

Spread, then Target, and Advertise in Waves: Optimal Capital Allocation Across Advertising Channels

Soheil Eshghi, Victor M. Preciado, Saswati Sarkar, Santosh S. Venkatesh,
Qing Zhao, Raissa D’Souza, and Ananthram Swami

Abstract—We obtain optimal strategies for the allocation of influence budget across multiple channels and across time for an external influencer, e.g., a political campaign, seeking to maximize its effect on an election given a network of agents with linear consensus-seeking opinion dynamics. We show that for a general set of objective functions, the optimal influence strategy at every time uses all channels at either their maximum rate or not at all. Furthermore, we prove that the number of switches between these extremes is bounded above both by a term that is typically much smaller than the number of agents. This means that the optimal influence strategy is to exert maximum effort in waves for every channel, and then cease effort and let the effects propagate. We also show that at the beginning, the total cost-adjusted reach of a channel determines its relative value, while targeting matters more closer to election time. We demonstrate that the optimal influence structures are easily computable in several practical cases. We explicitly characterize the optimal controls for the case of linear objective functions via a closed form. Finally, we see that in the canonical election example, identifying late-deciders approximately determines the optimal campaign resource allocation strategy.

I. INTRODUCTION

Opinions are important definers of real-world outcomes: they affect who is elected for political office [1], which policies are successful [2], and which products are bought by customers [3]. The proliferation of online media has complicated [4], sped up [5], and enhanced [6] opinion change processes. The opinion change process can be affected by interested parties through *advertising channels*, which are media by which messages are distributed to a target audience. Political campaigns and marketing departments apportion their advertising budgets between such channels (e.g., TV ads, website banner ads, billboards) in order to maximize some ultimate goal (e.g., votes, sales) [7]. The importance of this decision has increased in conjunction with the increasing resources devoted to these efforts: the 2016 US election campaigns, for example, were estimated to spend \$4.4 billion on TV advertising alone [8], while worldwide total marketing spend was estimated to be over \$1.6 trillion [9]. Thus, studying

Soheil Eshghi is with the School of Electrical Engineering, and Yale Institute of Network Science (YINS), Yale University, New Haven, CT 06520
soheil.esghghi@yale.edu

Victor M. Preciado, Saswati Sarkar, and Santosh S. Venkatesh are with the School of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 preciado, swasti, venkatesh@seas.upenn.edu

Qing Zhao is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 15853 qz16@cornell.edu

Raissa D’Souza is with the Departments of Computer Science and Mechanical and Aerospace Engineering, University of California, Davis, CA 95616 raissa@cse.ucdavis.edu

Ananthram Swami is with the Army Research Laboratory, Adelphi, MD 20783 awswami@arl.army.mil

The work of Eshghi, Zhao, D’Souza, and Swami were supported in part by the Army Research Laboratory Network Science CTA under Cooperative Agreement W911NF-09-2-0053.

the related multi-channel resource allocation problem is both timely and significant.

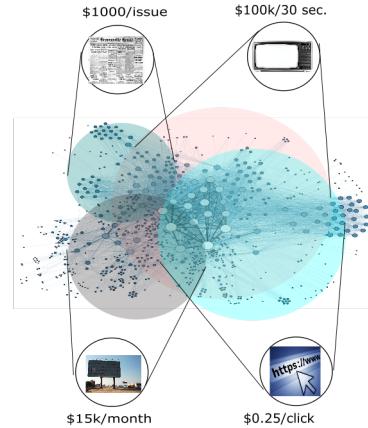


Fig. 1: Different advertising channels have differing, possibly overlapping, reaches and come at differing costs. Furthermore, the effect of a channel may be different on various individuals. These external influences are modulated by internal conversations within the network whereby agents integrate this information with that of their neighbors. The decision of the advertiser, therefore, is to apportion resources between these channels factoring in these complexities so as to get an optimal return on investment, which in the election example is votes cast in favor of the campaign.

In particular, influence on opinions can be classified into two types based on its direct provenance: in the *endogenous* influence mechanism (e.g., word-of-mouth), individuals process the expressed opinions of other individuals they meet, and consider their credibility and the level of acquaintance and trust in synthesizing a new opinion based on the information.¹ This leads to the notion of an endogenous influence weighted graph that captures these (consensus-seeking) processes. On the other hand, in the *exogenous* influence mechanism, an external influencer seeks to convince an individual to change their mind. This mechanism is facilitated by various advertising channels² in the opinion formation model that have a channel reach structure (which denotes the nodes that can be reached by a particular channel) that is not necessarily related to the endogenous neighborhood. The actions of other influencers also affect each individual’s opinion change process, which

¹These weights can, in general, be dynamic, even depending on the expressed opinion [10]. In this work, we consider static weights.

²Throughout this work, we use the word channel to represent both the medium (e.g., TV advertising) and the reach of the medium (e.g., people who watch TV) – the distinction is clear in context.

can in general be random (noisy) process. The influencer seeks to maximize a function of the *global state* (the vector of individual opinions) at a specific time (corresponding, for example, to an election or a new product launch) by controlling how much of their budget they will allocate to each costly advertising channel at their disposal at various times (see Fig. 1). In this paper, we study the nature of these optimal allocations and provide structures and algorithms for their computation.

However, this apportioning decision is complicated by several factors: 1- The reach of each channel is limited, and there are significant overlaps in the target audience of various channels [11]. 2- Different channels have differing costs, and attempts to influence opinions by external sources can affect individuals in different, and sometimes opposite, ways [12]. 3- The apportioning decision is dynamic (depends on time) and changes with the state of the network, and therefore constitutes an optimization in the space of functions.

Furthermore, the influencer faces several trade-offs: influencing a channel earlier allows the influenced individuals to spread the effect to their neighbors (*diffusion*), while lessening the impact on the influenced individuals as they moderate the effects of the external influence with the opinions of their neighbors (*dilution*). There is also a trade-off between utilizing cheap channels versus utilizing expensive but effective ones. These competing forces make the *a priori* determination of optimal influence allocation hard to determine.

There are also significant technical challenges to solving this problem, as characterizing the optimal allocation of budget to channels requires characterizing the structure of an optimal constrained vector of controls over a graph. There does not exist a general theory of the structure of vector optimal controls that would define the nature of these controls, much less so in the presence of a general graph structure. Furthermore, the work also requires computing the optimal control of the well-studied linear consensus dynamics [13], [14] in a novel setting, as the classical literature is concerned with reaching agreement among agents, while our objectives may incentivize agreement in some circumstances and disagreement in others. Thus, finding the new allocation requires a new synthesis of spectral graph theory and optimal control theory.

Contributions: In this work, we model the advertising influence problem as a constrained graphical consensus control process (without requiring the homogeneity and independence assumptions of epidemiological models for the same process [15]) with overlapping exogenous influence channels and the endogenous influence of agents upon each other [16]. Using Pontryagin's Maximum principle, spectral graph theory, and custom analytical arguments, we determine the structure of optimal the allocation of capital to the various influence channels at various times during the time horizon.

We show³ that for a general class of objectives, the optimal control for each of the channels is bang-bang (only takes its extreme values), with the number of switches being upper-

bounded by a term⁴ which is itself upper-bounded by the number of individuals. These structures mean that the search for optimal controls can be conducted on the space of vectors of a fixed size rather than on the space of functions.

Furthermore, for the case of a linear objective (i.e., when individuals make a decision in proportion to their opinion), we explicitly calculate the candidate optimal controls, providing an open-loop algorithm with respect to the dynamics that can compute the vector of optimal controls in a logarithmic number of steps. This allocation also implicitly determines the relative criticality of a particular channel to the global objective, and thus defines **an explicitly computable metric for the influence of a channel** at any given time that allows the influencer to compare and contrast the effects of different channels, as well as the effect of a channel at different times. We prove that *reaching likely voters* is useful for ranking the influence of channels *only close to decision/election time*, while the *cost-effectiveness* of a channel (defined as its total reach divided by its cost) is more important at earlier times.

For the case where the objective is a sum of sigmoids (e.g., individuals make a decision whether they are convinced or not by a specific party before deciding to vote), which is a relaxed version of voting between two alternatives, we show that the optimal control can be approximated just by knowing the agents who change their minds at the terminal time in the optimal allocation (*late-deciders* [17]).

In sum, the work represents a new confluence of the literature on consensus dynamics and optimal control theory, while providing significant novel structures, computational algorithms, metrics, and insights to the optimal capital allocation for advertising problem.

II. LITERATURE REVIEW

As this work draws upon the literature in multiple areas, we will discuss antecedents in each area in turn:

Consensus: Linear consensus-seeking dynamics are some of the oldest models used to model the spread of opinions [18]. In these models, opinions (states) are taken to be continuous, and each node uses a (weighted) average opinion of its neighbors' opinions in each time-step to update its opinion. Continuous-time variants of these dynamics, as well as the closely related literature on flocking [19] have been rigorously studied by control theorists [13], [14], [20]. Most of these results focus on asymptotic properties of these dynamics and their convergence, and not on their finite-time behavior and the effect of influence on such behavior.

Control of Opinion Propagation: The case of affecting the outcome of such opinion dynamics is a research question that has only recently been studied. Yildiz *et al.* [21] consider the case of stubborn agents who refuse to change their opinions in a two-opinion voter model. They show that the mean average opinion is only a function of the structure of the network and the placement of the stubborn nodes. They then investigate the optimal placement of these stubborn nodes. Kempe *et al.* [22] and Anshelevich *et al.* [23] also discuss how to

³for space reasons, these proofs are available in the technical report for this manuscript [?].

⁴The number of non-zero elements in the orthogonal projection of the channel influence vector on the eigenvectors of the consensus dynamics

contain the spread of an opinion through an optimal initial placement of stubborn nodes holding opposing opinions for differing models of opinion propagation. However, the focus of all of these papers has been on static optimization, i.e., actions that are taken at a specific point in time. On the other hand, social networks are naturally dynamic, i.e., their states are time-varying, and it is natural to assume that actions prescribed to affect them can also be dynamic. Such optimal actions (henceforth referred to as controls) can be derived using optimal control theory.

Linear Optimal Control: Optimal control theory has addressed issues of linear control in detail [24]–[26]. Results on the bang-bang nature and the finite number of switches in linear optimal control in the case where actions are not costly and the time horizon is not fixed are well known [24]. However these results do not apply directly in the case with costly actions and where the goal is not to drive the system to a known state in minimum time. Therefore, the current work represents a step beyond those results and provides a context-specific method of evaluating the relative value of influence on a channel at all points within a time horizon, a question of practical importance that was not addressed by this strain of work.

Optimal Control of Epidemic Spread and Diffusion: This work bears a similarity with the literature on the optimal control of information spread, in that both aim to optimize a terminal function subject to some spread dynamics. Most such work uses compartmental epidemic models (e.g., SI [15], [27]) and is thus dissimilar in dynamics to the one we consider. Furthermore, we show that when opinions can take continuous values (instead of the finite fixed values assumed in compartmental models), the optimal controls for influence maximization are significantly different to the structures proposed for information spread (which typically advocate some form of maximal spreading at the start of the time interval [27], [28]). The model also allows an even more explicit incorporation of graph structure than metapopulation models, e.g., [29], as their approximation breaks down when the population of each patch/type is small, and therefore provides a poor model for interactions at the scale of individuals.

Adversarial Sensor Network Deception: Finally, the problem discussed in this paper has a direct analog in the optimal deception of a sensor network by an adversary, as discussed in [30]. In this setting, a state-estimation sensor network [31] can be misdirected through local noise injection at a fixed number of points, that will effect a subset of nodes in the vicinity. The optimal locations and patterns for the noise to affect the conclusion of the network will depend also on the dynamic information fusion model of the sensor network and its relationship with the reach of each of the noise injection points. This problem, too, will require the same type of exogenous influence and endogenous processing model as the opinion influence problem, as well as having the same objective structure. Thus, any structural results obtained will have direct implications for the adversary's optimal deception policy.

In summary, the work integrates elements of the rich literature in linear consensus protocols, spectral graph theory, and

optimal control, and applies the synthesis to the problem of resource allocation in advertising, achieving strong structural guarantees and applied insights.

III. SYSTEM MODEL DESCRIPTION

In this section, after presenting our notation (§III-A), we outline our system model (§III-B) and its dynamics (§III-C). Then we outline the bounds on the actions of the influencer (§III-D) and describe their objective (§III-E). We finish the section by presenting a technical assumption (§III-F) and by stating the overall problem (§III-G).

A. Notation

$$\begin{aligned} x_i(t) &= \text{opinion of agent } i \text{ at time } t \\ u_k(t) &= \text{influence on channel } k \text{ at time } t \\ a_{ij} &= \text{magnitude of effect agent } j \text{'s opinion on the} \\ &\quad \text{opinion of agent } i \\ N_i &= \text{neighbors of agent } i \text{ in } G \\ b_{ik} &= \text{relative magnitude of the effect of influence on} \\ &\quad \text{channel } k \text{ on agent } i \text{ in channel } k \\ e_i(t) &= \text{sum effect of other influence on agent } i \text{ at time } t \\ T &= \text{terminal time} \\ c_k(\cdot) &= \text{cost of influence on channel } k \\ r &= \text{total resources of influencer over the time period} \\ J_i(\cdot) &= \text{Value of opinion of agent } i \text{ at time } T \text{ to influencer} \\ J(\cdot) &= \text{Value of opinion profile at time } T \text{ to influencer} \end{aligned}$$

For compactness, we use $[n]$ to represent $\{1, 2, \dots, n\}$. For a matrix \mathbf{W} , we denote the k -th column of \mathbf{W} as \overrightarrow{w}_k , and the k -th row of the same as \overleftarrow{W}_k . Furthermore, we use w_{ij} to denote the (i, j) -th element of the matrix \mathbf{W} .

B. System Model

We assume that the system contains n agents, with the (continuous valued) *state/opinion* of agent $i \in [n]$ at time t denoted by $x_i(t) \in \mathbb{R}$. Each agent communicates with other agents based on an edge-weighted **undirected connected** communication graph $G = (V, E, \mathbf{A})^5$. The (non-negative) weight on an edge between agents $i, j \in [n]$, which determines the relative weight agent i gives to agent j in its state update, is represented by a_{ij} , and the matrix of such weights is represented by \mathbf{A} . An agent j is said to be a *neighbor* of agent i (and vice versa) if $a_{ij} > 0$ (See Fig. 2).⁶

At each time t , each agent updates their state based on a *weighted average* of the difference of its current state and that of its neighbors, as well as the influence that will be described below, and a known drift signal (which may be due to the influence of other competing influencers), which we denote by $e_i(t)$ for $i \in [n]$. The linear averaging model of opinion updates has a long history in opinion and social influence

⁵However, the results can be generalized to weighted directed graphs in which the weighted Laplacian has a generalized Jordan form with real eigenvalues

⁶While negative weight updates are conceivable, they will not be considered in this paper.

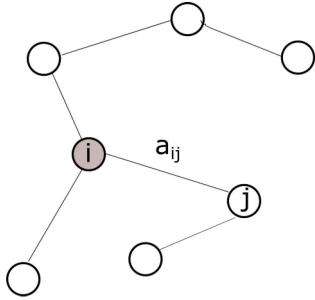


Fig. 2: Each agent takes into account the options of its neighbors in updating its own opinion. The weight given by node i to the opinion of its neighbor j is a measure of how much it trusts their appraisal. If agents i and j are not neighbors, we will have $a_{ij} = 0$ by default. In our model, we assume that neighborhood is a symmetric relation and that $a_{ij} = a_{ji}$, i.e., trust is a symmetric relation.

studies, both in continuous time (the celebrated model of Abelson [32], [33] which serves as a building block for this model) as well as in discrete time [18], [34], as well as being a well-studied model of consensus [13], [35]. Mathematically, it can be thought of as a gradient descent algorithm implemented by agents seeking to minimize disagreement (measured by Laplacian potential) [35].

An influencer aims to shape the *opinion profile* (i.e., the opinion vector of all agents) at a fixed terminal time T according to its objective function through the judicious use of its fixed, known *influence channels*. Each channel of influence (e.g., advertising medium) is limited in its reach.

We say that any channel only affects a specific subset of agents, hence called a *channel*. The structure of these channels is pre-specified and captured in a hyper-graph $G' = (V, H)$ on the same node set. There are $|H| = m$ hyper-edges (channels) that can be influenced, with the assumption that influencing a channel only directly affects the members within that channel. The influence exerted by the influencer on channel $i \in [m]$ is denoted by the scalar u_i .

In this model, the effect of influence on a channel can differ on all agents within the channel, potentially even having opposite effects. These effects are captured by the *influence gain*: b_{ik} determines the linear relative gain of influence on channel $i \in [m]$ on agent $k \in h_i$ within that channel. If agent k is not within channel i , we define b_{ik} to be zero. Without loss of generality, we assume $u_k \geq 0$ for $k \in [m]$, and encode the possible negative effects in the respective gains. Stacking these values into a matrix \mathbf{B}_{n*m} captures the hypergraph structure of the channels. The linear approximation of the effect of external influence on opinion dynamics is also a long-standing tradition [33], [36].

C. Dynamics

To understand the dynamics, we provide the following discrete-time intuition: an agent $i \in [m]$ constructs its change in state in the time interval $(t, t + \Delta)$ based on the weighted difference between its own state and that of its neighbors, as

well as the influence exerted on it in that time period and the drift signal:

$$x_i(t + \Delta) = x_i(t) + \Delta \left(\sum_{j \in N_i} a_{ij}(x_j(t) - x_i(t)) + \sum_{k:i \in h_k} b_{ik} u_k(t) + e_i(t) \right).$$

This simply states that agents attempt to align their state/opinion with that of their neighbors, and the influencer's effort acts as a hindrance to that process. Note that the above can be re-written to represent the classic discrete-time consensus model [18] with influence:⁷

$$x_i(t + \Delta) = (1 - \Delta \sum_{j \in N_i} a_{ij}) x_i(t) + \sum_{j \in N_i} \Delta a_{ij} x_j(t) + \Delta \left(\sum_{k:i \in h_k} b_{ik} u_k(t) + e_i(t) \right).$$

Subtracting $x_i(t)$ from both sides, dividing by Δ and taking the limit as Δ goes to zero, we arrive at the following continuous time agent-level dynamics:

$$\begin{aligned} \dot{x}_i(t) &= \sum_{j \in N_i} a_{ij}(x_j(t) - x_i(t)) + \sum_{k:i \in h_k} b_{ik} u_k(t) + e_i(t) \\ &= (-\sum_{j \in N_i} a_{ij}) x_i(t) + \sum_{j \in N_i} a_{ij} x_j(t) \\ &\quad + \sum_{k:i \in h_k} b_{ik} u_k(t) + e_i(t) \end{aligned}$$

We let \mathbf{L} be the weighted Laplacian matrix, where for all $i, j \in [n]$ such that $i \neq j$, $l_{ij} = -a_{ij}$, and $l_{ii} = \sum_{j \in N_i} a_{ij}$ for all $i \in [n]$. Stacking these equations, we arrive at the following system-level dynamics:

$$\vec{x}_{n*1}(t) = -\mathbf{L}_{n*n} \vec{x}_{n*1}(t) + \mathbf{B}_{n*m} \vec{u}_{m*1}(t) + \vec{e}_{n*1}(t) \quad (1)$$

We assume that the states/opinions at time 0 are known ($\vec{x}(0) = \vec{x}_0$), however we will see that the value of the states at time 0 has no direct bearing on our structural results.

D. Admissible Control Strategies

The total expenditure on all channels is bounded by $r > 0$, the amount of capital available to the influencer. This is captured through the following capital constraint:

$$\int_0^T \sum_{k=1}^m c_k(u_k(t)) dt. \leq r \quad (\text{budget constraint}) \quad (2)$$

Assumption 1. We assume that for all $k \in [m]$, $c_k(\cdot)$ is increasing, differentiable and concave. Furthermore, without loss of generality, we assume that $c_k(0) = 0$ for all $k \in [m]$.

This models the high start-up cost of utilizing a channel relative to the diminishing cost of additional influence once a

⁷The results derived in this paper would also apply to the Friedkin and Johnsen model of opinion updates [36] given uniform susceptibility to change across agents.

channel is already being influenced. Nevertheless, the above assumption allows the case where c_k is linear.

We assume that for all channels $k \in [m]$, the influence that can be exerted on channel k at any time t is bounded above by a time-varying value $u_k^{\max}(t)$.⁸ This can capture both physical limits on the influence (i.e., availability of media) and limits on the susceptibility of agents to the influence. Note that we only require $u_k^{\max}(t)$ to be differentiable, which is not a strong assumption. This includes the case where u_k^{\max} is a constant.

$$\forall k \in [m] \quad 0 \leq u_k(t) \leq u_k^{\max}(t) \text{ (influence constraint)} \quad (3)$$

We will restrict ourselves to \vec{u} that are piecewise continuous (i.e., that only have a finite number of discontinuities)⁹. We shall use \mathcal{U} to denote the set of such controls that fulfill (3):

$$\mathcal{U} := \{ \vec{u} : 0 \leq u_k(t) \leq u_k^{\max}(t), k \in [m], t \in [0, T] \}$$

In our model, employing channel $k \in [m]$ at effort level u at time $t \in [0, T]$ incurs a cost of $c_k(u)$, where for all $k \in [m]$, $c_k(\cdot) : [0, \max_{t \in [0, T]} u_k^{\max}(t)] \rightarrow [0, +\infty)$. Note that we assume that for all $k \in [m]$, $c_k(\cdot)$ is time-invariant.

E. Objective

The objective of the influencer is a function of the opinion profile at a fixed time T . The nature of the function will depend on the information aggregation method employed by the set of individuals. We consider the most general case, where any increase in the opinion of any particular individual at time T (keeping all other opinions the same) is not detrimental to the influencer. This is obviously the case in both political and marketing campaigns.

While our reasoning applies to a general family of objective functions, we will give special consideration to functions that model voting in an election between two options (relevant in the political campaign setting) and weighted averaging (relevant in estimating total returns from marketing efforts and in the sensor network setting).

Assumption 2. *The objective, $J(\vec{x}(T))$, is an increasing¹⁰, differentiable function of vector of terminal opinions ($\vec{x}(T)$).*

In particular, we will mention the application of our results to a particular family of objective functions that are separable in the elements of the vector of opinions:

$$J(\vec{x}(T)) = \sum_{i=1}^n J_i(x_i(T)). \quad (4)$$

Two specific types of separable functions are of interest:

1) Linear functions:

$$J_i(x_i(T)) = p_i x_i(T), \quad p_i > 0, \quad (5)$$

which model the simplest case, where the utility the influencer gains from an individual has a linear relationship with its state at time T . In the marketing example, this can model the amount of sales as a simple function of an individual's

⁸This rules out impulse controls.

⁹Note that this means, in particular, that the integral in (2) is well-defined

¹⁰Note that this encodes the non-negative benefit of an increase in opinions at time T to the influencer.

opinion of a product. This is also a useful approximation in the intruder detection case where the utility for the influencer is a simple function of a sensor's report level. The mapping of the election example to this utility is not direct: this can model the case where each agent votes with probability p_i , and if it does, chooses among two options with a probability that is linearly related to their opinion (i.e., they flip an appropriately weighted coin). However, finding the correct total weight of the coin to be considered depends on the assumptions made by the modeler, and multiple normalizations may be defensible. This ambiguity leads to the definition of a second type of utility for the specific case of the election example.

2) Sigmoid functions: Assume each individual $i \in [n]$ has to vote for one of two options/candidates/products/policies (0 and 1) at time T , and that the influencer backs option 1 (without loss of generality). Each individual is assumed to vote with probability $p_i > 0$, and to choose who to vote for among two options based on whether their state at time T is above or below a agent-specific threshold θ_i (which models the various biases for and against an option). Thus, the utility gained from each individual can be modeled using a Heaviside function with a jump at θ_i ($h(x_i(T) - \theta_i)$ (which is agent i 's vote):

$$J_i(x_i(T)) = p_i h(x_i(T) - \theta_i).$$

However, this utility is discontinuous at $x_i(T) = \theta_i$, which complicates analysis. The sigmoid (see Fig. 3):

$$J_i(x_i(T)) = \frac{p_i}{1 + e^{-\alpha_i(x_i(T) - \theta_i)}}, \quad (6)$$

is a smooth approximation to the Heaviside utility, with the closeness of the approximation being determined by the choice of the parameter α_i — the greater α_i is, the faster the transition. In the extreme of taking α_i to infinity, this function will indeed converge to the aforementioned Heaviside function.

F. Technical Assumption

We now add some technical assumption that will be needed in our arguments:

Assumption 3. *There exists a $j \in [n]$ such that $\frac{\partial J(\vec{x}(T))}{\partial x_j(T)} > 0$ for all $x(T) \in \mathbb{R}$.*

Note that this is equivalent to saying there exists at least one individual such that the influencer always values a marginal increase in its state. That is, holding all opinions the same, any increase in that agent's opinion will be translated to a strict increase in their likelihood of voting for the choice backed by the influencer.¹¹ The purpose of this assumption is to rule out a pathological case where the necessity conditions for the optimality of an allocation become so general that they apply for all controls and are thus uninformative.

¹¹Note that this rules out $J(\cdot)$ functions with stationary points, i.e., those for which $\nabla_{x(T)} J(x(T)) = \vec{0}$ for some $x(T) \in \mathbb{R}$.

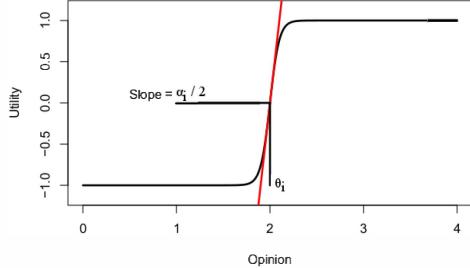


Fig. 3: The sigmoid function is a smooth approximation to a Heaviside step function. It maps the opinions of an agent into their probability of voting for the option favored by the influencer. The parameter θ_i determines the convincing threshold for individual i , where the probability the agent votes for the favored option increases significantly if they are convinced beyond their threshold. The parameter α_i determines the sharpness of this dependence on the bias - large α_i model an abrupt transition from not voting for the choice to voting for it given a small change of opinion, while smaller α_i represent a more gradual transition.

G. Problem Statement

We aim to find analytical and computational structures for the controls that maximize $J(\vec{x}(T))$ under the dynamics outlined in (1) and with constraints (3) and (2) on the controls:

$$\begin{aligned} \max & J(\vec{x}(T)) \\ \text{s.t.} & \dot{\vec{x}} = -\mathbf{L}\vec{x} + \mathbf{B}\vec{u} \\ & 0 \leq u_m(t) \leq u_m^{\max}, \quad \forall m \\ & \int_0^T \sum_{m=1}^M c_m(u_m(t)) dt \leq r \\ & \vec{u} \in \mathcal{U} \end{aligned}$$

In the most general case, this problem has a general-form, potentially non-convex objective function, as well as a potentially non-convex budget constraint (when any of the $c_i(\cdot)$'s is strictly concave). Furthermore, the number of competing influence channels, m , and therefore the number of optimization variables, can potentially be large. These factors complicate naive approaches to solving the problem.

We reformulate the problem with auxiliary variables to aid the analysis. We define the auxiliary variables γ and q such that,

$$\begin{aligned} \gamma(0) = 0, & \quad \dot{\gamma}(t) = -\sum_{k=1}^m c_k(u_k(t)), \quad (7) \\ q(0) = 0, & \quad \dot{q}(t) = 1, \quad (\text{proxy for time}). \quad (8) \end{aligned}$$

As can be seen, $\gamma(t)$ is the running cost of the influence up to time t , and q is a proxy for time. Thus, the capital constraint becomes $\gamma(T) \geq -r$, and the integral constraint has been

transformed to a terminal time one. So we can rewrite the optimization as:

$$\begin{aligned} \max & J(\vec{x}(T)) \\ \text{s.t.} & \dot{\vec{x}}(t) = -\mathbf{L}\vec{x}(t) + \mathbf{B}\vec{u}(t) + \vec{e}(t), \\ & \dot{\gamma}(t) = -\sum_{k=1}^m c_k(u_k(t)), \quad \dot{q}(t) = 1, \\ & \vec{u} \in \mathcal{U}, \quad \gamma(T) \geq -r. \end{aligned} \quad (9)$$

IV. RESULTS

In this section, we outline the analytical structures of the optimal controls. To show the nature of the results, we first explain some necessary priors §IV-A. Then we prove the existence of optimal controls §IV-B (under some conditions) and identify their structure using our main theorem §IV-C. A refinement is presented for the case of the linear objective §IV-D that allows the direct computation of the control and shows that the optimal control is unique, while providing insight into the logic of the allocation decision. Finally, the sigmoid approximation to voting is revisited §IV-E and an approximation to the optimal control is presented.

A. Preliminaries

For an **undirected, connected** graph G , the weighted Laplacian matrix \mathbf{L} will be real and symmetric (i.e., \mathbf{L} will be Hermitian), and therefore \mathbf{L} will be positive semi-definite and will have real, non-negative eigenvalues [37, page 13]. Thus, it will have an eigen-decomposition [38, Appendix A.5] (i.e., $\mathbf{L} = \mathbf{Q}\Xi\mathbf{Q}^T$, where \mathbf{Q} is an orthogonal matrix where all columns are eigenvectors of \mathbf{L} and Ξ is a diagonal matrix where the diagonal elements are eigenvalues of \mathbf{L}).¹² The smallest eigenvalue will be **zero** and will have multiplicity 1 (as \mathbf{L} is row-stochastic and G is connected) [37, page 13]. We will order the eigenvalues of Ξ smallest to largest ($\xi_1 = 0 < \xi_2 \leq \dots \leq \xi_n$) and therefore column $i \in [m]$ of \mathbf{Q} , \vec{Q}_i will be the i -th eigenvector of \mathbf{L} ¹³. Note that this means that $\vec{Q}_1 = \frac{1}{\sqrt{n}} \vec{1}_n$, where $\vec{1}_n = (1, \dots, 1)^T$.

We now state and prove a lemma that shows optimal controls for the main problem exist. We then provide the claim of our main result (Theorem 1) and then present a subcase where the bound can be significantly strengthened and the optimal control can be calculated open-loop (Theorem 2).

B. Existence Result for Optimal Solutions

Lemma 1. *Optimal controls for problem (9) exist for the case of linear $c_i(\cdot)$'s.*

This is because due to Filippov's theorem [25, page 119], the reachable set at T for the dynamics is compact and convex, and thus by the Weierstrass theorem [25, page 7] the objective will have a global maximum within the reachable set that is attainable by a valid control (for a full proof, see [40]).

¹²However, the reasoning below applies to any \mathbf{L} that has a generalized Jordan form with real eigenvalues.

¹³Notice this is possible because \mathbf{L} is Hermitian and therefore the geometric multiplicity of all eigenvalues is equal to their algebraic multiplicity [39, page 6-5].

C. Structural Results for the Optimal Control

Theorem 1. For all $i \in [m]$ and for the case where

- $c_i(\cdot)$ is strictly concave, $\langle \vec{B}_i, \vec{1}_n \rangle \neq 0$, and $u_i^{\max}(t)$ is constant
- $c_i(\cdot)$ is linear and the system $(\mathbf{L}, \mathbf{L}\vec{B}_i)$ is controllable [41, p. 144]

1) Optimal controls are **bang-bang**, taking on their maximum or minimum values at all times t ($u_i^*(t) \in \{u_i^{\max}, 0\}$).

2) The number of switches between these values is bounded above:

- In the general case, by one less than the number of non-zero elements in $\{\langle \vec{Q}_j, \vec{B}_i \rangle\}_{j=1}^N$.
- For $J(\vec{x}(T)) = \langle \vec{p}, \vec{x}(T) \rangle$, by the number of sign variations in $\{\sum_{k=1}^j s_k\}_{j=1}^N$, where $s_j := \langle \vec{Q}_j, \vec{p} \rangle \langle \vec{Q}_j, \vec{B}_i \rangle$.

The proof of this theorem relies upon Pontryagin's Maximum Principle [26, page 182], a necessary condition of optimality, spectral arguments for the Laplacian, and some real analysis (again, the full proof is available in [40]). A *switching function* is defined to characterize the optimal control, and through analysis, the number of switches of this function is upper-bounded.

An example of an optimal control with these characteristics is provided in Fig. 4. The actual number of switches of each optimal control can be significantly less than the fixed upper-bound of $N - 1$ (which can in general be very high). The conditions in the statement of the theorem rule out pathological cases where the necessary conditions for optimality derived from the Maximum Principle [26, page 182] do not rule out singular arcs [25, page 113] within the optimal controls.¹⁴

Note that the theorem conditions can apply for some $i \in [m]$ and not others, in which case the derived structure only applies to those i that fulfill the conditions of the theorem.

D. Water-filling: Optimal Budget Allocation for the Separable Linear Objective

In the proof of Theorem 1, we obtain explicit necessary conditions for the optimal control in the case of $J(\vec{x}(T)) = \langle \vec{p}, \vec{x}(T) \rangle$ (i.e., the separable, linear objective). Define

$$h_i(t) := \langle \vec{Q}_j, \vec{p} \rangle \langle \vec{Q}_j, \vec{B}_i \rangle e^{-\xi_j(t-T)}$$

. For the concave $c_i(\cdot)$ case, the necessary optimality condition becomes:

$$u_i^*(t) = \begin{cases} u_i^{\max}, & \text{if } h_i(t) > \beta^*(T) \frac{c_i(u_i^{\max})}{u_i^{\max}}, \\ 0, & \text{if } h_i(t) < \beta^*(T) \frac{c_i(u_i^{\max})}{u_i^{\max}}, \\ ?, & \text{if } h_i(t) = \beta^*(T) \frac{c_i(u_i^{\max})}{u_i^{\max}}. \end{cases} \quad (10)$$

¹⁴Singular arcs are those for which control does not affect the Hamiltonian, and thus the Hamiltonian-maximizing condition of the maximum principle is silent on the optimal value of the control over that arc.

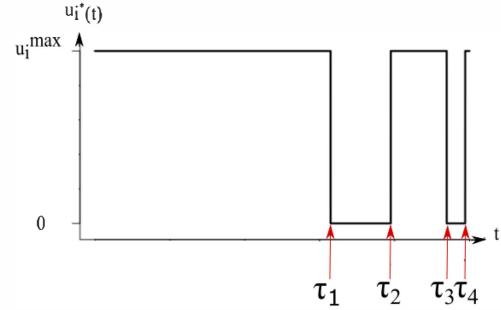


Fig. 4: For a $c_i(\cdot)$ and \vec{B}_i fulfilling the conditions of Theorem 1, the optimal control $u_i(t)$ will be bang-bang, only taking its minimum or maximum values and switching between them a bounded number of times. Note that the function can be fully described by the set of switching times $\{\tau_i\}_i$, making them easier to compute, store, and implement.

and for the linear $c_i(\cdot)$ case:

$$u_i^*(t) = \begin{cases} u_i^{\max}(t), & \text{if } h_i(t) > \beta^*(T)v_i, \\ 0, & \text{if } h_i(t) < \beta^*(T)v_i, \\ ?, & \text{if } h_i(t) = \beta^*(T)v_i. \end{cases} \quad (11)$$

In these equations, $\beta^*(T)$ is a positive *costate* variable that is associated with the budget constraint. In (10) and (11), it is the only unknown parameter (as all the elements in $h_i(t)$ are explicitly computable without solving the optimal control problem). Thus, determining $\beta^*(T)$ will determine $\vec{u}(t)$ for all t except for a finite, explicitly bounded number of points. As proved in [40], $\beta^*(T)$ has a monotonic relationship with the total budget spent, leading to the following theorem:

Theorem 2. For the case of separable, linear $J_i(\cdot)$ functions (5) for all i (i.e., $J(\vec{x}(T)) = \langle \vec{p}, \vec{x}(T) \rangle$) the unique optimal control can be explicitly calculated using a number of evaluations of (??) that is logarithmic in the range of considered $\hat{\beta}(T)$'s.

This is because the process outlined above results in a simple algorithm to find $\beta^*(T)$ and solve the optimal control problem using a single-shooting approach (which is akin to Newton's method of finding a zero of a function [38, page 484]). $\hat{\beta}(T)$ is adjusted so as to find the root of $\int_0^T \sum_{i=1}^m c_i(u_i(t, \hat{\beta}(T))) dt = r$. This significantly decreases the complexity of calculating the optimal control, as instead of evaluating and comparing potential optimal solutions that fulfill the necessary conditions in Theorem 1, one can simply evaluate $\int_0^T \sum_{i=1}^m c_i(u_i(t, \hat{\beta}(T))) dt$ using (??) and (??) over a number of iterations that is logarithmic in the range of $\hat{\beta}(T)$ under consideration to explicitly characterize the unique optimal control.

The procedure outlined above is also instructive in understanding the relative importance of different channels at different times graphically. In particular, we will be interested in comparing $\frac{u_i^{\max} h_i(t)}{c_i(u_i^{\max})}$ for concave $c(\cdot)$ and $\frac{h_i(t)}{v_i}$ for linear

$c(\cdot)$ with $\hat{\beta}(T)$ (as in (??) and (??)). One can think of the terms containing $h_i(t)$ as a topographic relief map, signifying hills and valleys. $\hat{\beta}(T)$ signifies a water-line, below which the valleys are flooded. The budget expenditure in this case is a monotone function of the area above water (see Fig. 5). Therefore, the algorithm outlined is equivalent to adjusting the water-line so that the budget expenditure (evaluated as a function of the land above water) matches the budget constraint.

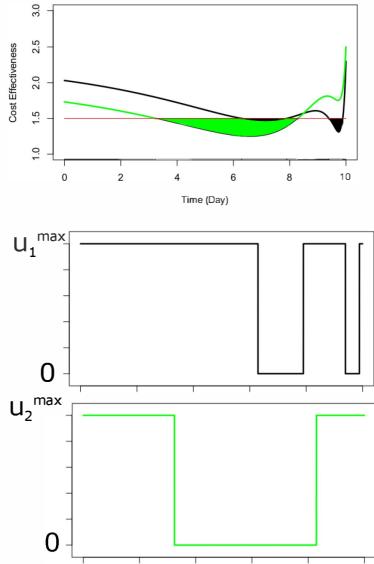


Fig. 5: Here, we demonstrate a case with two $\frac{h_i(t)}{v_i}$ functions for the case of linear $c_i(\cdot)$. Areas above the water-line ($\hat{\beta}(T)$) translate to $u_i(t, \hat{\beta}(T)) = u_i^{max}(t)$, while those below translate to $u_i(t, \hat{\beta}(T)) = 0$. The amount of budget spent for this $\hat{\beta}(T)$ can thus be calculated from the resulting $u_i(t, \hat{\beta}(T))$, and so $\hat{\beta}(T)$ can be adjusted to find $\beta^*(T)$.

Furthermore, the water-filling procedure shows the relative importance of the use of influence on any time and any channel. As the optimal water-level is a monotone decreasing function of the budget available, one can see that the peaks in $u_i^{max} h_i(t)$ and $\frac{h_i(t)}{v_i}$ signify the time intervals and channels that would be prioritized when the budget is tight, while if the budget is increased, more and more channels will be utilized at an increasing set of intervals. This lends itself nicely to considering these explicitly computable values as a *direct metric* for the *effect of advertising on any channel at a given time on the outcome of the election*, which we shall henceforth call **cost-effectiveness** of a channel.¹⁵

One can extract some more insight from the structure of this metric to compare the relative importance of channels at a given time t :

- 1) If $|T - t| \gg 1$ (i.e., early on), the deciding factor in comparing the cost-effectiveness of channels is their **total**

reach (across all nodes) per unit cost; for example, for the linear $c(\cdot)$ case:

$$\frac{h_i(t)}{v_i} \approx \left(\frac{1}{N} \sum_{j=1}^N p_j \right) \left(\frac{\sum_{j=1}^N b_{ji}}{v_i} \right),$$

- 2) However, if $|T - t| \ll 1$ (i.e., late on), **targeting** (e.g., how well a channel is aligned with the *a priori* likelihood of people to vote) is more important than total reach; for example, for the linear $c(\cdot)$ case:

$$\frac{h_i(t)}{v_i} \approx \frac{\langle \vec{p}, \vec{B}_i \rangle}{v_i},$$

This is instructive, as it shows that at the start of a campaign, cheap broadcast methods would be preferable to costly (premature) targeting of likely voters, while as election day approaches, the alignment of a channel with the target demographic gradually increases in importance.

E. Separable Sigmoid Objective

In this case, if we define:

$$Y_i(x_i(T)) = \frac{p_i \alpha_i}{(e^{\frac{\alpha_i}{2}(x_i(T) - \theta_i)} + e^{-\frac{\alpha_i}{2}(x_i(T) - \theta_i)})^2},$$

then we will have the optimality conditions (10) and (11) for:

$$h_i(t) := \langle \vec{Q}_j, \vec{Y} \rangle \langle \vec{Q}_j, \vec{B}_i \rangle e^{-\xi_j(t-T)},$$

when \vec{Y} is the vector of Y_i 's. The further away $x_i^*(T)$ is from τ_i (i.e., the farther they are from changing their mind, or alternatively the more convinced they are), the smaller the relevant $Y_i(T)$. For a given $\epsilon \ll 1$, define the set of **late-deciders** [17] under the optimal advertising action \vec{u}^* to be $\mathcal{L} := \{j \mid |x_i^*(T) - \theta_j| < \epsilon\}$. When $\mathcal{L} \neq \emptyset$ (i.e., there are late deciders), we can use the water-filling machinery in §IV-D with the changes outlined below to approximate the cost-effectiveness of channels and to calculate the optimal allocation using the much faster method described there-in.

We define $\vec{\bar{Y}}$ such that:

$$\bar{Y}_j := \begin{cases} 0 & j \notin \mathcal{L}, \\ \frac{\alpha_j p_j}{2} & j \in \mathcal{L} \end{cases}$$

Then, the approximate cost-effectiveness metric of channel i with linear $c(\cdot)$ becomes:

$$\frac{h_i(t)}{v_i} = \frac{\langle \vec{Q}_j, \vec{\bar{Y}} \rangle \langle \vec{Q}_j, \vec{B}_i \rangle e^{-\xi_j(t-T)}}{v_i}$$

This confirms the practical intuition that identifying the people who will decide late in the campaign can delineate the whole trajectory of the campaign.

¹⁵This persists in the presence of any type of noise in the state dynamics.

V. SIMULATION STUDIES

In this section, we look at a simple example to show that even in small networks, the apportioning of resources across channels can have complicated structures and can be counter-intuitive. Furthermore, we show that in many cases, the bound derived from Theorem 1 grows much slower than the number of agents (N).

We first examine a network of 7 agents with linear objectives, with $\vec{p} = (3\%, 2\%, 10\%, 100\%, 6\%, 7\%, 1\%)$. Note that under these conditions, agent 4 is the only reliable voter, with all other agents having small probabilities of voting. The connections within the network are represented in Fig. 6; the off-diagonal elements of the Laplacian \mathbf{L} are such that $l_{ij} = 1$ if there is an edge in the figure between nodes i and j and zero otherwise. Assume two equal (linear) cost channels are available to the advertiser: Channel 1, $\vec{B}_1 = (1, -1, 1, 0, 1, 0, 0)^T$ has a positive impact on agents 1, 3, and 5, but a negative impact on agent 2. It has no effect on the likely voter, agent 4. In contrast, channel 2, $\vec{B}_2 = (-1, 1, 0, 1, 0, 0, 0)^T$, has a positive effect on the likely voter, but it has more limited effects on the rest of the graph. We solve this allocation problem in Fig. 7 using the waterfilling methodology of §IV-D. As conjectured, at times $t \ll T = 10$, the cost-effectiveness of the two equal cost channels is measured by $h_i(t) \approx (\frac{1}{N} \sum_{j=1}^N p_j)(\sum_{j=1}^N b_{ji})$, which is larger for channel 1, even though it does not do a good job of targeting the likely voter. However, for t close to T , we can see that the cost-effectiveness ranking is flipped, as now $h_i(t) \approx \langle \vec{p}, \vec{B}_i \rangle$, leading to the primacy of targeting. Note that the optimal control is bang-bang with bounded numbers of transitions, as proven in Theorem 1.

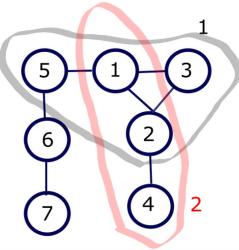


Fig. 6: We examine a network of 7 agents with 2 influence channels. The lines in solid black represent the underlying communication network \mathbf{L} . The grey and red box delineate the two channels that are available for influence in terms of agents affected (but not intensity).

One important question, especially from a computational point of view, is how tight the upper-bounds on the number of switches are. The most general bound (Theorem 2a) grows with the number of agents in the system, potentially leading to a large computational burden. On the other hand, knowing $\Lambda^*(T)$ will allow us to use tighter bounds, like that in Theorem 2b. We simulated 1000 random geometric graphs with a correspondingly uniformly random linear objective function for the case of $\vec{B}_1 = (1, 0_{1*(N-1)})^T$, and plotted the mean and variance of the bound in Theorem 2b as the number of

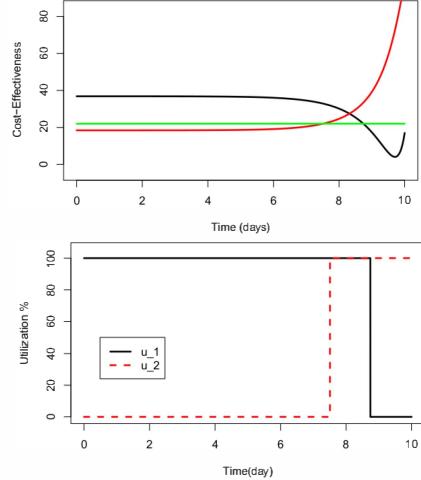


Fig. 7: We plot the cost-effectiveness of these two channels over a time horizon of $T = 10$ days when they have equal cost ($v_1 = v_2 = 1$). We then derive the water-level $\hat{\beta}$ for $r = 11$ (in green). This determines the optimal utilization rate of the two channels at different times. As can be seen, the channel with the most reach (channel 1) is prioritized at small t , and the one that is most aligned with the likelihood to vote (channel 2) is prioritized late as the election draws near.

agents was varied. As can be seen in Fig. 8, this latter bound is much smaller (around 10 for 200 agents), and its growth with respect to the number of agents is negligible. This is significant because it means that from an applied perspective, the advertiser can enumerate and evaluate a much smaller set of candidate optimals, and yet can be reasonably sure that the best such policy is globally optimal.

VI. SUMMARY

We showed that the optimal influence strategy each channel in an explicitly bounded number of waves, then ceases effort to let influence propagate (it is a bang-bang vector control). This bound is less than the number of agents, and the difference between these two values can be significant. Furthermore, we showed that the exact optimal control can be calculated using a water-filling procedure for a linear objective. These methods simplify the search for optimal controls, which in general would be on the space of vector functionals. From the water-filling procedure, we rigorously defined “cost-effectiveness” a metric for ranking and comparing the influence of different channels at differing times on outcomes.

Applying our results to the sigmoid approximation of the electoral campaign model confirmed the intuitive notion that identifying last-deciders determines the campaign strategy.

REFERENCES

- [1] A. S. Gerber, D. Karlan, and D. Bergan, “Does the media matter? a field experiment measuring the effect of newspapers on voting behavior and political opinions,” *American Economic Journal: Applied Economics*, vol. 1, no. 2, pp. 35–52, 2009.

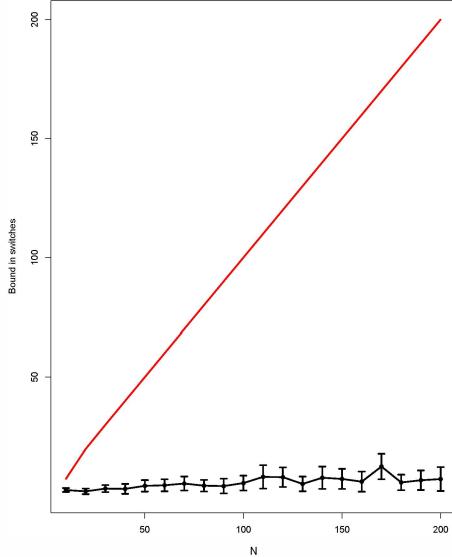


Fig. 8: We plot the upper-bounds on the number of switches of the optimal resource allocation derived from Theorem 1 for 1000 random geometric graphs for a channel that only affects the first agent as the size of the network is varied. The bound in red is from Theorem 2a, and can be seen to grow with the number of agents. For a linear objective (Theorem 2b), the black line shows that the bound on the number of switches is significantly smaller, and does not seem to grow with the size of the network.

- [2] M. Brodie, C. Deane, and S. Cho, “Regional variations in public opinion on the affordable care act,” *Journal of health politics, policy and law*, vol. 36, no. 6, pp. 1097–1103, 2011.
- [3] J. A. Chevalier and D. Mayzlin, “The effect of word of mouth on sales: Online book reviews,” *Journal of marketing research*, vol. 43, no. 3, pp. 345–354, 2006.
- [4] C. Campbell, J. Cohen, and J. Ma, “Advertisements just aren’t advertisements anymore:: a new typology for evolving forms of online advertising”, *Journal of Advertising Research*, vol. 54, no. 1, pp. 7–10, 2014.
- [5] J. Pfeffer, T. Zorbach, and K. M. Carley, “Understanding online firestorms: Negative word-of-mouth dynamics in social media networks,” *Journal of Marketing Communications*, vol. 20, no. 1-2, pp. 117–128, 2014.
- [6] S. Valenzuela, “Unpacking the use of social media for protest behavior the roles of information, opinion expression, and activism,” *American Behavioral Scientist*, vol. 57, no. 7, pp. 920–942, 2013.
- [7] D. J. Chung and Z. Lingling, “The air war versus the ground game: An analysis of multi-channel marketing in us presidential elections,” 2014.
- [8] D. Kurtzleben, “2016 campaigns will spend \$4.4 billion on tv ads, but why?” 2015. [Online]. Available: www.npr.org/sections/itsallpolitics/2015/08/19/43275931/2016-campaign-tv-ad-spending
- [9] M. Bagchi, S. Murdoch, and J. Scanlan, “The state of global media spending,” *McKinsey Global Media Report 2015*, 2015.
- [10] R. Hegselmann and U. Krause, “Opinion dynamics and bounded confidence models, analysis, and simulation,” *Journal of Artificial Societies and Social Simulation*, vol. 5, no. 3, 2002.
- [11] P. Baines, E. K. Macdonald, H. Wilson, and F. Blades, “Measuring communication channel experiences and their influence on voting in the 2010 british general election,” *Journal of Marketing Management*, vol. 27, no. 7-8, pp. 691–717, 2011.
- [12] D. M. Kahan, E. Peters, M. Wittlin, P. Slovic, L. L. Ouellette, D. Braman, and G. Mandel, “The polarizing impact of science literacy and numeracy on perceived climate change risks,” *Nature Climate Change*, vol. 2, no. 10, pp. 732–735, 2012.
- [13] R. Olfati-Saber, J. A. Fax, and R. M. Murray, “Consensus and cooperation in networked multi-agent systems,” *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.
- [14] D. P. Spanos, R. Olfati-Saber, and R. M. Murray, “Dynamic consensus on mobile networks,” in *IFAC world congress*. Citeseer, 2005, pp. 1–6.
- [15] K. Kandhway and J. Kuri, “Campaigning in heterogeneous social networks: Optimal control of si information epidemics,” *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 383–396, 2016.
- [16] P. A. Naik and K. Raman, “Understanding the impact of synergy in multimedia communications,” *Journal of Marketing Research*, vol. 40, no. 4, pp. 375–388, 2003.
- [17] B. C. Hayes and I. McAllister, “Marketing politics to voters: Late deciders in the 1992 british election,” *European Journal of Marketing*, vol. 30, no. 10/11, pp. 127–139, 1996.
- [18] M. DeGroot, “Reaching a consensus,” *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [19] R. Olfati-Saber, “Flocking for multi-agent dynamic systems: Algorithms and theory,” *Automatic Control, IEEE Transactions on*, vol. 51, no. 3, pp. 401–420, 2006.
- [20] W. Ren, R. W. Beard, and E. M. Atkins, “A survey of consensus problems in multi-agent coordination,” in *Proceedings of the 2005 American Control Conference, 2005*. IEEE, 2005, pp. 1859–1864.
- [21] E. Yıldız, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione, “Discrete opinion dynamics with stubborn agents,” Available at SSRN 1744113, 2011.
- [22] D. Kempe, J. Kleinberg, and É. Tardos, “Influential nodes in a diffusion model for social networks,” in *Automata, languages and programming*. Springer, 2005, pp. 1127–1138.
- [23] E. Anshelevich, D. Chakrabarty, A. Hate, and C. Swamy, “Approximation algorithms for the firefighter problem: Cuts over time and submodularity,” in *Algorithms and Computation*. Springer, 2009, pp. 974–983.
- [24] M. Athans and P. L. Falb, *Optimal control: an introduction to the theory and its applications*. Courier Corporation, 1966.
- [25] D. Liberzon, *Calculus of variations and optimal control theory: a concise introduction*. Princeton University Press, 2012.
- [26] A. Seierstad and K. Sydsæter, *Optimal control theory with economic applications*. Elsevier North-Holland, Inc., 1986.
- [27] S. Eshghi, M. Khouzani, S. Sarkar, N. B. Shroff, and S. S. Venkatesh, “Optimal energy-aware epidemic routing in dtns,” *Automatic Control, IEEE Transactions on*, vol. 60, no. 6, pp. 1554–1569, 2015.
- [28] 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.
- [29] S. Eshghi, M. Khouzani, S. Sarkar, and S. S. Venkatesh, “Optimal patching in clustered malware epidemics,” *IEEE/ACM Transactions on Networking (TON)*, vol. 24, no. 1, pp. 283–298, 2016.
- [30] A. A. Cárdenas, S. Amin, and S. Sastry, “Research challenges for the security of control systems,” in *HotSec*, 2008.
- [31] Y. Mo, E. Garone, A. Casavola, and B. Sinopoli, “False data injection attacks against state estimation in wireless sensor networks,” in *Decision and Control (CDC), 2010 49th IEEE Conference on*. IEEE, 2010, pp. 5967–5972.
- [32] R. P. Abelson, “Mathematical models of the distribution of attitudes under controversy,” *Contributions to mathematical psychology*, vol. 14, pp. 1–160, 1964.
- [33] M. Taylor, “Towards a mathematical theory of influence and attitude change,” *Human Relations*, vol. 21, no. 2, pp. 121–139, 1968.
- [34] P. V. Marsden and N. E. Friedkin, “Network studies of social influence,” *Sociological Methods & Research*, vol. 22, no. 1, pp. 127–151, 1993.
- [35] R. Olfati-Saber and R. M. Murray, “Consensus problems in networks of agents with switching topology and time-delays,” *IEEE Transactions on automatic control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [36] N. E. Friedkin and E. C. Johnsen, “Social influence and opinions,” *Journal of Mathematical Sociology*, vol. 15, no. 3-4, pp. 193–206, 1990.
- [37] F. R. Chung, *Spectral graph theory*. American Mathematical Soc., 1997, vol. 92.
- [38] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [39] L. Hogben, *Handbook of linear algebra*. CRC Press, 2006.
- [40] S. Eshghi, V. Preciado, S. Sarkar, 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*, 2017.
- [41] C.-T. Chen, *Linear system theory and design*. Oxford University Press, Inc., 1995.