

Spread, Then Target, and Advertise in Waves: Optimal Budget Allocation Across Advertising Channels

Soheil Eshghi¹, Member, IEEE, Victor M. Preciado², Member, IEEE, Saswati Sarkar,
Santosh S. Venkatesh³, Member, IEEE, Qing Zhao⁴, Fellow, IEEE,
Raissa D'Souza, and Ananthram Swami, Fellow, IEEE

Abstract—We analyze optimal strategies for the allocation of a finite budget that can be invested in different advertising channels over time with the objective of influencing social opinions in a network of individuals. In our analysis, we consider both exogenous influence mechanisms, such as advertising campaigns, as well as endogenous mechanisms of social influence, such as word-of-mouth and peer-pressure, which are modeled using diffusion dynamics. We show that for a broad family of objective functions, the optimal influence strategy at every time uses all channels at either their maximum rate or not at all, i.e., a bang-bang strategy. Furthermore, we prove that the number of switches between these extremes is bounded above 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 of the campaign, the total cost-adjusted reach of an exogenous advertising channel determines its relative value. In contrast, as we approach our investment horizon (e.g., election day), the optimal strategy is to invest in channels able to target individuals instead of broad-reaching channels. We demonstrate that the optimal influence strategies are easily computable in several practical cases, and explicitly characterize the optimal controls for the case of linear

objective functions in closed form. Finally, we see that, in the canonical example of designing an election campaign, identifying late-deciders is a critical component in the optimal design.

Index Terms—Marketing management, optimal control, bang-bang control, election, advertising.

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 formation processes. The opinion formation 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], though the extent of the effect of these efforts is a matter of debate [1], [8]. The importance of this decision has increased in conjunction with the increasing resources devoted to these efforts: In 2017, over \$1 trillion was spent on marketing globally [9], while \$9.8 billion was spent on advertising in the 2016 US elections alone [10]. Thus, studying the related multi-channel resource allocation problem is both timely and significant.

In particular, the mechanisms of opinion influence can be classified into two types based on its direct provenance. First, there are *endogenous* influence mechanisms (e.g., word-of-mouth), in which 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 capturing endogenous influence between individuals. On the other hand, there are *exogenous* influence mechanisms, in which an external influencer seeks to shape the opinions of an individual. This

1. These weights can, in general, be dynamic, even depending on the expressed opinion [11]. In this work, we consider static weights.

Manuscript received January 31, 2018; revised July 23, 2018; accepted August 23, 2018. Date of publication October 2, 2018; date of current version June 4, 2020. The work of S. Eshghi, Q. Zhao, R. D'Souza, and A. Swami was supported in part by the Army Research Laboratory Network Science CTA under Cooperative Agreement W911NF-09-2-0053. The work of V. M. Preciado was supported in part by the U.S. National Science Foundation under Grant CAREER-ECCS-1651433. The work of S. Sarkar was supported in part by the U.S. National Science Foundation under Grant CNS-1525457. This article was presented in part at the IEEE Information Theory and Applications Workshop (ITA), San Diego, CA, USA, February 2017, and in part as a poster at the Yale Day of Data, New Haven, CT, USA, December 2017. Recommended for acceptance by X. Li. (*Corresponding author: Soheil Eshghi.*)

Soheil Eshghi is with the School of Electrical Engineering, Yale Institute of Network Science (YINS), Yale University, New Haven, CT 06520 USA (e-mail: soheil.eshghi@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 USA (e-mail: preciado@seas.upenn.edu; swati@seas.upenn.edu; venkates@seas.upenn.edu).

Qing Zhao is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853 (e-mail: qz16@cornell.edu).

Raissa D'Souza is with the Departments of Computer Science & Mechanical and Aerospace Engineering, University of California, Davis, Davis, CA 95616 USA (e-mail: raissa@cse.ucdavis.edu).

Ananthram Swami is with the Army Research Laboratory, Adelphi, MD 20783 USA (e-mail: ananthram.swami.civ@mail.mil).

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Digital Object Identifier 10.1109/TNSE.2018.2873281

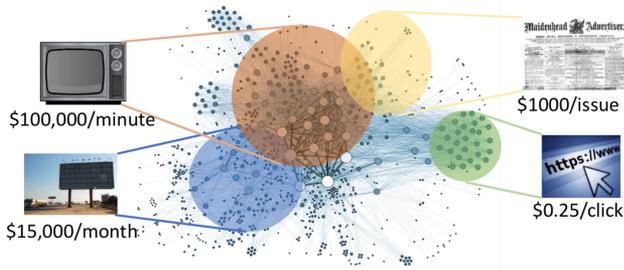


Fig. 1. Different advertising channels have differing, possibly overlapping, reaches and come with differing costs. Furthermore, the effect of a channel may differ across individuals. These external influences are modulated by internal conversations within the network whereby agents integrate this information with that of their neighbors. The decision of the advertiser 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.

mechanism is facilitated by various advertising channels.² In our opinion formation model, each channel has a reach structure, i.e., the individuals that can be reached by that channel, that is not necessarily related to the endogenous neighborhood. The actions of other external influencers also affect each individual's opinion formation process, which can in general be a random (noisy) process. The external influencer seeks to maximize a function of the *global state* (the vector of individual opinions) at a specific time (e.g., election day) by allocating their budget across several advertising channels in a given time interval (see Fig. 1). In this paper, we study the nature of the optimal budget allocation and provide structures and algorithms for their computation.

Finding the optimal budget allocation is complicated by several factors: (i) The reach of each channel is limited, and there are significant overlaps between the target audiences of various channels [12]; (ii) different channels have differing costs, and attempts to influence opinions by external sources can affect individuals in different, and sometimes opposite, ways [13]; (iii) the budget allocation decision is dynamic (depends on time) and changes with the state of the network. Furthermore, the influencer faces several trade-offs: utilizing an advertising channel early 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 the optimal budget allocation hard to determine.

There are also significant technical challenges to solving this problem, since characterizing the optimal budget allocation across channels and throughout the time interval requires characterizing the structure of an optimal constrained vector of controls over a graph. Furthermore, the work also requires computing the optimal control of the well-studied linear consensus dynamics [14], [15] in a novel setting, as the classical literature is

2. Throughout 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.

concerned with reaching agreement among agents, while our objectives may incentivize agreement in some circumstances and disagreement in others. As we show in this paper, finding the optimal allocation in our problem 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 consensus control process in an arbitrary network with overlapping influence channels and endogenous influence of agents on each other [16]. Using Pontryagin's maximum principle, spectral graph theory, and custom analytical arguments, we determine the structure of the optimal budget allocation to the various influence channels along a given time horizon.

We show that for a broad family of objectives, the optimal control for each channel is bang-bang (only takes its extreme values), with the number of switches being upper-bounded by a term which is smaller than the number of individuals. Therefore, the search for optimal controls can be conducted on the space of vectors of a fixed size whose entries represent times of switching between extreme values 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 value), we explicitly calculate the optimal budget allocation over time, providing an open-loop algorithm that can compute the vector of optimal controls in a logarithmic number of steps. This allocation also implicitly determines the relative importance 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. This metric allows the influencer to compare and contrast the effects of different channels, as well as the effect of a channel at different times. Finally, our results show that investing in an influence channel reaching likely voters is important as we get *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, 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, our 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 budget allocation for the multi-channel 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 and Opinion Dynamics: Linear consensus-seeking dynamics are some of the oldest models used to model the spread of opinions and social influence, first proposed by French [18] and expounded upon by DeGroot [19]. In these models, opinions (states) are taken to be continuous real variables, and each node uses a (weighted) average opinion of its neighbors' opinions in each time-step to update its opinion.

Abelson [20] provided a continuous-time variant of these dynamics, which is the model we base our work upon. This work was generalized first by Taylor [21] to also incorporate individual-specific prejudices, leading to the desirable persistence of opinion cleavage within such simple models (leading to a dynamics very similar to ours). Other closely related continuous-time variants of these dynamics have been rigorously studied by control theorists [14], [15], [22], [23]. 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. While more complex models of opinion dynamics have been proposed and studied in detail [11], [24], [25], the linear consensus dynamics remains a baseline for comparison. Recent detailed overviews of the developments in the field of opinion dynamics make the above distinctions and limitations clearer [26], [27]. Finally, the linear approximation of the effect of external influence on opinion dynamics also follows a long-standing tradition [21], [24]. Our work covers finite-time budget-constrained opinion change with a specific goal, while the focus of these papers is understanding asymptotic properties of these systems (without strategic interventions and goals).

Control of Opinion Propagation: The case of influencing opinion dynamics is a research question of current interest. The problem of Influence Maximization (IM) consists of finding the set of individuals that must be initially influenced in order to maximize the final effect of endogenous spreading mechanisms [28]. Variants of this problem, under multiple models of opinion propagation, have been the subject of much study (e.g., [29]–[32]). Among this line of work, budgeted influence maximization with partial incentives [33] is the closest to our setting, as it relaxes the artificial binary assumptions on the success of influence efforts. While this literature is closely related to work on epidemic control [32], its more immediate analog is work on control of social learning. For example, Yildiz et al. [34] 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. 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. In this paper, we analyze such optimal actions (henceforth referred to as controls) using tools from optimal control theory.

Linear Optimal Control: In linear optimal control problems, a controller seeks to optimize the time integration of a linear objective depending on the states and inputs of a linear dynamical system with linear bounded controls.

In the case where actions are not costly and the time horizon is not fixed, the optimal control signal has a bang-bang structure with a finite number of switches [35]–[37]. 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. In contrast, our work takes a

step beyond those results and provides a context-specific method for evaluating the relative influence value on a channel within a time horizon.

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 [38], [39]) 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 strategies derived for information spread (which typically advocate some form of maximal spreading at the start of the time interval [39], [40]). The model also allows an even more explicit incorporation of graph structure than metapopulation models, e.g., [41], 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 [42]. In this setting, a state-estimation sensor network [43] can be misdirected through local noise injection at a fixed number of points, that will affect 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. The modeling approach employed in our work is, to the best of our knowledge, novel for this setting.

In summary, our 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, we present our notation (Section III) and outline our system model (Section III) and its dynamics (Section III). Then we outline the bounds on the actions of the influencer (Section III) and describe their objective (Section III). We finish the section by presenting a technical assumption (Section III) and by stating the overall problem (Section III).

A. Notation

- n = number of agents
- m = number of channels
- $x_i(t)$ = opinion of agent i at time t , $i = 1, \dots, n$
- $u_k(t)$ = utilization of channel k at time t , $k = 1, \dots, m$
- $u_k^{\max}(t)$ = maximum utilization of channel k at time t

- a_{ij} = magnitude of effect agent j 's opinion on the opinion of agent i
- N_i = neighbors of agent i in communication graph G
- H_k = set of agents within the reach of channel k
- b_{ik} = relative magnitude of the effect of utilization of channel k on agent i
- $e_i(t)$ = sum effect of other influences on agent i at time t
- T = terminal time
- $c_k(\cdot)$ = cost of influence on channel k
- r = total resources of influencer over the time period
- \mathcal{U} = set of feasible influence allocations
- $J_i(\cdot)$ = value of opinion of agent i at time T to influencer
- $J(\cdot)$ = value of opinion profile at time T to influencer

We use bold lower case letters to denote vectors and bold upper case letters to denote matrices, $[n]$ to represent $\{1, 2, \dots, n\}$, and $\langle a, b \rangle$ to represent $\mathbf{a}^T \mathbf{b}$. For a matrix \mathbf{W} , we denote the k th column of \mathbf{W} as $\mathbf{W}(:, k)$, and the k th row of the same as $\mathbf{W}(k, :)$. Furthermore, we use w_{ij} to denote the (i, j) th element of the matrix \mathbf{W} .

B. System Model

We consider a social system with n agents. The *state/opinion* of agent $i \in [n]$ at time t is denoted by $x_i(t) \in \mathbb{R}$. Each agent communicates with other agents based on an *edge-weighted, undirected, and connected* communication graph $G = (V, E, \mathbf{A})$. The (non-negative) weight on an edge between agents $i, j \in [n]$, which determines the relative influence agent j has on agent i 's 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 $a_{ij} = a_{ji} > 0$ (see Fig. 2).³ The assumption of symmetric interaction weights, that a difference of opinion can have the same magnitude effect on both sides of an interaction, is common in the Influence Maximization literature, most commonly being present in the Independent Cascade (IC) model [28], [44], [45].

Remark 1. However, weighted *directed* communication graphs can also be considered in our framework, in which case some of our results apply to cases where the weighted Laplacian of the graph has real eigenvalues.⁴ In particular, this includes the set of quasi-strongly connected weighted Directed Acyclic Graphs (DAGs).

At each time t , each agent updates its state based on a *weighted average* of the difference of its current state from those of its neighbors, as well as on an external 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 $i \in [n]$.

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

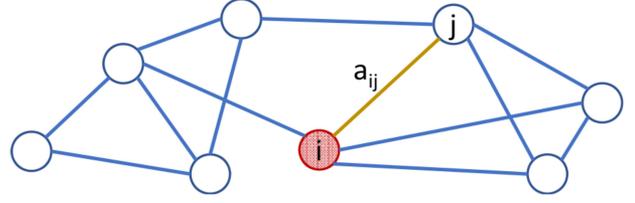


Fig. 2. Each agent takes into account the opinions of its neighbors in updating its own opinion. The weight given by node i to the opinion of its neighbor j is a measure of how much i trusts j 's appraisal. If agents i and j are not neighbors, $a_{ij} = 0$ by default. In our model, $a_{ij} = a_{ji}$, i.e., trust is a symmetric relation, though our results apply even if that is not the case (see conditions in Remark 1).

of agents (denoted H_k for channel k). The structure of these m influence channels is pre-specified, with the assumption that influencing a channel only directly affects the members within that channel. The influence exerted by the influencer on channel $i \in [m]$ at time t is denoted by the scalar $u_i(t)$.

In this model, the effect of influence on a channel can differ across agents within the channel, potentially even having opposite effects. These effects are captured by the *influence gain*, denoted by b_{ik} , which determines the linear relative gain of influence of channel $k \in [m]$ on agent i within that channel. For example, if billboard advertising (say, channel k) has a more positive effect on the opinion of individual i than advertising on the radio (say, channel l), we will have $b_{ik} > b_{il}$.⁵ If agent i is not within channel k , we define b_{ik} to be zero. Without loss of generality, we assume $u_i \geq 0$ for $i \in [m]$, and encode the possible negative effects of channel i on agent k within its reach through the sign of b_{ik} . Stacking these values into a matrix $\mathbf{B}_{n \times m}$ captures the structure of the channels.

C. Dynamics

To understand the dynamics, we provide the following discrete-time intuition: an agent $i \in [n]$ 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 external 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 may act as a hindrance to that process. Mathematically, it can be thought of as a gradient descent algorithm implemented by agents seeking to minimize disagreement (measured by a Laplacian potential) [22]. The above can be re-written to represent the classic discrete-time consensus model [19] with influence:⁶

³ While negative weight updates are conceivable, they will not be considered in this paper.

⁴ Lemma 1 and Theorem 1.1 carry over, as does a modification of the water-filling procedure in Section IV. For specifics, see [46, Section VI].

⁵ Note that the magnitude of b_{ik} can be determined by comparing the size of the effect of channel k on individual i 's opinion with that of one of i 's neighbors having the same amount of difference in opinion with i . At scale, these orderings and values can possibly be inferred from demographic information.

⁶ The results derived in this paper would also apply to the Friedkin and Johnsen model of opinion updates [24] given uniform susceptibility to change across agents.

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

Note also that, from this formulation, it is evident that the influencer's effect on the state of any individual is dissipative, as their prior state is discounted by a factor of $(1 - \Delta \sum_{j \in N_i} a_{ij}) < 1$ at every time-step.⁷

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:⁸

$$\dot{x}_i(t) = \sum_{j \in N_i} a_{ij} (x_j(t) - x_i(t)) + \sum_{k: i \in H_k} b_{ik} u_k(t) + e_i(t).$$

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 the n equations, we arrive at the following system-level dynamics:

$$\dot{\mathbf{x}}(t) = -\mathbf{L}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) + \mathbf{e}(t). \quad (1)$$

We assume that the states/opinions at time 0 are known ($\mathbf{x}(0) = \mathbf{x}_0$), 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 $r > 0$, which is the budget available to the influencer. This is captured through the following budget constraint:

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

where $c_k(\cdot)$ represents the time-independent cost-function which maps the utilization of channel k to its associated cost to the influencer.

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

This assumption models the diminishing cost of additional utilization of a channel once it is already in use. The above assumption allows the case where c_k is linear (that is $c_k(x) = v_k x$). 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 $u_k^{\max}(t)$.⁹ This can

capture both physical limits on the influence (i.e., availability of media) and limits on the susceptibility of agents to the influence. We impose the modest assumption that $u_k^{\max}(t)$ is differentiable. Hence, we have that:

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

We will restrict our analysis to control signals \mathbf{u} that are piecewise continuous with only a finite number of discontinuities.¹⁰ We shall use \mathcal{U} to denote the set of such controls that fulfill (3):

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

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 $k \in [m]$, $c_k : [0, +\infty) \rightarrow [0, +\infty)$. Note that we assume that for all $k \in [m]$, the cost function $c_k(\cdot)$ 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, $J(\mathbf{x}(T))$, is an increasing, differentiable function of the n components of the vector of terminal opinions $\mathbf{x}(T)$.

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

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

Two specific types of separable functions are of particular 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 opinion of a product. This is also a useful approximation in the adversarial sensor network deception 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:

7. This can be seen by looking at the explicit effect after two time-steps (i.e., time $t + 2\Delta$), where the effect of direct influence exerted at time t , $\mathbf{u}(t)$, is multiplied by $(1 - \Delta \sum_{j \in N_i} a_{ij})$: $x_i(t + 2\Delta) = \dots + \Delta \sum_{k: i \in H_k} b_{ik} u_k(t + \Delta) + (1 - \Delta \sum_{j \in N_i} a_{ij}) \Delta \sum_{k: i \in H_k} b_{ik} u_k(t) + \Delta^2 (\sum_{j \in N_i} \sum_{m: j \in H_m} a_{ij} b_{jm} u_m(t)) + \dots$

8. While discrete-time dynamics are more commonly used for the modeling of opinions, discretization is typically a simplifying assumption for analytic purposes. In this paper, we work with the continuous-time dynamics directly, which allows the use of mathematical tools new to the domain. However, all derived structures and insights for the continuous case can be discretized and applied to the discrete-time case as well.

9. This rules out impulse controls.

10. This means, in particular, that the integral in (2) is well-defined.

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 (e.g., candidates, products, policies), encoded by 0 and 1, at time T . Assume 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 an 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 , which is agent i 's vote. However, this utility is discontinuous at $x_i(T) = \theta_i$, which complicates analysis. The sigmoid function:

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

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 a technical assumption that will be needed in our arguments:

Assumption 3. There exists a $j \in [n]$ such that $\frac{\partial J(\mathbf{z})}{\partial z_j(T)} > 0$ for all $\mathbf{z} \in \mathbb{R}^n$.

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 to all controls and are thus uninformative.

G. Problem Statement

We aim to characterize the control inputs $\mathbf{u}(t)$ that maximize $J(\mathbf{x}(T))$ under the dynamics outlined in (1) and constraints (3) and (2). Mathematically, we state our problem as:

¹¹ This rules out $J(\cdot)$ functions with stationary points, i.e., those for which $\nabla_{\mathbf{z}} J(\mathbf{z}) = \mathbf{0}$ for some $\mathbf{z} \in \mathbb{R}^n$. For example, this rules out an objective which is a sum of Heaviside functions.

$$\begin{aligned} \max_{\mathbf{u} \in \mathcal{U}} \quad & J(\mathbf{x}(T)) \\ \text{s.t.} \quad & \dot{\mathbf{x}}(t) = -\mathbf{L}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) + \mathbf{e}(t), \quad \mathbf{x}(0) = \mathbf{x}_0 \in \mathbb{R}^n \\ & \int_0^T \sum_{k=1}^m c_k(u_k(t)) dt \leq r, \\ & \mathbf{e}(t) \text{ given.} \end{aligned}$$

Note that the above problem is non-convex in general owing to the potentially non-convex objective function, as well as the potentially non-convex budget constraint (when any of the $c_i(\cdot)$'s are strictly concave). In this paper, we solve it using tools from optimal control theory. It should be observed that 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 an auxiliary variable to aid the analysis. We define the auxiliary function γ such that,

$$\gamma(0) = 0, \quad \dot{\gamma}(t) = -\sum_{k=1}^m c_k(u_k(t)). \quad (7)$$

As can be seen, $\gamma(t)$ is the accumulated cost of the influence up to time t . Thus, the budget 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:

$$\begin{aligned} \max_{\mathbf{u} \in \mathcal{U}} \quad & J(\mathbf{x}(T)) \\ \text{s.t.} \quad & \dot{\mathbf{x}}(t) = -\mathbf{L}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) + \mathbf{e}(t), \quad \mathbf{x}(0) = \mathbf{x}_0 \in \mathbb{R}^n, \\ & \dot{\gamma}(t) = -\sum_{k=1}^m c_k(u_k(t)), \quad \gamma(T) \geq -r, \\ & \mathbf{e}(t) \text{ given, } \gamma(0) = 0. \end{aligned} \quad (8)$$

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 in Section IV. Then, we prove the existence of optimal controls (under some conditions) and identify their structure using our main theorem in Section IV (with proofs in Appendices A and B, respectively, which can be found at 10.1109/TNSE.2018.2873281). A refinement is presented for the case of the linear objective Section IV that allows the direct computation of the control input \mathbf{u} and shows that the optimal control is unique, while providing insights into the logic of the allocation decision. Finally, the sigmoid approximation to voting is covered in Section IV and an approximation to the optimal control is presented.

A. Preliminaries

For an *undirected, connected* graph G , the weighted Laplacian matrix \mathbf{L} is real, symmetric, and positive semi-definite; hence it has real, non-negative eigenvalues [47, page 13]. Thus, \mathbf{L} has an eigen-decomposition $\mathbf{L} = \mathbf{Q}\mathbf{\Xi}\mathbf{Q}^T$, where \mathbf{Q} is a real

orthogonal matrix whose columns are the eigenvectors of \mathbf{L} , and Ξ is the diagonal matrix of eigenvalues of \mathbf{L} [48, p. 393, Theorem 8.1.1].¹² The smallest eigenvalue of \mathbf{L} is always **zero**, its multiplicity is 1, and its associated eigenvector is $\frac{1}{\sqrt{n}}\mathbf{1}_n$, where $\mathbf{1}_n = (1, \dots, 1)^T$ (as G is connected) [47, 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} , $\mathbf{Q}(:, i)$ will be the i th eigenvector of \mathbf{L} . This means that $\mathbf{Q}(:, 1) = \frac{1}{\sqrt{n}}\mathbf{1}_n$.

We now state a lemma that shows that an optimal control \mathbf{u} for the main problem exists. We then state our main result (Theorem 1) and present a subcase where the bound can be significantly strengthened and the optimal control can be calculated in open-loop (Theorem 1). We provide proofs of these results in Appendices A and B, available in the online supplemental material, respectively.

B. Existence of Optimal Solutions and Structural Results for the Optimal Control

We prove (in Appendix A, available in the online supplemental material, that optimal controls for (8), exist when $c_i(\cdot)$'s are linear.

Lemma 1. Optimal controls for problem (8) exist for linear costs, i.e., $c_i(u_i) = v_i u_i$ for all $i \in [m]$.

We are now ready to state our main theorem. We will provide results for a large natural class of channels that we shall call *disciplined*. We first formally define the set of disciplined channels \mathcal{I} before stating the theorem:

Definition 1. The set of disciplined channels, $\mathcal{I} \subseteq [m]$, is such that for all $i \in \mathcal{I}$, one of the two following conditions holds:

- $c_i(\cdot)$ is strictly concave, $\langle \mathbf{B}(:, i)\mathbf{1}_n \rangle \neq 0$, and $u_i^{\max}(t) = u_i^{\max}$ for all t .
- $c_i(\cdot)$ is linear and the system $(\mathbf{L}, \mathbf{L}\mathbf{B}(:, i))$ is controllable [49, p. 144] (with $u_i^{\max}(t)$ being any differentiable function).

Theorem 1. For all $i \in \mathcal{I}$:

- 1) Optimal controls are *bang-bang*, taking on their maximum or minimum values at all times t (i.e., $u_i^*(t) \in \{u_i^{\max}, 0\}$).
- 2) The number of switches between these values is bounded above:
 - a) In the general case, by one less than the number of non-zero elements in $\{\langle \mathbf{Q}(:, j)\mathbf{B}(:, i) \rangle\}_{j=1}^n$.
 - b) For $J(\mathbf{x}(T)) = \langle \mathbf{p}, \mathbf{x}(T) \rangle$, by the number of sign variations in $\{\sum_{k=1}^j s_k\}_{j=1}^n$, where $s_j := \langle \mathbf{Q}(:, j)\mathbf{p} \rangle \langle \mathbf{Q}(:, j)\mathbf{B}(:, i) \rangle$.

An example of an optimal control with these characteristics is provided in Fig. 3. The proof of this theorem is presented in Appendix B, available in the online supplemental material.

This theorem means that the optimal strategy uses each channel in waves (see Fig. 3), stopping between them to let influence propagate. From a computational stand-point, this result simplifies the space of possible optimal controls for

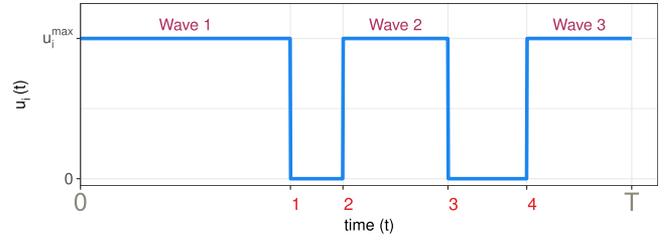


Fig. 3. For a $c_i(\cdot)$ and $\mathbf{B}(:, i)$ fulfilling the conditions of Theorem 1, the optimal control $u_i(t)$ will be bang-bang, only taking its minimum or maximum values and switching between them a bounded number of times. Thus, the function can be fully described by the set of switching times $\{\tau_i\}_i$, making them easier to compute, store, and implement.

each channel, since the optimal control is characterized by the bounded number of switching times for each channel. 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), as we will see in Section V.

Remark 2. The conditions in Definition 1 rule out pathological cases where the necessary conditions for optimality derived from the Maximum Principle [37, page 182] cannot directly determine the optimal value of the control (i.e., singular arcs [36, page 113] exist).

Remark 3. The set of disciplined channels may be a proper subset of the set of channels ($\mathcal{I} \subset [m]$), in which case the derived structure only applies to disciplined channels. This means that even if $c_k(\cdot)$ is non-concave or the conditions around \mathbf{B}_k in Definition 1 do not hold for some k , Theorem 1 will remain valid for disciplined channels. Note that optimal controls for undisciplined channels may also abide by the bang-bang structures stated in Theorem 1.

C. Water-Filling: Optimal Budget Allocation for Separable Linear Objectives

In this section, for separable linear objectives, we will derive a detailed cost-effectiveness metric for channel i 's utilization that depends on the eigenvalues and eigenvectors of the Laplacian (\mathbf{L}), the channel influence gain vector ($\mathbf{B}(:, i)$), and the weights of the linear objective (\mathbf{p}). The variation of this metric across time will result in *hills* and *valleys* that represent the variations in the effectiveness of the channel across time. Choosing a water-line for this topography (see Fig. 4), we will show, leads to the description of a candidate control which takes its maximum values when a hill is above water, and will be set to zero when a valley is under water. This waterline is varied using the bisection/binary search method so that the cost of the total area above water matches the budget constraint (2). We will further show how this approach generalizes for more varied objective functions.

From the proof of Theorem 1 in Appendix B, available in the online supplemental material, we can define:

$$h_i(t) = \langle \lambda^*(t)\mathbf{B}(:, i) \rangle \\ = \sum_{j=1}^N \langle \mathbf{Q}(:, j)\mathbf{p} \rangle \langle \mathbf{Q}(:, j)\mathbf{B}(:, i) \rangle e^{-\xi_j(T-t)}, \quad (9)$$

12. However, the reasoning below applies to any \mathbf{L} that has real eigenvalues for the case of linear costs $c_i(\cdot)$ and a quasi-strongly connected communication digraph. For details, see [46, Section VI].

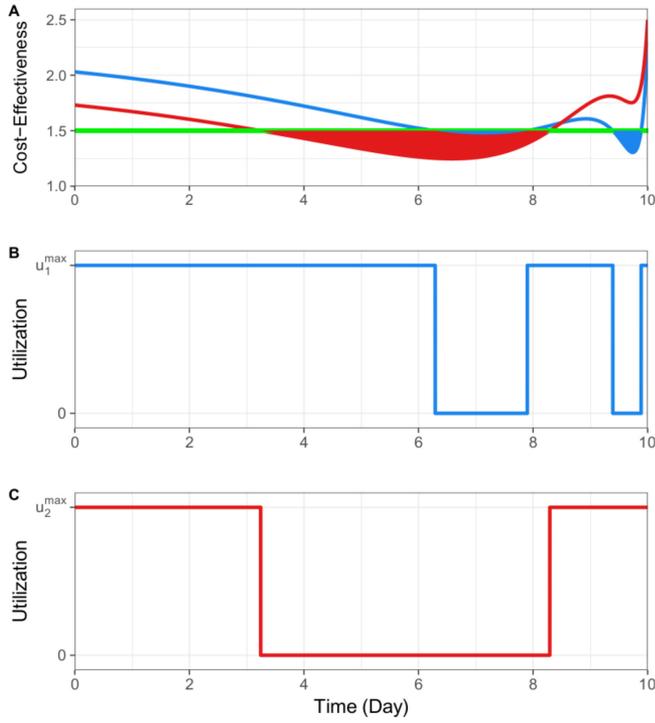


Fig. 4. We demonstrate a case with two $h_i(t)/v_i$ functions for the case of linear $c_i(\cdot)$. Areas above the water-line ($\hat{\beta}(T)$) translate to $u_i(t, \hat{\beta}(T)) = u_i^{\max}(t)$, while those below translate to $u_i(t, \hat{\beta}(T)) = 0$. The amount of budget spent for this $\hat{\beta}(T)$ can thus be calculated from the resulting $u_i(t, \hat{\beta}(T))$, and so $\hat{\beta}(T)$ can be adjusted to find $\beta^*(T)$.

where $\mathbf{Q}(:, j)$ is the j th eigenvector of the Laplacian matrix \mathbf{L} with associated eigenvalue ξ_j , and \mathbf{p} is the vector of weights of the linear objective, i.e., $J(\mathbf{x}(T)) = \langle \mathbf{p}, \mathbf{x}(T) \rangle$, such that the necessary condition for optimal controls in the concave $c_i(\cdot)$ case becomes:¹³

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

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

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

for some optimal *a priori* unknown parameter $\beta^*(T)$. All other terms in (10) and (11) are explicitly computable without solving the optimal control problem. Thus, determining $\beta^*(T)$ will determine $\mathbf{u}(t)$ for all t except for a finite, explicitly bounded number of points (notice that the existence of singular controls was ruled out in the proof of Theorem 1). However, as we shall see in (A4) of the appendix, available in the online supplemental material, $\beta^*(T) > 0$ if and only if $\int_0^T \sum_{k=1}^m c_k(u_k(t)) dt = r$. This last equation is the budget constraint.

13. The question mark denoting the fact that PMP does not uniquely determine the optimal u_i^* at times t when $\varphi_i(t, u_i)$ does not change with u_i .

Define the equivalent of (10) and (11) as functions of a variable $\hat{\beta}(T)$, an estimate for $\beta^*(T)$:

$$u_i(t, \hat{\beta}(T)) = \begin{cases} u_i^{\max}, & \text{if } h_i(t) > \hat{\beta}(T) \frac{c_i(u_i^{\max})}{u_i^{\max}}, \\ 0, & \text{if } h_i(t) < \hat{\beta}(T) \frac{c_i(u_i^{\max})}{u_i^{\max}}, \\ ?, & \text{if } h_i(t) = \hat{\beta}(T) \frac{c_i(u_i^{\max})}{u_i^{\max}}. \end{cases} \quad (12)$$

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

$$u_i(t, \hat{\beta}(T)) = \begin{cases} u_i^{\max}(t), & \text{if } h_i(t) > \hat{\beta}(T)v_i, \\ 0, & \text{if } h_i(t) < \hat{\beta}(T)v_i, \\ ?, & \text{if } h_i(t) = \hat{\beta}(T)v_i. \end{cases} \quad (13)$$

One can see that in both cases, if $\hat{\beta}_1(T) > \hat{\beta}_2(T) \geq 0$, $u_i(t, \hat{\beta}_2(T)) \geq u_i(t, \hat{\beta}_1(T))$ for all i and all t . This, along with Assumption (1), leads to $c_i(u_i(t, \hat{\beta}_2(T))) \geq c_i(u_i(t, \hat{\beta}_1(T)))$ for all i and all t , culminating in:

$$\int_0^T \sum_{i=1}^m c_i(u_i(t, \hat{\beta}_2(T))) dt \geq \int_0^T \sum_{i=1}^m c_i(u_i(t, \hat{\beta}_1(T))) dt. \quad (14)$$

As a corollary, (14) holds with equality if and only if $\mathbf{u}(t, \hat{\beta}_2(T)) = \mathbf{u}(t, \hat{\beta}_1(T))$ for all t (excluding any switching points). Thus, if

$$\int_0^T \sum_{i=1}^m c_i(u_i(t, \hat{\beta}(T))) dt = r, \quad (15)$$

then $\mathbf{u}(t, \hat{\beta}(T)) = \mathbf{u}^*(t)$ also for all t . Therefore, we have the following result:

Proposition 1. For the case of separable, linear objective functions, i.e., $J(\mathbf{x}(T)) = \langle \mathbf{p}, \mathbf{x}(T) \rangle$, the unique optimal control can be explicitly calculated using a number of evaluations of (15) that is logarithmic in the range of considered $\hat{\beta}(T)$'s.¹⁴

Using the process outlined above, we can use a simple bisection algorithm to find $\beta^*(T)$ and to solve the optimal control problem using a *single-shooting* approach. $\hat{\beta}(T)$ 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, since 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 (12) and (13) 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 $u_i^{\max} h_i(t) / c_i(u_i^{\max})$ for concave $c(\cdot)$ and $h_i(t)/v_i$ for linear $c(\cdot)$ with $\hat{\beta}(T)$ (as in (12) and (13)). One can think of

14. This proposition does not apply for the general $J(\mathbf{x}(T))$, as the equivalent definition of $h_i(t)$ in (14) would have to replace \mathbf{p} with $[\partial J(\mathbf{x}) / \partial \mathbf{x}]_{\mathbf{x}=\mathbf{x}^*(T)}$, which can only be evaluated with knowledge of the optimal terminal opinion vector $\mathbf{x}^*(T)$.

the terms containing $h_i(t)$ as a topographic relief map, signifying hills and valleys. $\hat{\beta}(T)$ represents 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. 4). 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.

Furthermore, the water-filling procedure shows the relative importance of channels over time with respect to external influence. As the optimal water-level is a monotone decreasing function of the budget available, one can see that the peaks in $u_i^{\max}h_i(t)/c_i(u_i^{\max})$ and $h_i(t)/v_i$ 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. Therefore, we can consider the explicitly computable result of $u_i^{\max}h_i(t)/c_i(u_i^{\max})$ (for the linear cost case, $h_i(t)/v_i$) to be a *direct metric* total order for the *effect of advertising on channel i at time t on the outcome of the election*, which we shall henceforth call *cost-effectiveness* of a channel.

One can extract some more insight from the structure of this metric to compare the relative importance of channels by considering (9) at extreme values of t :

Remark 4. If $t \ll T$ (i.e., early on in the time horizon) and $T \gg \frac{1}{\xi_2}$, the deciding factor in comparing the cost-effectiveness of channels is their *total reach* (e.g., $\sum_{j=1}^n b_{ji}$ for channel i) per unit cost; for example, for the linear $c(\cdot)$ case:

$$\frac{h_i(t)}{v_i} \approx \left(\frac{1}{n} \sum_{j=1}^n p_j \right) \left(\frac{\sum_{j=1}^n b_{ji}}{v_i} \right), \quad (16)$$

as $\frac{1}{n} \sum_{j=1}^n p_j$ is the same for all channels.

Proof. From Section IV, we know that $\xi_1 = 0 < \xi_2 \leq \dots \leq \xi_n$, so when $t \ll T$ and $T \gg \frac{1}{\xi_2}$, then $e^{-\xi_j(T-t)} \approx 0$ for $j \geq 2$ and $e^{-\xi_1(T-t)} = e^0 = 1$. Replacing these values, and $\mathbf{Q}(:, 1) = \frac{1}{\sqrt{n}} \mathbf{1}_n$, in (9) completes the argument. ■

Remark 5. However, if $1 - \frac{t}{T} \ll \frac{1}{\xi_n T}$ (i.e., late on in the time horizon), *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:

$$\frac{h_i(t)}{v_i} \approx \frac{\langle \mathbf{p}, \mathbf{B}(:, i) \rangle}{v_i}, \quad (17)$$

Proof. When $1 - \frac{t}{T} \ll \frac{1}{\xi_n T}$, then $e^{-\xi_j(T-t)} \approx 1$ for all $j \in [n]$. Replacing these values in (9) results in:

$$h_i(t) = \sum_{j=1}^N \langle \mathbf{Q}(:, j) \mathbf{p} \rangle \langle \mathbf{Q}(:, j) \mathbf{B}(:, i) \rangle = \langle \mathbf{p} \mathbf{B}(:, i) \rangle,$$

due to the orthonormality of the eigenvectors in \mathbf{Q} and the definition of an inner product. ■

This is instructive, as it shows that at the start of a campaign, cheap broadcast methods (that maximize total reach per unit cost) would be preferable to costly (premature) targeting of likely voters, while as election day approaches, the alignment of a channel with the likelihood of voting among its targets gradually increases in importance.

D. Separable Sigmoid Objective

In this case, as shown in Appendix B, available in the online supplemental material, the equivalent $h_i(t)$ expression (9) will feature a term $\lambda^*(T)$, instead of \mathbf{p} , that depends strongly on $|x_i^*(T) - \theta_i|$, how far agent i is from changing their mind (6), for all i . The further away $x_i^*(T)$ is from θ_i (i.e., the farther they are from changing their mind, or alternatively the more convinced they are), the smaller the relevant $\lambda_i^*(T)$. For a given $\epsilon \ll 1$, define the set of *late-deciders* [17] under the optimal advertising action \mathbf{u}^* to be $\mathcal{L} := \{j : |x_j^*(T) - \theta_j| < \epsilon\}$. When $\mathcal{L} \neq \emptyset$ (i.e., there are late deciders), we can use the water-filling machinery in Section IV with the changes outlined below to approximate the cost-effectiveness of channels and to calculate the optimal allocation using the much faster method described therein.

We define λ such that:

$$\bar{\lambda}_j = \begin{cases} 0 & \text{for } j \notin \mathcal{L}, \\ \frac{\alpha_j p_j}{2} & \text{for } j \in \mathcal{L}. \end{cases}$$

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

$$\frac{h_i(t)}{v_i} = \frac{\langle \mathbf{Q}(:, j) \lambda \rangle \langle \mathbf{Q}(:, j) \mathbf{B}(:, i) \rangle e^{-\xi_j(T-t)}}{v_i}.$$

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

V. SIMULATION STUDIES

In this section, we first study a simple example to show that even in small networks, the optimal budget allocation across channels can have complicated, sometimes counter-intuitive, structures. Furthermore, we show that in many cases, the bound derived from Theorem 1 grows much slower than the number of agents, n . Then, we study the performance of our algorithm on a real network derived from political discussions between MIT students prior to the 2008 US general election, and compare it to policies that use more simple centrality metrics that do not consider the temporal degrees of freedom of advertising policies.

We first examine a network of 7 agents with linear objectives, with $\mathbf{p} = (3\%, 2\%, 10\%, 100\%, 6\%, 7\%, 1\%)$. 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. 5; 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, $\mathbf{B}(:, 1) = (1, -1, 1, 0, 1, 0, 0)^T$ 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, $\mathbf{B}(:, 2) = (-1, 1, 0, 1, 0, 0, 0)^T$, has a positive effect on the likely voter, but it has more limited effects on the rest of the agents. We solve the optimal budget allocation problem in Fig. 6 using the waterfilling methodology of Section IV. As noted in

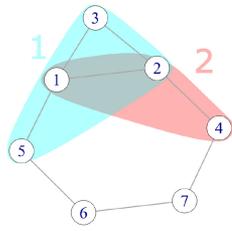


Fig. 5. A network of $n = 7$ agents with $m = 2$ influence channels. The lines in solid black represent the underlying communication network L . The blue and red boxes delineate the two channels that are available for influence in terms of agents affected (but not intensity).

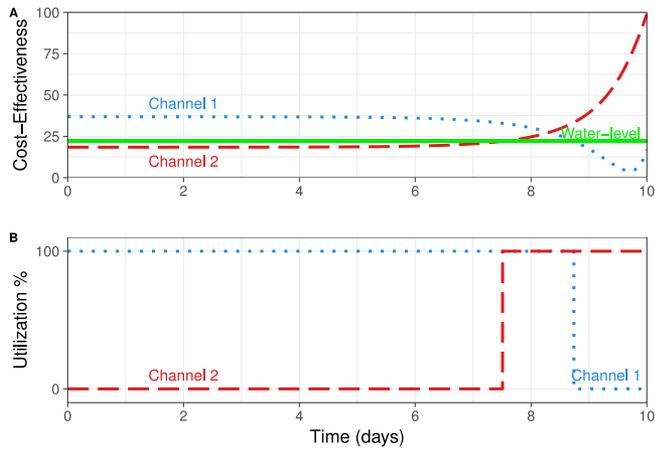


Fig. 6. (A) We plot the cost-effectiveness of the two channels over a time horizon of $T = 10$ days when they have equal cost ($c_i(u_i) = v_i u_i$ for $i = 1, 2$ with $v_1 = v_2 = 1$). We then derive the water-level β for $r = 11$ (in green). (B) The optimal water-level determines the optimal utilization 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.

Remarks 4 and 5, at times $t \ll T = 10$, the cost-effectiveness of the two equal cost channels is measured by (16) 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 depends on the match between the reach of the channel and the likelihood of agents to vote (17), and therefore the cost-effectiveness of channel 2 is higher. This optimal control is bang-bang with bounded numbers of transitions, as proven in Theorem 1.

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 1.2.a) 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 1.2.b. We simulated 1000 random connected Erdos-Renyi graphs with uniformly random linear objective functions for the case of $\mathbf{B}(:, 1) = (1, \mathbf{0}_{1 \times (n-1)})^T$, and plotted the mean, variance, and maximum value of the bound in Theorem 1.2.b as the number of agents was varied. As can be seen in Fig. 7, this latter bound is much smaller (around 10 for 200 agents), and its growth with respect to the number of agents is very slow. This is significant since, from an applied perspective, the

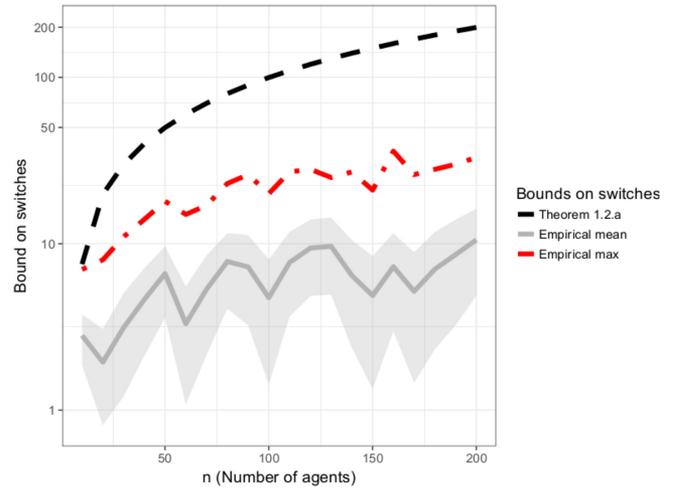


Fig. 7. We plot the upper-bounds on the number of switches of the optimal resource allocation derived from Theorem 1 for 1,000 random Erdos-Renyi graphs for a channel that only affects the first agent as the size of the network is varied. The dashed bound is from Theorem 1.2.a, and can be seen to grow with the number of agents. For a linear objective (Theorem 1.2.b), the gray line (with the related standard deviation band) shows the mean bound on the number of switches, while the dashed red line shows the empirical maximum of the bound over 1,000 runs. We can see that both these values are significantly smaller than the bound from Theorem 1.2.a and increase at a much slower rate with the size of the network.

advertiser can enumerate and evaluate a much smaller set of candidate optimal solutions, and yet can be reasonably sure that the best such policy is globally optimal.

We now study the performance of our algorithm on a test scenario derived from the MIT Social Evolution data-set [50]. In this data-set, among other data, the political opinions and communication patterns of 84 MIT students are recorded in the period prior to, and following, the 2008 US presidential election. Furthermore, the living sector and year of the students was recorded. We consider the problem of deciding how the campaign of Barack Obama should have invested its resources to disseminate campaign literature in order to guarantee the best electoral outcome. While this is admittedly a stylized and somewhat simplistic, it adequately demonstrates how the model could be specified and identified.

In particular, we focus on a social network derived from the reported political discussions between students conducted on 2008-09-09 and 2008-10-19, the only two surveys conducted before the November 4th election. We consider “discussion” to be an undirected communication between individuals, and thus we aggregate communications that are flagged by both participants. However, we sum distinct communications between two individuals to denote a stronger bond. This process is used to generate the Laplacian communication matrix L , including by normalizing discussions by time-frame $T = 66$ days. The discussion graph is plotted in Fig. 8.

We then calculate the channel matrix \mathbf{B} for channels that represent the 8 dorm floors and 5 seniority levels (freshman, sophomore, junior, senior, graduate), leading to a total m of 13. The channels mapping to dorm floors capture possible advertising on bulletin boards, for example, while the 5 seniority-based channels could represent e-mail lists targeting

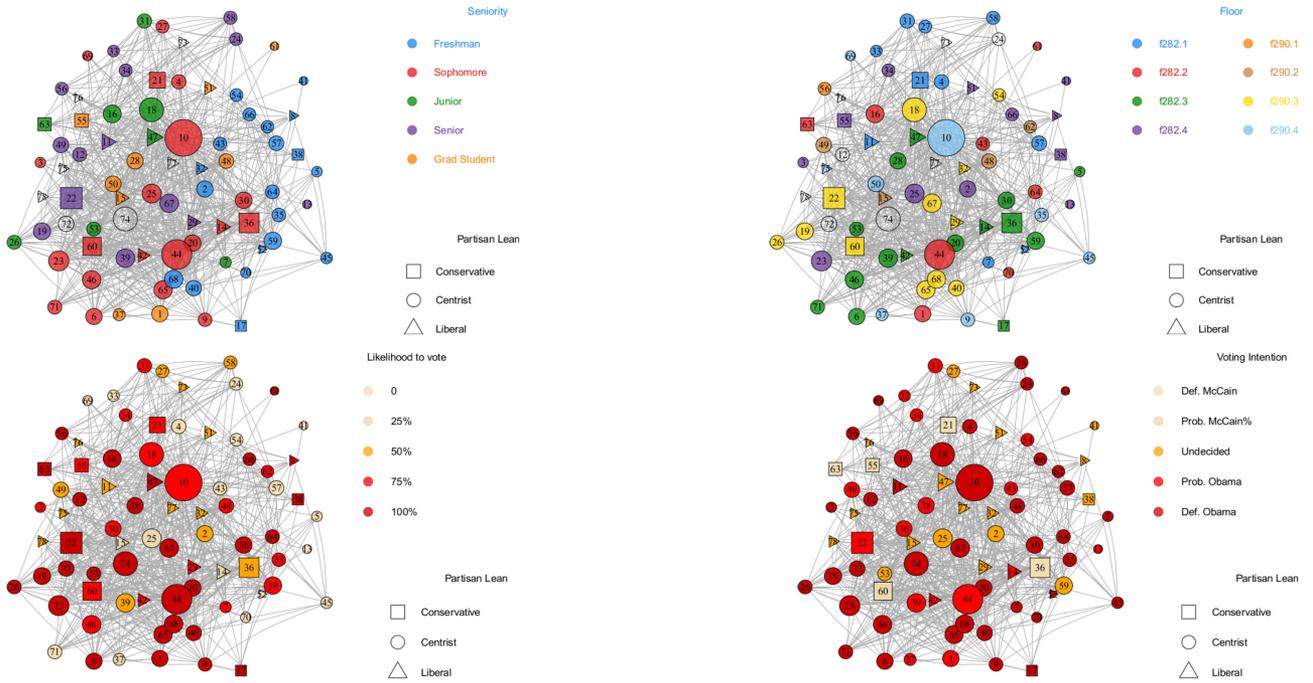


Fig. 8. In this figure, political discussions between students are mapped as a graph, with the weight of links being derived from the frequency of discussions. In the left-hand graph, nodes are colored according to the seniority of the students, while in the right-hand graph, they are colored in according to their residence, which are the two determinants of advertising channels.

specific graduating years. The weight of the effect of each channel is derived from the self-reported liberal or conservative initial bias of the individuals, as the effect of advertising depends on its alignment with the values of the target [51]. In this example, we consider the propagation of campaign literature targeted at liberals, and thus off-putting to conservatives. This can model any of the wedge issues of the campaign (e.g., the Iraq war [52]). Thus, advertising can have a negative effect on outcomes for the campaign, making some individuals less likely to vote for the candidate. Thus, we assign a non-zero value to b_{ik} if i is in the k th dorm floor/seniority group, with the sign being determined by individual i 's self-described "liberal" (from the perspective of the Obama campaign, positive) or "conservative" (respectively, negative) affiliation, and the magnitude being determined by their self-described strength of identification with that affiliation.¹⁵

Furthermore, and again for simplicity, we assume that all the channels have a similar linear cost, $\mathbf{v} = \mathbf{1}$, and have similar small effects on the voting intentions of participants $\mathbf{u}_{max} = 0.01$. The channels are also shown in Fig. 8.

We consider a linear objective for the campaign. While the more complex sigmoid objective functions are a better model for decision-making, we consider the simpler linear case for tractability. We map the self-reported likelihood of voting of participants in September '08 to a $[0, 1]$ scale, taking 5 equally spaced values, and constituting the vector \mathbf{p} . While self-reported turnout

has been shown to be an unreliable predictor of voting behavior [53], we operate under the reasonable assumption that more reliable information is not available to the political campaign. These likelihoods of voting can be seen in Fig. 8.

Finally, we instantiate the opinions of individuals $\mathbf{x}(0)$ (representing their voting intentions, as viewed by the Obama campaign) with the self-reported voting intention of individuals in September '08, which takes 8 values, mapped to values between $[-1, 1]$. Again, for simplicity, we pool third-party voters with undecided voters. A more realistic scenario with vectors of opinions would be able to more accurately capture the diversity in opinions, but would not be as instructive as the current example for the performance of our policy and the resulting centralities. These initial preferences can be seen in Fig. 8.

In Fig. 9, we show the water-filling procedure and the resulting optimal utilization of the channels for this problem for a budget of $r = 52$. We observe that the optimal budget allocation only uses three of the dorm floor channels, with time-variations in the use of channels f290.2 and f290.4. This is somewhat counter-intuitive given the significantly higher reach of the seniority-based channels, which are unused by the optimal allocation, while dorm floor f290.1, which only includes one solitary individual, is used throughout the time period at the maximum possible rate. However, as the cost of utilizing a channel is taken to be proportional to its reach, f290.1 is utilized because it is very effective *relative to its cost*.

To benchmark our results, we compared the results of the optimal budget allocation policy on electoral outcomes to policies based on different types of static centralities: betweenness centrality, eigen-centrality, Page-rank, and degree centrality. For the comparison, we ranked channels according to a

15. The mapping for b_{ik} within channel k was as follows: "Extremely conservative": -1, "Conservative": -0.66, "Slightly conservative": -0.33, "Moderate middle of the road": 0, "Slightly liberal": 0.33, "Liberal": 0.66, and "Extremely liberal": 1.

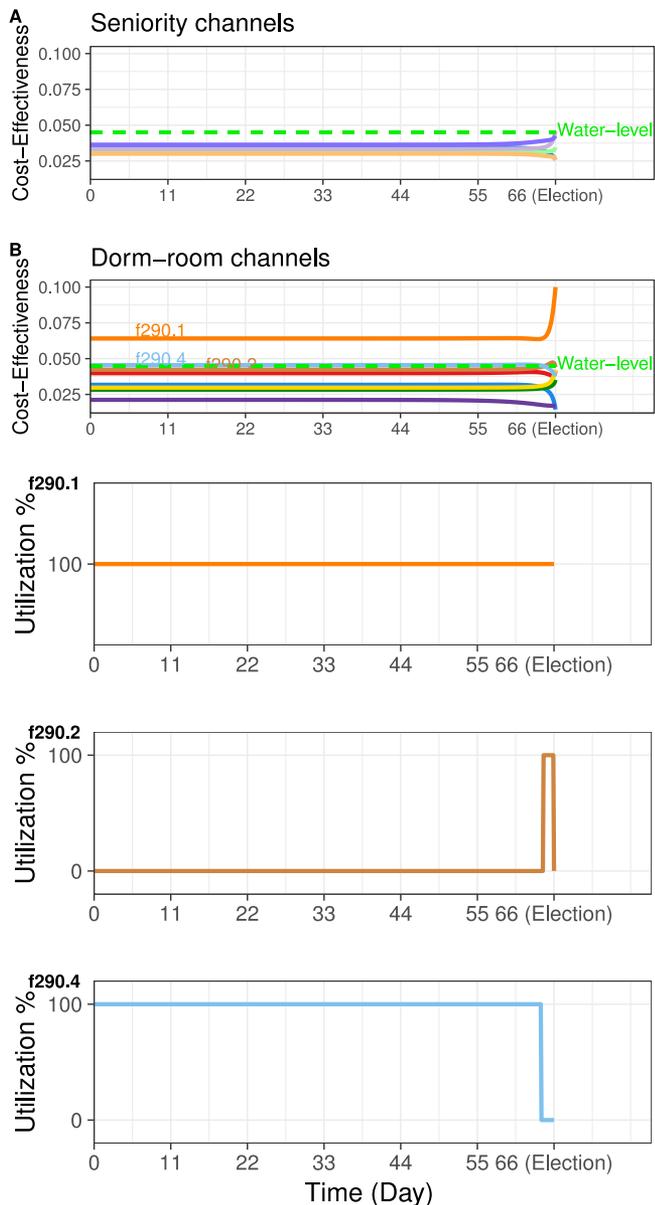


Fig. 9. We demonstrate a case with two $h_i(t)/v_i$ functions for the case of linear $c_i(\cdot)$. Areas above the water-line ($\hat{\beta}(T)$) translate to $u_i(t, \hat{\beta}(T)) = u_i^{max}(t)$, while those below translate to $u_i(t, \hat{\beta}(T)) = 0$. The amount of budget spent for this $\hat{\beta}(T)$ can thus be calculated from the resulting $u_i(t, \hat{\beta}(T))$, and so $\hat{\beta}(T)$ can be adjusted to find $\beta^*(T)$.

channel-weighted sum of the centrality in question (to account for possible negative effects of a channel on an individual), and allocated our budget to the highest-ranked channels at the maximum possible static rate until the exhaustion of our budget ($r = 52$). Table I summarizes the relative differences in outcomes between the optimal dynamic budget allocation policy and the static benchmarks. We see that using our optimal water-filling algorithm based on our novel cost-effectiveness metric leads to a 26 percent increase in the expected number of votes compared to the best static policy based on common centrality measures, a significant improvement.

Note that due to our results outlined in Section IV, the total budget does not affect the relative priority assigned to the

TABLE I
EXPECTED NUMBER OF VOTES GAINED BASED ON THE GREEDY ALLOCATION OF BUDGET TO THE HIGHEST-RANKED CHANNELS ACCORDING TO DIFFERENT TYPES OF CENTRALITIES

Basis of decision-making	Expected number of votes
Optimal Allocation	39.03
Betweenness centrality	29.86
Eigen-centrality	29.16
Page-rank	29.89
Degree centrality	30.03

channels by either the optimal algorithm (as it does not change the cost-effectiveness of channels) or the heuristics. Rather, it determines how many of the channels with high priorities can be used without going over the budget. The results presented in Table I are representative in the regime where the budget is a binding constraint in the choice of advertising channels.

VI. SUMMARY AND DISCUSSION

We consider the problem of optimally allocating a finite budget over time across several advertising channels. We showed, using Pontryagin’s maximum principle, that the optimal allocation follows a bang-bang structure, in which we either invest fully in a channel or not at all. In other words, to maximize the effectiveness of our budget, we should invest fully over a number of waves, and let the effect of the waves propagate in between waves. Furthermore, we show that the number of advertising waves during which we invest fully is, in practice, much smaller than the number of agents in the network. This result greatly facilitates the explicit computation of the optimal allocation policy over time. Furthermore, we showed that the exact optimal control can be calculated using an efficient water-filling procedure for a linear objective. From this water-filling procedure, we rigorously defined “cost-effectiveness” as a metric for ranking and comparing the influence of different channels at differing times on outcomes. Finally, applying our results to the sigmoid approximation of the electoral campaign/voting model confirmed the intuitive notion that identifying last-deciders determines the campaign strategy.

These results can be generalized in various ways. The notion of channel interaction in this work did not come with any constraints on the presence or attention of the channel members. Adding such a constraint can more clearly model real-world interactions. Furthermore, the linear model of influence is also a constraint that may be relaxed to obtain more general structures on influence control. Finally, this work looked at a single issue where each agents opinion was represented with a scalar—the same methodology can be extended to find optimal advertising strategies with vectors of opinions.

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,” *Amer. Econ. J.: Appl. Econ.*, vol. 1, no. 2, pp. 35–52, 2009.
- [2] M. Brodie, C. Deane, and S. Cho, “Regional variations in public opinion on the affordable care act,” *J. Health Politics Policy 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," *J. Marketing Res.*, 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'," *J. Advertising Res.*, 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," *J. Marketing Commun.*, 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," *Amer. Behavioral Scientist*, vol. 57, no. 7, pp. 920–942, 2013.
- [7] L. Zhang and D. J. Chung, "The air war versus the ground game: An analysis of multi-channel marketing in U.S. presidential elections," Harvard Business School Working Paper, no. 15-033, Oct. 2014, (Revised March 2018).
- [8] G. A. Huber and K. Arceneaux, "Identifying the persuasive effects of presidential advertising," *Amer. J. Political Sci.*, vol. 51, no. 4, pp. 957–977, 2007.
- [9] This Year Next Year: Worldwide Media and Marketing Forecasts, *GroupM*, 2017. [Online]. Available: <https://www.groupm.com/news/this-year-next-year-2017-global-advertising-to-reach-547b>
- [10] K. Kaye, "Data-driven targeting creates huge 2016 political ad shift: Broadcast tv down 20%, cable and digital way up," 2017. [Online]. Available: <http://adage.com/article/media/2016-political-broadcast-tv-spend-20-cable-52/307346>
- [11] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *J. Artif. Soc. Social Simul.*, vol. 5, no. 3, pp. 1–33, 2002.
- [12] 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," *J. Marketing Manag.*, vol. 27, no. 7–8, pp. 691–717, 2011.
- [13] 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.
- [14] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.
- [15] D. P. Spanos, R. Olfati-Saber, and R. M. Murray, "Dynamic consensus on mobile networks," in *Proc. IFAC World Congress*, 2005, pp. 1–6.
- [16] P. A. Naik and K. Raman, "Understanding the impact of synergy in multimedia communications," *J. Marketing Res.*, 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," *Eur. J. Marketing*, vol. 30, no. 10/11, pp. 127–139, 1996.
- [18] J. R. French Jr, "A formal theory of social power," *Psychological Rev.*, vol. 63, no. 3, 1956, Art. no. 181.
- [19] M. DeGroot, "Reaching a consensus," *J. Amer. Statistical Assoc.*, vol. 69, no. 345, pp. 118–121, 1974.
- [20] R. P. Abelson, "Mathematical models of the distribution of attitudes under controversy," *Contributions Math. Psychology*, vol. 14, pp. 1–160, 1964.
- [21] M. Taylor, "Towards a mathematical theory of influence and attitude change," *Human Relations*, vol. 21, no. 2, pp. 121–139, 1968.
- [22] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, Sep. 2004.
- [23] W. Ren, R. W. Beard, and E. M. Atkins, "A survey of consensus problems in multi-agent coordination," in *Proc. IEEE Amer. Control Conf.*, 2005, pp. 1859–1864.
- [24] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *J. Math. Sociology*, vol. 15, no. 3–4, pp. 193–206, 1990.
- [25] P. V. Marsden and N. E. Friedkin, "Network studies of social influence," *Sociol. Methods Res.*, vol. 22, no. 1, pp. 127–151, 1993.
- [26] N. E. Friedkin, "The problem of social control and coordination of complex systems in sociology: A look at the community cleavage problem," *IEEE Control Syst.*, vol. 35, no. 3, pp. 40–51, Jun. 2015.
- [27] A. V. Proskurnikov and R. Tempo, "A tutorial on modeling and analysis of dynamic social networks: Part i," *Annu. Rev. Control*, vol. 43, pp. 65–79, 2017.
- [28] D. Kempe, J. Kleinberg, and É. Tardos, "Influential nodes in a diffusion model for social networks," in *Automata, Languages and Programming*. Berlin, Germany: Springer, 2005, pp. 1127–1138.
- [29] E. Mossel and S. Roch, "On the submodularity of influence in social networks," in *Proc. 39th Annu. ACM Symp. Theory Comput.*, pp. 128–134, 2007.
- [30] E. Anshelevich, D. Chakrabarty, A. Hate, and C. Swamy, "Approximation algorithms for the firefighter problem: Cuts over time and submodularity," in *Algorithms and Computation*. Berlin, Germany: Springer, 2009, pp. 974–983.
- [31] C. Borgs, M. Brautbar, J. Chayes, and B. Lucier, "Maximizing social influence in nearly optimal time," in *Proc. 25th Annu. ACM-SIAM Symp. Discrete Algorithms*, 2014, pp. 946–957.
- [32] F. Morone and H. A. Makse, "Influence maximization in complex networks through optimal percolation," *Nature*, vol. 524, no. 7563, pp. 65–68, 2015.
- [33] E. D. Demaine, M. Hajiaghayi, H. Mahini, D. L. Malec, S. Raghavan, A. Sawant, and M. Zadimoghadam, "How to influence people with partial incentives," in *Proc. 23rd Int. Conf. World Wide Web*, 2014, pp. 937–948.
- [34] E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione, "Binary opinion dynamics with stubborn agents," *ACM Trans. Econ. Comput. (TEAC)*, vol. 1, no. 4, Dec. 2013, Art. no. 19.
- [35] M. Athans and P. L. Falb, *Optimal Control: An Introduction to the Theory and Its Applications*. North Chelmsford, MA, USA: Courier Corporation, 1966.
- [36] D. Liberzon, *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton, NJ, USA: Princeton Univ. Press, 2012.
- [37] A. Seierstad and K. Sydsaeter, *Optimal Control Theory with Economic Applications*. Amsterdam, The Netherlands: Elsevier, 1986.
- [38] K. Kandhway and J. Kuri, "Campaigning in heterogeneous social networks: Optimal control of si information epidemics," *IEEE/ACM Trans. Netw.*, vol. 24, no. 1, pp. 383–396, Feb. 2016.
- [39] S. Eshghi, M. Khouzani, S. Sarkar, N. B. Shroff, and S. S. Venkatesh, "Optimal energy-aware epidemic routing in DTNs," *IEEE Trans. Autom. Control*, vol. 60, no. 6, pp. 1554–1569, Jun. 2015.
- [40] K. Kandhway and J. Kuri, "How to run a campaign: Optimal control of sis and sir information epidemics," *Appl. Math. Comput.*, vol. 231, pp. 79–92, 2014.
- [41] S. Eshghi, M. Khouzani, S. Sarkar, and S. S. Venkatesh, "Optimal patching in clustered malware epidemics," *IEEE/ACM Trans. Netw.*, vol. 24, no. 1, pp. 283–298, Feb. 2016.
- [42] A. A. Cárdenas, S. Amin, and S. Sastry, "Research challenges for the security of control systems," in *Proc. 3rd Conf. Hot Topics Security*, 2008, Art. no. 6.
- [43] Y. Mo, E. Garone, A. Casavola, and B. Sinopoli, "False data injection attacks against state estimation in wireless sensor networks," in *Proc. 49th IEEE Conf. Decision Control*, 2010, pp. 5967–5972.
- [44] J. Goldenberg, B. Libai, and E. Muller, "Talk of the network: A complex systems look at the underlying process of word-of-mouth," *Marketing Lett.*, vol. 12, no. 3, pp. 211–223, 2001.
- [45] W. Chen, Y. Wang, and S. Yang, "Efficient influence maximization in social networks," in *Proc. 15th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2009, pp. 199–208.
- [46] 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:1702.03432, 2017.
- [47] F. R. Chung, *Spectral Graph Theory*. Providence, RI, USA: American Mathematical Society, 1997, vol. 92.
- [48] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD, USA: JHU Press, 1996, vol. 3.
- [49] C.-T. Chen, *Linear System Theory and Design*. London, U.K.: Oxford Univ. Press, 1995.
- [50] A. Madan, M. Cebrían, S. Moturu, K. Farrahi, et al., "Sensing the 'health state' of a community," *IEEE Pervasive Comput.*, vol. 11, no. 4, pp. 36–45, Oct. 2012.
- [51] C. K. Atkin, L. Bowen, O. B. Nayman, and K. G. Sheinkopf, "Quality versus quantity in televised political ads," *Public Opinion Quart.*, vol. 37, no. 2, pp. 209–224, 1973.
- [52] "Issues: Iraq," *CNN Election Center*, 2008. [Online]. Available: <http://www.cnn.com/ELECTION/2008/issues/issues.iraq.html>
- [53] D. Granberg and S. Holmberg, "Self-reported turnout and voter validation," *Amer. J. Political Sci.*, vol. 35, pp. 448–459, 1991.



Soheil Eshghi (Member, IEEE) received the Ph.D. degree in electrical and systems engineering from the University of Pennsylvania, in 2015, Philadelphia, PA, USA. He is currently a Postdoctoral Associate with the Yale Institute for Network Science (YINS), Department of Electrical Engineering, Yale University, New Haven, CT, USA. His research interests include mathematical models of influence in social systems, as well as applications of optimal control and game theory to social, biological, and electrical networks.



Victor M. Preciado (Member, IEEE) received the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA. He is currently an Associate Professor of electrical and systems engineering with the University of Pennsylvania, Philadelphia, PA, USA. His main research interests include the intersection of big data and network science, in particular, in using innovative mathematical and computational approaches to capture the essence of complex, high-dimensional dynamical systems.

Relevant applications of this line of research can be found in the context of socio-technical networks, brain dynamical networks, healthcare operations, biological systems, and critical technological infrastructure. He is a member of the Networked & Social Systems Engineering Program, the Warren Center for Network & Data Sciences, and the Applied Math & Computational Science (AMCS) Program. He was the recipient of the 2017 National Science Foundation Faculty Early Career Development (CAREER) Award.



Saswati Sarkar received the M.E. from the Electrical Communication Engineering Department, Indian Institute of Science, Bangalore, India, in 1996 and Ph.D. degree from the Electrical and Computer Engineering Department at the University of Maryland, College Park, MD, USA, in 2000. In 2000, she joined as an Assistant Professor the Electrical and Systems Engineering Department, University of Pennsylvania, Philadelphia, PA, USA, where she is currently a Professor. Her research interests are in stochastic control, resource allocation, dynamic games, and economics

of networks. She received the Motorola gold medal for the best masters student in the Division of Electrical Sciences at the Indian Institute of Science and a National Science Foundation (NSF) Faculty Early Career Development Award in 2003. She was an Associate Editor for the IEEE TRANSACTION ON WIRELESS COMMUNICATIONS from 2001 to 2006, an Associate Editor for the IEEE/ACM TRANSACTIONS ON NETWORKS from 2006 to 2012, and is currently a Senior Editor of IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS.



Santosh S. Venkatesh (Member, IEEE) received the PhD degree in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 1986. He is currently a Professor with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA. His research interests include probabilistic models, random graphs, epidemiology, network, and information theory.



Qing Zhao (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, USA, in 2001. She is currently a Professor and Gordon Lankton Sesquicentennial Faculty Fellow with the School of Electrical and Computer Engineering, Cornell University. Prior to joining Cornell University in 2015, she was a Professor with the University of California, Davis (UC Davis). Her research interests include decision theory and stochastic optimization, machine learning, statistical inference, and algorithmic theory with applications

in infrastructure and communication systems and social economic networks. She was the recipient of the 2010 IEEE Signal Processing Magazine Best Paper Award and the 2000 Young Author Best Paper Award from the IEEE Signal Processing Society. While on the faculty of UC Davis, she held the title of UC Davis Chancellors Fellow. She is a Marie Skłodowska-Curie Fellow, awarded by the European Unions Horizon 2020 Research and Innovation Program, and a Jubilee Chair Professor of Chalmers University, Sweden.



Raissa D'Souza received the Ph.D. degree in statistical physics from the Massachusetts Institute of Technology, Cambridge, MA, in 1999. She is currently a professor of computer science and of mechanical engineering with the University of California Davis, Davis, CA, USA, as well as an External Professor with the Santa Fe Institute, Santa Fe, NM, USA. She is a Fellow of the American Physical Society, is on the editorial board of numerous international mathematics and physics journals, has organized key scientific meetings like NetSci 2014, was a member of the

World Economic Forum's Global Agenda Council on Complex Systems, and is currently the President of the Network Science Society.



Ananthram Swami (Fellow, IEEE) received the B. Tech. degree in electrical engineering from Indian Institute of Technology Bombay, Mumbai, India, the M.S. degree in electrical engineering from Rice University, Houston, TX, USA, and the Ph.D. degree in electrical engineering from University of Southern California (USC), Los Angeles, CA, USA. He is with the U.S. Army Research Laboratory (ARL) as the Army's ST (Senior Research Scientist) for Network Science. He has held positions with Unocal Corporation, USC, CS-3, and Malgudi Systems. He was a

Statistical Consultant to the California Lottery, developed a MATLAB-based toolbox for non-Gaussian signal processing, and has held visiting faculty positions at INP, Toulouse, and Imperial College, London. His research interests include the broad area of network science with applications in composite tactical networks. He is an ARL Fellow.