon}_1}{\approx} \frac{\tilde{q}}{\tilde{q}+\tilde{p}}\left(1-(1-\tilde{\alpha})^{bT}\right) + \frac{r(0)}{N}(1-\tilde{\alpha})^{bT}, \tag{16}$$

  where $\tilde{\epsilon}_1 = \epsilon_1/P_{bT}$.

- With probability at least $1-2\exp\left\{\frac{\tilde{\epsilon}_2^2 N^2\left(1-\alpha\right)^{2(1-b)T}}{2(1-b)T(p+q+2)^2}\right\}$,

$$\frac{r(T)}{N} \overset{\tilde{\epsilon}_2}{\approx} \frac{q}{p+q}\left(1-(1-\alpha)^{(1-b)T}\right) + \frac{r(bT)}{N}(1-\alpha)^{(1-b)T}, \tag{17}$$

  where $\tilde{\epsilon}_2 = \epsilon/P_{(1-b)T}$.

Combining (16) and (17) and using the union bound, we get for $\delta = 2\exp\left\{\frac{-\tilde{\epsilon}_1^2 N^2\left(1-\tilde{\alpha}\right)^{2bT}}{2bT(\tilde{p}+\tilde{q}+2)^2}\right\} + 2\exp\left\{\frac{-\tilde{\epsilon}_2^2 N^2\left(1-\alpha\right)^{2(1-b)T}}{2(1-b)T(p+q+2)^2}\right\}$, with probability at least $1-\delta$:

$$\begin{aligned}\left.\frac{r(T)}{N}\right|_{\text{first}} \overset{\epsilon}{\approx} & \frac{q}{p+q}\left[1-(1-\alpha)^{(1-b)T}\right] \\ & + (1-\alpha)^{(1-b)T}\left[\frac{\tilde{q}}{\tilde{q}+\tilde{p}}\left(1-(1-\tilde{\alpha})^{bT}\right)\right] \\ & + (1-\alpha)^{(1-b)T}\left[\frac{r(0)}{N}(1-\tilde{\alpha})^{bT}\right].\end{aligned}$$

□

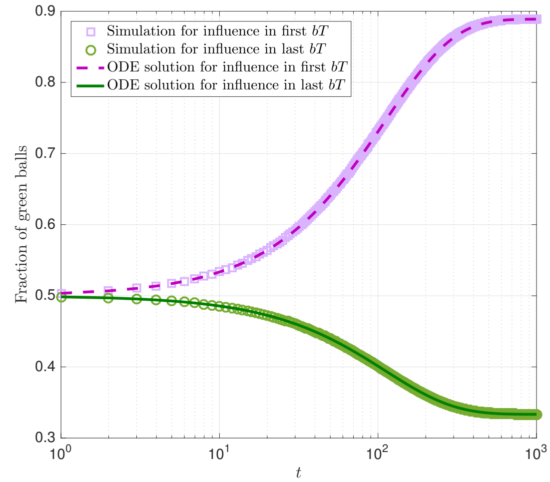

Fig. 6. Expected fraction of green balls at the end of time-slot t for the ODE and for the simulation of the model under Assumption 1, for both the policies, namely, influence in the first bT timeslots and influence in the last bT time-slots. The parameters are $p = 0.3, q = 0.6, \tilde{p} = 0.8, \tilde{q} = 0.1, b = 0.3, T = 1000, N = 100$. The ODE solution tracks the simulation in both cases.

Remark 4. *Note that δ also depends on $p, q, \tilde{p}, \tilde{q}, b$. In a large urn scenario, i.e. where $\alpha, \tilde{\alpha}$ are small $(1-\alpha)^{bT}$ can be approximated as $e^{-\alpha bT}$. This gives us an approximate solution $M(t)$ to the ODE corresponding to the stochastic approximation scheme. The proposition claims that the $r(T)$ obtained via simulation of the actual stochastic process is close to $M(T)$ with high probability (Figure 6). In fact for any t such that $0 < t < bT$ or $(1-b)T < t < T$, using the same concentration inequality, we get that $r(t)$ is close to $M(t)$ with a probability at least $1-\delta$, where $\delta = \delta(\epsilon, t)$. For a constant b and large time horizon T, $\left.\frac{r(T)}{N}\right|_{first} = \frac{q}{p+q}$, which is the stable fixed point of the ODE (4).*