

Opinion Dynamics with Varying Susceptibility to Persuasion

Rediet Abebe
Cornell University
red@cs.cornell.edu

David Parkes
Harvard University
parkes@eecs.harvard.edu

Jon Kleinberg
Cornell University
kleinber@cs.cornell.edu

Charalampos E. Tsourakakis
Boston University
ctsourak@bu.edu

ABSTRACT

A long line of work in social psychology has studied variations in people’s susceptibility to persuasion – the extent to which they are willing to modify their opinions on a topic. This body of literature suggests an interesting perspective on theoretical models of opinion formation on social networks: in addition to considering interventions that directly modify people’s intrinsic opinions, it is also natural to consider those that modify people’s susceptibility to persuasion. Here, we adopt a popular model for social opinion dynamics, and formalize the opinion maximization and minimization problems where interventions happen at the level of susceptibility.

We show that modeling interventions at the level of susceptibility leads to an interesting family of new questions in network opinion dynamics. We find that the questions are quite different depending on whether there is an overall budget constraining the number of agents we can target or not. We give a polynomial-time algorithm for finding the optimal target-set to optimize the sum of opinions when there are no budget constraints on the size of the target-set. We show that this problem is NP-hard when there is a budget, and that the objective function is neither submodular nor supermodular. Finally, we propose a heuristic for the budgeted opinion optimization problem and show its efficacy at finding target-sets that optimize the sum of opinions on real world networks, including a Twitter network with real opinion estimates.

CCS CONCEPTS

• **Theory of computation** → **Social networks**; *Graph algorithms analysis; Exact and approximate computation of equilibria;*

KEYWORDS

social networks, opinion dynamics, influence models

ACM Reference Format:

Rediet Abebe, Jon Kleinberg, David Parkes, and Charalampos E. Tsourakakis. 2018. Opinion Dynamics with Varying Susceptibility to Persuasion. In *KDD ’18: The 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, August 19–23, 2018, London, United Kingdom*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3219819.3219983>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD ’18, August 19–23, 2018, London, United Kingdom

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5552-0/18/08...\$15.00

<https://doi.org/10.1145/3219819.3219983>

1 INTRODUCTION

A rich line of empirical work in social psychology has studied people’s susceptibility to persuasion. This property measures the extent to which individuals are willing to modify their opinions in reaction to the opinions expressed by those around them, and it is distinct from the opinions they express. Research in the area has ranged from adolescent susceptibility to peer pressure related to risky and antisocial behavior [3, 15, 17, 40, 43, 44] to the role of susceptibility to persuasion in politics [19, 34, 36]. Individuals’ susceptibility to persuasion can be affected by specific strategies and framings aimed at increasing susceptibility [7–9, 28, 32, 41, 42]. For instance, if it is known that an individual is receptive to *persuasion by authority*, one can adopt a strategy that utilizes arguments from official sources and authority figures to increase that individuals’ susceptibility to persuasion with respect to a particular topic.

Modifying network opinions has far-reaching implications including product marketing, public health campaigns, the success of political candidates, and public opinions on issues of global interest. In recent years, there has also been work in Human Computer Interaction focusing on *persuasive technologies*, which are designed with the goal of changing a person’s attitude or behavior [18, 26, 28]. This work has shown that not only do people differ in their susceptibility to persuasion, but that persuasive technologies can also be adapted to each individual to change their susceptibility to persuasion.

Despite the long line of empirical work emphasizing the importance of individuals’ susceptibility to persuasion, theoretical studies of opinion formation models have not focused on interventions at the level of susceptibility. Social influence studies have considered interventions that directly act on the opinions themselves, both in discrete models, as in the work of Domingos and Richardson [16], Kempe et al. [29], and related applications, and more recently in a model with continuous opinions by Gionis et al. [23].

In this work, we adopt an opinion formation model inspired by the work of DeGroot [14] and Friedkin and Johnsen [20], and initiate a study of the impact of interventions at the level of susceptibility. In this model, each agent i is endowed with an *innate opinion* $s_i \in [0, 1]$ and a parameter representing susceptibility to persuasion, which we will call the *resistance parameter*, $\alpha_i \in (0, 1]$. The innate opinion s_i reflects the intrinsic position of agent i on a certain topic, while α_i reflects the agent’s willingness, or lack thereof, to conform with the opinions of neighbors in the social network. Higher values of the resistance parameter α_i indicate a lower tendency to conform with neighboring opinions. According to the opinion dynamics model, the final opinion of each agent i is a function of the social network, the set of innate opinions, and the resistance parameters,

determined by computing the *equilibrium* state of a dynamic process of opinion updating. For more details on the model, see Section 2. We study the following natural question:

Problem 1. Given an opinion dynamics model, and a set of agents, each of whom has an innate opinion that reflects the agent's intrinsic position on a topic, and a resistance parameter measuring the agent's propensity for changing his or her opinion, how should we change the agents' susceptibility in order to maximize (or minimize) the total sum of opinions at equilibrium?

We call the set of agents whose susceptibility parameters we change the *target-set*. In most opinion dynamics settings, the case where there are no constraints on the size of the target-set is generally not interesting since it often involves simple solutions that modify a parameter to a specific value across essentially all agents. This is not the case in our setting – the unbudgeted version remains technically interesting, and the optimal solution does not generally involve targeting all agents. We will therefore consider both the unbudgeted and the budgeted versions of this problem.

In this paper, we formalize the problem of optimizing the total sum of opinions at equilibrium by modifying the agents' susceptibility to persuasion. We formalize both the maximization and minimization versions of the problems, as well as an unbudgeted, and a budgeted variation of these problems.

For the unbudgeted maximization and minimization problems we provide efficient polynomial time algorithms. In contrast to the unbudgeted problem, the budgeted problem is NP-hard. We show that our objective is neither sub- nor super-modular. We prove that there always exists an optimal solution where the resistance parameters α_i are set to extreme values (see Theorem 2.1). We use this theorem as our guide to provide a heuristic that greedily optimizes the agents' susceptibility to persuasion.

We evaluate our unbudgeted and budgeted methods on real world network topologies, and on both synthetic and real opinions. Our findings indicate that intervening at the susceptibility level can significantly affect the total sum of opinions in the network.

2 MODEL

Let $G = (V, E)$ be a simple, undirected graph, where $V = [n]$ is the set of agents and E is the set of edges. Each agent $i \in V$ is associated with an *innate opinion* $s_i \in [0, 1]$, where higher values correspond to more favorable opinions towards a given topic and a parameter measuring an agent's susceptibility to persuasion $\alpha_i \in (0, 1]$, where higher values signify agents who are less susceptible to changing their opinion. We call α_i the *resistance parameter*. The opinion dynamics evolve in discrete time according to the following model, inspired by the work of [14, 20]:

$$x_i(t+1) = \alpha_i s_i + (1 - \alpha_i) \frac{\sum_{j \in N(i)} x_j(t)}{\deg(i)}. \quad (1)$$

Here, $N(i) = \{j : (i, j) \in E\}$ is the set of neighbors of i , and $\deg(i) = |N(i)|$. It is known that this dynamics converges to a unique equilibrium if $\alpha_i > 0$ for all $i \in V$ [12]. The equilibrium opinion vector z is the solution to a linear system of equations:

$$z = (I - (I - A)P)^{-1}As, \quad (2)$$

where $A = \text{Diag}(\alpha)$ is a diagonal matrix where entry $A_{i,i}$ corresponds to α_i and P is the random walk matrix, i.e., $P_{i,j} = \frac{1}{\deg(i)}$ for each edge $(i, j) \in E(G)$, and zero otherwise. We call z_i the *expressed opinion* of agent i .

Definition 2. Given the opinion dynamics model (1), a social network $G = (V, E)$, and the set of parameters $\{s_i\}_{i \in V}$, $\{\alpha_i\}_{i \in V}$, we define the *total sum of opinions* at the equilibrium as:

$$f(s, \alpha) \stackrel{\text{def}}{=} \vec{1}^T z = \vec{1}^T (I - (I - A)P)^{-1}As.$$

Our goal is to optimize the objective $f(s, \alpha)$, by choosing certain α_i parameters in a range $[\ell, u]$, where $0 < \ell \leq u$. We consider the following problems:

- (1) *Opinion Minimization*: find $S \subseteq V$ such that we can set the $\alpha_i \in [\ell, u]$ for $i \in S$ in order to minimize $f(s, \alpha)$.
- (2) *Opinion Maximization*: find $S \subseteq V$ such that we can set the $\alpha_i \in [\ell, u]$ for $i \in S$ in order to maximize $f(s, \alpha)$.

Previous work has focused on modifying the agents' innate or expressed opinions. We observe that targeting agents at the level of susceptibility versus at the level of innate opinions can lead to very different results, as the next example shows.

Example 3. Consider the star graph on n nodes, where the center node has (s_i, α_i) given by $(1, \epsilon/c)$ and the other nodes have values $(0, \epsilon)$. We want to maximize the sum of expressed opinions.

Suppose ϵ is a value arbitrarily close to 0. As we increase c , the total sum of opinions in the network tends to 0. If we can target a single agent, then an optimal intervention at the level of innate opinions changes the innate opinions of one of the leaves to 1. At equilibrium, the total sum of opinions (for small ϵ and large c) is close to $\frac{n}{n-1} \approx 1$. However, if we can target at the level of susceptibility, then we change the central node's α_i from $\frac{\epsilon}{c}$ to 1. Then, at equilibrium, the total sum of opinions becomes $f(s, \alpha) = n(1 - \phi(\epsilon))$ for a function $\phi(\cdot)$ going to 0 as ϵ goes to 0. The improvement is $\Omega(n)$, and hence very large compared to the intervention at the level of innate opinions.

We can construct an instance on the same graph structure to show the contrast in the opposite direction as well, where targeting at the level of innate opinions causes a large improvement compared to targeting at the level of susceptibility. Starting with the example above, change the innate opinion of the center node to be 0 and the resistance parameter to be 1. In equilibrium, the sum of expressed opinions is 0. If we can target agents at the level of innate opinions, then changing the innate opinion of the center node to be 1 will lead to a sum of expressed opinions of n , whereas we cannot obtain a sum of expressed opinions greater than 0 by targeting agents at the level of susceptibility.

This above example also shows that interventions at the level of susceptibility can have substantial impact on network opinions. It gives a way of modifying the resistance parameters to create an unbounded increase in the sum of expressed opinions.

We now prove that there always exists an optimal solution that involves extreme values, ℓ and u , for the resistance parameters α_i .

THEOREM 2.1. *There always exists an optimal solution for the opinion optimization problems where $\alpha_i = \ell$ or $\alpha_i = u$.*

PROOF. Our proof relies on a random-walk interpretation of opinion dynamics [21, 24]. Before we give our proof, we describe this connection between absorbing random walks and opinion dynamics for completeness. Gionis et al. [24] consider a random walk with absorbing states on an augmented graph $H = (X, R)$ that is constructed as follows:

- (i) The set of vertices X of H is defined as $X = V \cup \bar{V}$, where \bar{V} is a set of n new nodes such that for each node $i \in V$ there is a copy $\sigma(i) \in \bar{V}$;
- (ii) The set of edges R of H includes all the edges E of G , plus a new set of edges between each node $i \in V$ and its copy $\sigma(i) \in \bar{V}$. That is, $R = E \cup \bar{E}$, and $\bar{E} = \{(i, \sigma(i)) \mid i \in V\}$;
- (iii) Finally, the weights of all edges in R are set to 1, assuming that G is unweighted.

A key observation in [24] is that the opinion vector at the equilibrium can be computed by performing the following absorbing random walk on the graph H . To determine the expressed opinion of node $i \in V$, we start a random walk from i ; with probability $1 - \alpha_i$ we follow an incident edge chosen uniformly at random from E , and with probability α_i we follow the edge from i to $\sigma(i)$, ending the walk. If the walk continues to a node j , we perform the same type of step — choosing an edge of E incident to j with probability $1 - \alpha_j$, and moving to $\sigma(j)$ with probability α_j — and iterate. The expressed opinion z_i^* of node i in the equilibrium of the process is equal to the expected value of the innate opinion s_j over all endpoints $\sigma(j)$ of the walk starting from i .

Therefore, the agents' equilibrium opinion is given by:

$$z_i^* = \alpha_i s_i + (1 - \alpha_i) p_{ii} z_i^* + (1 - \alpha_i)(1 - p_{ii}) y_i, \quad (3)$$

where p_{ii} is the probability that the random walk that starts at i returns again to i before absorption, and y_i is the expected opinion if the random walk gets absorbed to a node $\sigma(j) \neq \sigma(i)$.

Without loss of generality, we assume that the opinion optimization problem at hand is the maximization problem, although the proof can easily be adopted to the opinion minimization case. Suppose we have a solution for the opinion maximization problem such that α_i lies strictly between ℓ_i and u_i . We wish to show that setting α_i to be either ℓ_i or u_i will yield a sum of expressed opinions that is at least as large. Proceeding in this way one node at a time, we will get a solution in which all α_i are either ℓ_i or u_i .

By solving Equation (3) for z_i^* we obtain,

$$z_i^* = \frac{\alpha_i}{1 - (1 - \alpha_i)p_{ii}} s_i + \frac{(1 - \alpha_i)(1 - p_{ii})}{1 - (1 - \alpha_i)p_{ii}} y_i. \quad (4)$$

Therefore z_i^* is the convex combination of two probabilities s_i, y_i . Since we wish to maximize z_i^* , if $s_i \leq y_i$ we need to minimize $\frac{\alpha_i}{1 - (1 - \alpha_i)p_{ii}}$; otherwise we will minimize $\frac{(1 - \alpha_i)(1 - p_{ii})}{1 - (1 - \alpha_i)p_{ii}}$. Note that the function $g(x) = \frac{x}{1 - (1 - x)p_{ii}}$ for $x \in [\ell_i, u_i]$ is monotone increasing, and therefore maximized at $x = u_i$ and minimized at $x = \ell_i$. Therefore, optimality is obtained by setting α_i to be either ℓ_i or u_i . This yields a sum of expressed opinions that is at least as large as the original solution. \square

3 UNBUDGETED OPINION OPTIMIZATION

In this section, we show that when there are no constraints on the size of the target-set, both the opinion maximization and minimization problems can be solved in polynomial time.

3.1 Unbudgeted Opinion Minimization

Recall from Section 2, $A = \text{Diag}(\alpha)$ is a diagonal matrix where entry $A_{i,i}$ corresponds to α_i and P is the random walk matrix.

We formalize the unbudgeted opinion minimization problem as:

$$\begin{aligned} & \text{minimize} && \vec{1}^T z \\ & \text{subject to} && z = (I - (I - A)P)^{-1} A s \\ & && A \geq \ell I \\ & && A \leq u I \end{aligned} \quad (5)$$

THEOREM 3.1. *The Opinion Minimization formulation (5) is solvable in polynomial time.*

Before we proceed with the proof of Theorem 3.1, we rewrite the objective of formulation (5) in a way that is convenient for algebraic manipulation. Define $X = A^{-1}$. Then,

$$\begin{aligned} (I - (I - A)P)^{-1} A &= (I - (I - A)P)^{-1} (A^{-1})^{-1} = (I - (I - A)P)^{-1} X^{-1} \\ &= (X - (X - XA)P)^{-1} = (X - (X - I)P)^{-1}. \end{aligned}$$

Therefore, we rewrite our original formulation to optimize over the set of diagonal matrices X whose entries lie in $[\frac{1}{u}, \frac{1}{\ell}]$.

$$\begin{aligned} & \text{minimize} && \vec{1}^T (X - (X - I)P)^{-1} s \\ & \text{subject to} && X \geq \frac{1}{u} I \\ & && X \leq \frac{1}{\ell} I \\ & && X \text{ diagonal} \end{aligned} \quad (6)$$

LEMMA 3.2. *The following set is convex:*

$$\mathcal{K} = \{X \in \mathbb{R}^{n \times n} : X \text{ diagonal}, \frac{1}{u} \leq x_{ii} \leq \frac{1}{\ell} \text{ for all } i \in [n]\}.$$

PROOF. Let $X_1, X_2 \in \mathcal{K}$, $\lambda \in (0, 1)$. Matrix $Y = \lambda X_1 + (1 - \lambda) X_2$ is clearly diagonal, and any diagonal entry $y_{ii} = \lambda x_{ii}^1 + (1 - \lambda) x_{ii}^2$ satisfies,

$$\frac{1}{u} \leq y_{ii} \leq \frac{1}{\ell}.$$

Therefore, $Y \in \mathcal{K}$, which implies the claim. \square

LEMMA 3.3. *All eigenvalues of matrix $Z = (X - (X - I)P)$, $X \in \mathcal{K}$ are strictly positive reals (in general Z is not symmetric). Furthermore, Z is positive definite, i.e., for any $y \in \mathbb{R}^n$ the quadratic form $y^T Z y$ is non-negative.*

Notice that the above two statements are not simultaneously true for non-symmetric matrices: a non-symmetric matrix with all eigenvalues positive is not necessarily positive definite. On the other hand, if a non-symmetric matrix is positive definite, all its eigenvalues have to be positive.

PROOF. First, we prove that all eigenvalues are real. We provide a sequence of similarity transformations that prove $Z = (X - (X - I)P)$ is similar to a symmetric matrix. To simplify notation, when $U = R^{-1}AR$, where R is a diagonal matrix, we write $U \equiv A$.

$$\begin{aligned}
X - (X - I)P &\equiv X - D^{\frac{1}{2}} \left((X - I)D^{-\frac{1}{2}}A \right) D^{-\frac{1}{2}} \\
&= X - (X - I)D^{-\frac{1}{2}}AD^{-\frac{1}{2}} \\
&\equiv X - (X - I)^{-\frac{1}{2}}(X - I)D^{-\frac{1}{2}}AD^{-\frac{1}{2}}(X - I)^{\frac{1}{2}} \\
&= X - (X - I)^{\frac{1}{2}}D^{-\frac{1}{2}}AD^{-\frac{1}{2}}(X - I)^{\frac{1}{2}}
\end{aligned}$$

Since the adjacency matrix A is symmetric, the latter matrix is also symmetric, and therefore has real eigenvalues.

To show all eigenvalues are positive, we invoke the Gershgorin circle theorem. Notice that $Z_{ii} = x_{ii}$ (under our assumption that the graph is simple) and the sum of non-diagonal entries in row i is $-(x_{ii} - 1)$. Therefore, every eigenvalue of Z lies in at least one of the Gershgorin discs:

$$|\lambda - x_{ii}| \leq (x_{ii} - 1) \Leftrightarrow 1 \leq \lambda \leq 2x_i - 1.$$

We now show that the quadratic form associated with Z is non-negative. Recall first that $x_{ii} \geq 1$ and $\deg(i) \geq 1$, therefore

$$\frac{1}{\deg(i)} + \frac{1}{\deg(j)} \geq 2.$$

For any $y \in \mathbb{R}^n$,

$$\begin{aligned}
y^T Z y &= y^T (X - (X - I)P) y = y^T X y - y^T (X - I) P y \\
&= \sum_{i=1}^n x_{ii} y_i^2 - \sum_{i=1}^n (x_{ii} - 1) \frac{y_i}{\deg(i)} \sum_{j \sim i} y_j \\
&= \sum_{i=1}^n y_i^2 + \sum_{i=1}^n (x_{ii} - 1) \left(y_i^2 - \frac{y_i}{\deg(i)} \sum_{j \sim i} y_j \right) \\
&= \sum_{i=1}^n y_i^2 + \sum_{(i,j) \in E(G)} (x_{ii} - 1) \left(y_i^2 + y_j^2 - \left(\frac{y_i y_j}{\deg(i)} + \frac{y_i y_j}{\deg(j)} \right) \right) \\
&\geq \sum_{i=1}^n y_i^2 + \sum_{(i,j) \in E(G), y_i y_j \geq 0} (x_{ii} - 1) (y_i^2 + y_j^2 - 2y_i y_j) + \\
&\quad \sum_{(i,j) \in E(G), y_i y_j < 0} (x_{ii} - 1) \left(y_i^2 + y_j^2 + \left(\frac{|y_i y_j|}{\deg(i)} + \frac{|y_i y_j|}{\deg(j)} \right) \right) \\
&= \sum_{i=1}^n y_i^2 + \sum_{(i,j) \in E(G), y_i y_j \geq 0} (x_{ii} - 1) (y_i - y_j)^2 + \\
&\quad \sum_{(i,j) \in E(G), y_i y_j < 0} (x_{ii} - 1) \left(y_i^2 + y_j^2 + \left(\frac{|y_i y_j|}{\deg(i)} + \frac{|y_i y_j|}{\deg(j)} \right) \right) \\
&\geq 0
\end{aligned}$$

□

We proceed to prove that the objective is convex.

THEOREM 3.4. *The matrix $(X - (X - I)P)^{-1}$ is a convex function of $X \in \mathcal{K}$.*

PROOF. Let $X_1, X_2 \in \mathcal{K}$, and let $Z_i := (X_i - (X_i - I)P)$ for $i = 1, 2$. Also, fix $\lambda \in (0, 1)$, and let $Z(\lambda) = \lambda Z_1 + (1 - \lambda)Z_2$. We prove the following statement, that implies that $(X - (X - I)P)^{-1}$:

$$\lambda Z_1^{-1} + (1 - \lambda)Z_2^{-1} \geq (\lambda Z_1 + (1 - \lambda)Z_2)^{-1}.$$

Since Z_1, Z_2 are positive definite, we obtain that for any $y \in \mathbb{R}^n$

$$y^T Z(\lambda) y = \lambda y^T Z_1 y + (1 - \lambda) y^T Z_2 y \geq 0,$$

so $Z(\lambda)$ is positive definite, and so is $Z(\lambda)^{-1}$,

$$Z Z^{-1} = I \Rightarrow \frac{dZ}{d\lambda} Z^{-1} + Z \frac{dZ^{-1}}{d\lambda} = 0 \Rightarrow \frac{dZ^{-1}}{d\lambda} = -Z^{-1} \frac{dZ}{d\lambda} Z^{-1}.$$

By differentiating twice $Z(\lambda)$ with respect to λ , we also obtain $\frac{d^2 Z}{d\lambda^2} = 0$. So, the second derivative of the inverse is given by,

$$\frac{d^2 Z^{-1}}{d\lambda^2} = -\frac{dZ^{-1}}{d\lambda} \frac{dZ}{d\lambda} Z^{-1} - Z^{-1} \frac{dZ}{d\lambda} \frac{dZ^{-1}}{d\lambda} = 2Z^{-1} \frac{dZ}{d\lambda} Z^{-1} \frac{dZ}{d\lambda} Z^{-1}.$$

Pick any non-zero vector u and consider following pair of vector/-matrix valued functions:

$$v(\lambda) = \frac{dZ}{d\lambda} Z^{-1} u \quad \text{and} \quad \varphi(\lambda) = u^T Z^{-1}(\lambda) u.$$

The second derivative of $\varphi(\lambda)$ satisfies the following inequality:

$$\varphi(\lambda) = u^T \frac{d^2 Z^{-1}}{d\lambda^2} u = 2v^T(\lambda) Z^{-1}(\lambda) v(\lambda) \geq 0,$$

since Z^{-1} is positive definite. Therefore $\varphi(\lambda)$ is a convex function of λ for $\lambda \in (0, 1)$. Thus, for any such λ

$$(1 - \lambda)\varphi(0) + \lambda\varphi(1) - \varphi(\lambda) \geq 0 \quad (7)$$

$$\Leftrightarrow u^T \left[(1 - \lambda)X^{-1} + \lambda Y^{-1} - ((1 - \lambda)X + \lambda Y)^{-1} \right] u \geq 0. \quad (8)$$

This implies that $\lambda Z_1^{-1} + (1 - \lambda)Z_2^{-1} - (\lambda Z_1 + (1 - \lambda)Z_2)^{-1}$ is positive semidefinite, and thus our claim that $(X - (X - I)P)^{-1}$ is a convex function holds. □

We can prove that $\left((X - (X - I)P)^{-1} \right)^T$ is also a convex function in a similar way. We therefore obtain the following corollary:

COROLLARY 3.5. *For any fixed $y \in \mathbb{R}^n$, we can optimize the quadratic forms,*

$$y^T (X - (X - I)P)^{-1} y, \quad y^T \left((X - (X - I)P)^{-1} + \left(y^T \left((X - (X - I)P)^{-1} \right)^T \right) \right) y$$

over all diagonal matrices $X \in \mathcal{K}$ in polynomial time.

We can use off-the-shelf interior point methods for convex optimization. To conclude the proof of Theorem 3.1, we need to prove that our objective can be brought in the quadratic form of Corollary 3.5.

PROOF OF THEOREM 3.1. For simplicity, let $Z = (X - (X - I)P)$. We consider the quadratic,

$$\begin{aligned}
(s + \vec{1})^T Z^{-1} (s + \vec{1}) &= s^T Z^{-1} s + \vec{1}^T Z^{-1} \vec{1} + \vec{1}^T Z^{-1} s + s^T Z^{-1} \vec{1} \\
&= s^T Z^{-1} s + n + \vec{1}^T Z^{-1} s + s^T \vec{1} \Leftrightarrow \\
\vec{1}^T Z^{-1} s &= (s + \vec{1})^T Z^{-1} (s + \vec{1}) - s^T Z^{-1} s - C \\
&= \underbrace{\langle (s + \vec{1})(s + \vec{1})^T - ss^T, Z^{-1} \rangle}_{Q} - C,
\end{aligned}$$

where $C = n + \sum_{i=1}^n s_i$ is a constant that does not depend on the variable diagonal matrix X . Furthermore, Q is symmetric, i.e., $Q = Q^T$. Let $Q = \sum_{i=1}^n v_i q_i q_i^T$ be the eigendecomposition of Q . We observe that:

$$\begin{aligned} \text{trace}(Z^{-1}Q) &= \text{trace}\left(Z^{-1}\left(\sum_{i=1}^n v_i q_i q_i^T\right)\right) = \text{trace}\left(\sum_{i=1}^n v_i Z^{-1} q_i q_i^T\right) \\ &= \sum_{i=1}^n v_i \text{trace}\left(Z^{-1} q_i q_i^T\right) = \sum_{i=1}^n v_i \text{trace}\left(q_i^T Z^{-1} q_i\right) \\ &= \sum_{i=1}^n v_i \left(q_i^T Z^{-1} q_i\right) = \sum_{i=1}^n v_i \left(q_i^T \frac{1}{2}(Z^{-1} + (Z^{-1})^T) q_i\right) \\ &= \frac{1}{2} \text{trace}\left((Z^{-1} + (Z^{-1})^T)Q\right) \end{aligned} \quad (*)$$

In (*), we used the fact that for any square matrix A

$$x^T A x = \frac{1}{2} x^T (A + A^T) x.$$

Therefore, we conclude that our objective is equivalent to minimizing

$$\frac{1}{2} \langle (s + \vec{1})(s + \vec{1})^T - s s^T, Z^{-1} + (Z^{-1})^T \rangle,$$

subject to $X \in \mathcal{K}$. By Corollary 3.5, the opinion minimization problem is convex, and thus solvable in polynomial time. \square

3.2 Unbudgeted Opinion Maximization

We also prove that the opinion maximization problem can also be solved in polynomial time using machinery that we have already developed. We state the problem as:

$$\begin{array}{ll} \text{maximize} & \vec{1}^T (X - (X - I)P)^{-1} s \\ \text{subject to} & X \geq \frac{1}{u} I \\ & X \leq \frac{1}{\ell} I \\ & X \text{ diagonal} \end{array}$$

THEOREM 3.6. *We can solve the opinion maximization problem by invoking our opinion minimization algorithm using as our vector of initial opinions $\vec{1} - s$.*

PROOF. By Lemma 3.3 and since $P\vec{1} = \vec{1}$, for any $X \in \mathcal{K}$,

$$\begin{aligned} (X - (X - I)P)\vec{1} &= X\vec{1} - (X - I)\vec{1} = \vec{1} \\ \Leftrightarrow (X - (X - I)P)^{-1}\vec{1} &= \vec{1}. \end{aligned}$$

Therefore, the opinion minimization objective that uses as an initial opinions the vector $\vec{1} - s$ becomes

$$\begin{aligned} \vec{1}^T (X - (X - I)P)^{-1} (\vec{1} - s) &= \vec{1}^T (X - (X - I)P)^{-1} \vec{1} - \\ \vec{1}^T (X - (X - I)P)^{-1} s &= n - \vec{1}^T (X - (X - I)P)^{-1} s. \end{aligned}$$

Therefore, minimizing this objective is equivalent to maximizing $\vec{1}^T * (X - (X - I)P)^{-1} s$ subject to $X \in \mathcal{K}$. \square

4 BUDGETED OPINION OPTIMIZATION

We now consider the setting where there is a constraint on the size of the target-set. That is, we want to identify a set $T \subseteq V$ of size k such that changing the resistance parameters of agents in T optimally maximizes (resp. minimizes) the sum of expressed opinions in equilibrium. Recall from Section 2 that we denote by $f(s, \alpha)$ the sum of expressed opinions under the given vector of innate values and resistance parameters. We denote by $f(T)$ the sum of expressed opinions where we can set the α_i of the $i \in T$ optimally.

The objective function for finding the optimal target-set for opinion maximization is therefore $\max_{|T|=k} f(T)$, and, respectively, $\min_{|T|=k} f(T)$ for opinion minimization. Note, the objective function is monotone in the target-set since we can always set the choice of α_i to be the original α_i for any node $i \in T$. But, can we use submodular optimization tools for our budgeted opinion optimization problems? We show that, unlike in the case in other standard opinion optimization or influence maximization settings, neither submodularity nor supermodularity holds. We prove this result for the case of opinion maximization, but similar claims can be made for the minimization problem too by setting $s = 1 - s$.

Example 4. Consider the complete graph on three nodes where $s = (1, 0, 0)$ and $\alpha = (0.1, 0.1, 0.1)$. Suppose we are allowed to set resistance values in $[0.001, 1]$.

To note non-submodularity, set $T = \{1, 2\}$, $T' = \{1\}$, and $x = \{3\}$. For each of these sets, the optimal solution would set $\alpha_1 = 1$, $\alpha_2 = 0.001$ and $\alpha_3 = 0.001$. Submodularity would require:

$$f(1, 2, 3) - f(1, 2) \leq f(1, 3) - f(1),$$

which would imply that $0.191 \leq 0.168$.

To show that supermodularity does not hold, set $T = \{1\}$, $T' = \emptyset$, and $x = \{3\}$. Then, supermodularity would require that

$$f(3) - f(\emptyset) \leq f(1, 3) - f(1),$$

which would imply that $1.493 \leq 0.168$.

Highlighting a difference with the unbudgeted version, we show that the opinion maximization and minimization problems are NP-hard by a reduction from the vertex cover problem on d -regular graphs. We adapt a construction for a different opinion maximization problem given by [24] in the context of their model. We only present the proof for the opinion maximization problem; the proof can easily be adapted for opinion minimization as well.

THEOREM 4.1. *The budgeted opinion maximization (and similarly opinion minimization) problem is NP-hard.*

PROOF. Given a d -regular graph $G = (V, E)$ and an integer K , the vertex cover problem asks whether there exists a set of nodes S of size at most K such that S is a vertex cover. An instance of the decision version of opinion maximization consists of a graph G' , where each agent i has an internal opinion s_i and resistance parameter α_i , an integer k and a threshold θ . The solution to this decision version is “yes” if and only if there exists a set T of size at most k whose resistance values we can optimally set such that $\ell \leq \alpha \leq u$ and $f(T) \geq \theta$. We first prove this hardness result by letting the α_i of the target-set be anywhere in $[0, 1]$. We then prove that the reduction continues to hold for the case where $\alpha_i \in (0, 1]$.

Suppose we are given an instance of the vertex cover problem for regular graphs. We construct an instance of the decision version of the opinion maximization problem as follows: set $G' = (V \cup V', E \cup E')$ where V' consists of $2\sqrt{d}$ nodes for each node $i \in V$ (called duplicate nodes). The set E' is the edges of the form (i, i') joining nodes in V to their corresponding duplicate nodes. We set $s_i = 1$ and $\alpha_i = 0$ for all $i \in V$. We also set $s_{i'} = 0$ for all $i' \in V'$. For the set of duplicate nodes, we set the resistance value of half (that is, \sqrt{d} nodes) to be 1 and the other half to be 0. We call the former the stubborn duplicate nodes, and the latter the flexible duplicate nodes. We set $k = K$ and $\theta = (\sqrt{d} + 1)k + \sqrt{d}(n - k)\frac{d}{d + \sqrt{d}}$. This gives us a decision version of the opinion maximization problem.

We now want to show that a set T of at most k nodes in V is a vertex cover of G if and only if we can set the resistance values of the nodes in T (as a subset of G') to satisfy the decision version of the opinion maximization problem.

Suppose that T is a vertex cover of G . Then, it satisfies the decision problem above by setting the α_i of the agents in T to be 1. Indeed, the agents $i \in T$ will have expressed opinions of 1, as will all the flexible duplicate nodes neighboring the agents in T . All other nodes $j \in V$ will have all d of their neighbors in V having opinion 1 since T is a vertex cover. Therefore, they will have a final expressed opinion $\frac{d}{d + \sqrt{d}}$. Their \sqrt{d} flexible duplicate nodes will take on this value as well, while their stubborn duplicate nodes will maintain their innate opinion of 0. Therefore, we get a final sum of $(\sqrt{d} + 1)k + \sqrt{d}(n - k)\frac{d}{d + \sqrt{d}}$. Therefore, T corresponds to a target-set whose resistance we can set optimally such that $f(T) \geq \theta$.

To see the reverse direction, we want to show that if T is not a vertex cover of G , then we cannot set the resistance values of agents in T such that $f(T) \geq \theta$. If T is not a vertex cover, then there exists an edge $(i, j) \in E$ such that $i, j \notin T$. Therefore, the expressed opinion of both i and j will be at most $\frac{d-1}{d + \sqrt{d}}$. All the agents in T will maintain their high opinion of 1, and all other nodes $\ell \neq i, j$ will have an expressed opinion of at most $\frac{d}{d + \sqrt{d}}$, by the same argument as above. Taking into account the duplicate nodes, we have,

$$f(T) \leq (\sqrt{d} + 1)k + \sqrt{d}(n - k - 2)\frac{d}{d + \sqrt{d}} + 2\sqrt{d}\frac{d - 1}{d + \sqrt{d}} < \theta.$$

Let θ' denote the middle term in this chain of inequalities; we have thus shown that if G does not have a k -node vertex cover, then for any set T of at most k nodes, we have $f(T) \leq \theta' < \theta$.

What remains to show is there exists an optimal target-set that consists of only nodes in V and not in V' , where we assume that this target-set is at most of size $k \leq n$ for $n = |V|$. Denote the optimal target-set in by T^* . Suppose there is a node $i' \in T$ which is a node in V' . Let i' be adjacent to a node $i \in V$. If i' is a flexible duplicate node, then it is clear that the expressed opinions of both i' and i will be 1, if we include i in the target-set instead. (Or simply remove i' from the target-set if i is already in T .) Suppose i' is a stubborn duplicate node. First, assume that $i \in T$. Consider $T = T \setminus \{i'\} \cup \{j\}$, where $j \notin T$. The expressed opinion of j , and therefore of all its corresponding flexible duplicate nodes, with T as the target-set is at most $\frac{d}{d + \sqrt{d}}$. With j in the target-set, the expressed opinion goes to 1. We therefore need to show that the loss in opinion from removing i' from the target-set, is at most the

gain by adding j to the target-set. Indeed, it is straightforward to verify that $1 \leq \left(1 - \frac{d}{d + \sqrt{d}}\right)(\sqrt{d} + 1)$.

Now, consider i' is a stubborn duplicate node and $i \notin T$. We compare T to $T' = T \setminus \{i'\} \cup \{i\}$. Suppose there are h nodes adjacent to i in V' with resistance parameter 0. (In particular, $h \geq \sqrt{d} + 1$.) The expressed opinion of i with T as the target-set is at most $\frac{d}{d + 2\sqrt{d} - h}$, and is shared by all h of the duplicate nodes. Under T' , i will have an expressed opinion of 1, as will all h of the nodes with $\alpha = 0$, except for node i' which will now have resistance value of 1. Thus, we need to show that

$$1 - \frac{d}{d + 2\sqrt{d} - h} \leq h \left(1 - \frac{d}{d + \sqrt{d} - h}\right).$$

This holds since $h \geq 1$.

We now show how to modify the proof for the case in which the resistance values belong to the interval $[\varepsilon, 1]$ rather than $[0, 1]$, for some $\varepsilon > 0$. If we view $f(T)$ as a function of the resistance parameters in the network G' constructed in the reduction, it is a continuous function. Thus, if we let $\gamma = \theta - \theta' > 0$, we can choose $\varepsilon > 0$ small enough that the following holds. If we take all nodes i in the previous reduction with $\alpha_i = 0$, and we push their resistance values up to $\alpha_i = \varepsilon$, then (i) if T is a k -node vertex cover in G , we have $f(T) \geq \theta - \gamma/3$, and (ii) if G does not have a k -node vertex cover, for any set T of at most k nodes, we have $f(T) \leq \theta' + \gamma/3$. Since $\theta - \gamma/2$ lies between these two bounds, it follows that with resistance values in the interval $[\varepsilon, 1]$, there is a k -node vertex cover T in the original instance if and only if there is a set T in the reduction to opinion maximization with $f(T) \geq \theta - \gamma/2$. Therefore, the proof holds for resistance values in $[\varepsilon, 1]$. \square

In light of this hardness result, and inspired by Theorem 2.1, we propose a greedy heuristic for the budgeted opinion optimization problems. The version shown below is for opinion maximization, but can be converted to the opinion minimization problem by setting $s = 1 - s$.

Greedy Heuristic

Require: $k \geq 1$

```

1:  $S \leftarrow \emptyset$ 
2: Compute  $f(\emptyset) = (I - (I - A)P)^{-1}As$ 
3: for  $i \leftarrow 1$  to  $k$  do
4:    $max \leftarrow -\infty$ 
5:   for  $v \in V \setminus T$  do
6:     Compute  $f(T \cup \{v\})$ 
7:     if  $f(T \cup \{v\}) - f(T) \geq max$  then
8:        $winner \leftarrow v$ 
9:        $max \leftarrow f(T \cup \{v\}) - f(T)$ 
10:    end if
11:  end for
12:   $T \leftarrow T \cup winner$ 
13: end for
14: Return  $T$ .
```

In each round, the heuristic chooses the node that gives the largest marginal gain $f(T \cup \{v\}) - f(T)$. The algorithm requires

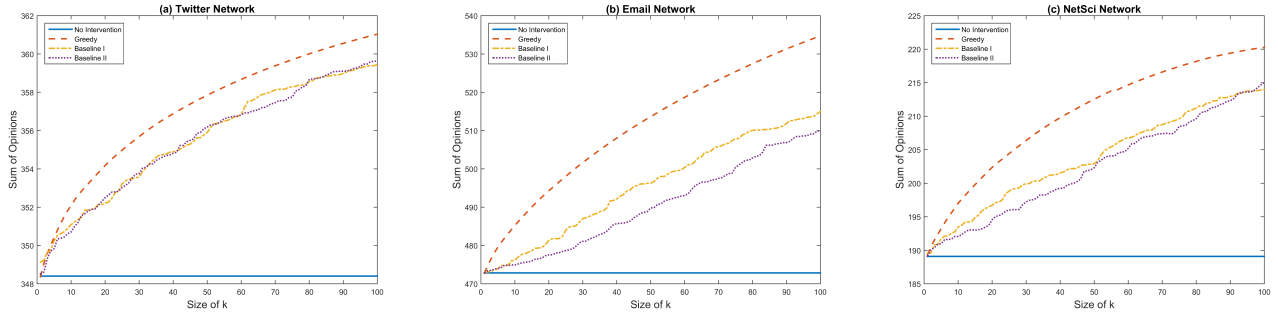


Figure 1: Total sum of opinions at equilibrium vs. budget k for the budgeted opinion maximization problem using no intervention, Baseline I, Baseline II, and our greedy heuristic for (a) Twitter, (b) Email, and (c) NetScience networks respectively.

$O(n^3)$ time to compute $f(0)$ using standard Gauss-Jordan elimination for matrix inversion, and then it performs $O(nk)$ loops, that each requires $O(n^2)$ time. To see why $O(n^2)$ time suffices instead of the straightforward $O(n^3)$ time of matrix inversion, suppose that we know $(I - (I - A)P)^{-1}$ and we wish to change the resistance of node i from α_i to be 1. The new resistance matrix is $A' = A + x_i x_i^T$, where $x_i \in \mathbb{R}^n$ has $x_i(j) = 0$ for all $j \neq i$, and $x_i(i) = \sqrt{1 - \alpha_i}$. Instead of computing the inverse $(I - (I - A')P)^{-1}$, we can use the Sherman-Morrison formula for the rank-1 update $(I - (I - A)P + x_i(x_i^T P))^{-1}$, which leads to a run-time of $O(kn^3)$ instead of $O(kn^4)$.

In the next section, we implement this heuristic and compare its performance against some natural baselines on real and synthetic values on real-world networks

5 EXPERIMENTAL RESULTS

Experimental setup. We use several datasets as sources of real network topologies: a social network of monks [39], social network of members in a karate club [46], co-appearance network of characters in Les Misérables [30], co-authorship network of network scientists [37], email communication network at a research institution [31, 45], and social links on Twitter [13].

The Twitter dataset is associated with real-world opinions obtained as described by De et al. [13]. For our purposes, we take the first expressed opinion by each agent to be the innate opinion, and we normalize these values, which fall in $[-1, 1]$ in the original dataset, to fall in $[0, 1]$. The sum of the innate opinions is 347.042 with a mean of 0.633. For a detailed description of the dataset collection, see [13].

For the remaining networks, we set the innate opinions uniformly at random in $[0, 1]$ for the budgeted experiments and uniformly at random as well as power law for the unbudgeted experiments. We set the slope for the power law distribution to be 2. The resistance parameters are chosen uniformly at random in $[0.001, 1]$.

For the budgeted version of our problem, we present the performance of the heuristics for the three largest networks (i.e., Twitter, Email Network, and NetScience). The results are qualitatively similar for the rest of the networks. We compare the performance of our greedy heuristic with two baselines that output a set of k nodes whose resistance parameter is set to 1. Baseline I chooses k agents with the highest innate opinions (ties broken arbitrarily). Baseline

II assigns to each node i a score $\text{score}(i) = \frac{\deg(i)}{2m} \times \frac{s_i}{\sum_{j \in N(i)} s_j}$ and chooses the k agents with the highest score. The intuition is that we would like to pick a node that has a high centrality and relatively high opinion compared to their neighbors and fix this agent to have an α of 1. We consider budget values k in $\{1, 2, \dots, 100\}$. We compare the sum of expressed opinions without intervention as well as the outputs of the greedy heuristic and the two baselines. All simulations run on a laptop with 1.7 GHz Intel Core i7 processor and 8GB of main memory.¹

Experimental findings. Figures 1(a), (b), (c) show the results obtained by our greedy heuristic, and the two baseline methods for the Twitter, Email, and NetScience networks respectively. All methods achieve a significant improvement compared to having no interventions even for small values of k . Recall that for the Twitter network, we use estimates of real opinions for the agents. Homophily by opinions is observed and impacts the relatively higher performance of these baselines compared to the Email and NetScience networks where the opinion values are randomly generated. Baseline I and Baseline II show similar performance for the Twitter and NetScience networks. One underlying explanation for the relative under-performance of Baseline II compared to Baseline I in the Email network is that this network has skewed degree distribution which might inflate the scores of higher degree nodes. We note that the run times for the Twitter, Email, and NetScience networks are 12.0, 3.4, and 111.5 seconds per k , respectively.

Baseline I

Require: $k \geq 1$

- 1: Let $s_1 \geq s_2 \geq \dots \geq s_n$.
 - 2: Return $T \leftarrow \{1, \dots, k\}$
-

Baseline II

Require: $k \geq 1$

- 1: Let $\text{score}(i) \leftarrow \frac{\deg(i)}{2m} \times \frac{s_i}{\sum_{j \in N(i)} s_j}$ for all $i \in V$.
 - 2: Sort the nodes such that $\text{score}(1) \geq \text{score}(2) \geq \dots \geq \text{score}(n)$
 - 3: Return $T \leftarrow \{1, \dots, k\}$
-

¹ The code for the budgeted and unbudgeted experiments can be found in <https://github.com/redabebe/opinion-optimization> and <https://github.com/tsourolampis/opdyn-social-influence>, respectively.

G	n	m	$s \sim U[0, 1]$				$s \sim \text{Power Law}(\alpha)$			
			$\bar{1}^T s$	$\bar{1}^T z$	$\bar{1}^T z_{\min}$	$\bar{1}^T z_{\max}$	$\bar{1}^T s$	$\bar{1}^T z$	$\bar{1}^T z_{\min}$	$\bar{1}^T z_{\max}$
Monks	18	41	9.34	8.12	1.22	16.38	11.88	11.29	5.62	16.68
Karate	33	78	16.70	17.83	1.97	32.47	22.29	21.07	5.27	33.56
LesMis	77	254	38.81	40.5	4.21	74.00	48.68	52.11	8.44	75.84
NetScience	379	914	188.39	191.79	17.19	319.59	257.58	249.88	64.43	340.54
Email	986	16064	414.62	440.96	388.39	495.44	283.62	565.86	91.18	610.72
Twitter	548	3638	289.29	260.96	241.86	270.21	193.09	359.85	62.00	379.00

Table 1: Results for unbudgeted algorithm for uniform and power law distributions (with slope $\alpha = 2$) of innate opinions s with output: sum of innate opinions, sum expressed opinions without intervention, sum of expressed opinions after interventions with minimization algorithm, and sum of expressed opinions with intervention by maximization algorithm.

Table 1 shows our experimental findings for the unbudgeted opinion minimization, and maximization problems. We report the total sum $\bar{1}^T s$ of innate opinions, the total sum $\bar{1}^T z$ of opinions at equilibrium (no intervention), the total sum $\bar{1}^T z_{\min}$ of opinions produced by our minimization algorithm, and the total sum $\bar{1}^T z_{\max}$ of opinions produced by our maximization algorithm. The resistance parameters for each of the networks are chosen uniformly at random from $[0, 1]$. We then let $\ell = 0.001$ and $u = 1$, which serve as an upper and lower bound for the resistance parameters. Note that the reported values are the average of 5 trials. In addition to the synthetically generated opinion values reported above, we also run the opinion optimization algorithms on the Twitter dataset using the opinions from [13] and get a sum of opinions 117.31 for the opinion minimization problem and 363.87 for the opinion maximization problem. The main observation here is that the effect of our intervention can be significant for both the opinion maximization and opinion minimization problems. Indeed, the sum of opinions in the case of the former is over an order of magnitude larger than the case of the latter for many of these networks.

An open direction for both the unbudgeted and budgeted settings is to find heuristics and algorithms that scale better with the size of the network, and how this trades off with the optimization problem for the sum of opinions. Furthermore, it may also be interesting to explore how well the proposed solutions in this paper perform in practice. That is, we assume that the model we consider represents the opinion formation process on social networks. Existing work in the literature explores how well the model considered here, as well as other models, approximate real evolution of opinions [11, 33]. It would be interesting to explore how well our solutions perform when the model we consider is only an approximation. There are also outstanding questions associated with estimation of parameters, including the susceptibility parameter in our paper.

6 RELATED WORK

To our knowledge, we are the first to consider an optimization framework based on opinion dynamics with varying susceptibility to persuasion. Below, we review some work that lies close to ours.

Asch’s conformity experiments are perhaps the most famous study on the impact of agents’ susceptibility to change their opinions [4]. This study shows how agents have different propensities for conforming with others. These propensities are modeled in our context by the set of parameters α . Since the work of Asch, there

have been various theories on peoples’ susceptibility to persuasion and how these can be affected. A notable example is Cialdini’s Six Principles of Persuasion, which highlight reciprocity, commitment and consistency, social proof, authority, liking, and scarcity, as key principles which can be utilized to alter peoples’ susceptibility to persuasion [7, 8]. This framework, and others, have been discussed in the context of altering susceptibility to persuasion in a variety of contexts. Crowley and Hoyer [9], and McGuire [32] discuss the ‘optimal arousal theory’, i.e., how novel stimuli can be utilized for persuasion when discussing arguments.

Opinion dynamics model social learning processes. DeGroot introduced a continuous opinion dynamics model in his seminal work on consensus formation [14]. A set of n individuals in society start with initial opinions on a subject. Individual opinions are updated using the average of the neighborhood of a fixed social network. Friedkin and Johnsen [20] extended the DeGroot model to include both disagreement and consensus by mixing each individual’s *innate belief* with some weight into the averaging process. This has inspired a lot of follow up work, including [2, 6, 11, 22, 23].

Bindel et al. use the Friedkin-Johnsen model as a framework for understanding the price of anarchy in society when individuals selfishly update their opinions in order to minimize the stress they experience [6]. They also consider network design questions: given a budget of k edges, and a node u , how should we add those k edges to u to optimize an objective related to the stress? Gionis et al. [23] use the same model to identify a set of target nodes whose innate opinions can be modified to optimize the sum of expressed opinions. Finally, in work concurrent with ours, Musco et al. adopt the same model to understand which graph topologies minimize the sum of disagreement and polarization [35].

While the expressed opinion of an agent is readily observable in a social network, both the agent’s innate opinion and conformity parameter are hidden. Das et al. [10, 12] have studied inferring these values. Specifically, they give a near-optimal sampling algorithm for estimating the true average innate opinion of the social network and justify the algorithm both analytically and experimentally [12]. Das et al. view the problem of susceptibility parameter estimation as a problem in constrained optimization and give efficient algorithms, which they validate on real-world data [10]. There are also several studies that perform experiments to study what phenomena influence opinion formation and how well these are captured by existing models [11, 33].

The problem we focus on is related to influence maximization, a question that draws on work of Domingos and Richardson [16, 38] and Kempe et al. [29]. Recent work by Abebe et al. considers the problem of influence maximization where agents who reject a behavior contribute to a negative reputational effect against it [1]. There is also a rich line of empirical work at the intersection of opinion dynamics and influence maximization including [5, 25, 27] that argues that agents adopt products to boost their status.

7 CONCLUSION

Inspired by a long line of empirical work on modifying agents' propensity for changing their opinions, we initiate the study of the effect on social influence of interventions on the susceptibility to persuasion. For this purpose, we adopt a popular opinion dynamics model, and seek to optimize the total sum of opinions at equilibrium in a social network. We consider both opinion maximization and minimization, and focus on two variations: a budgeted and an unbudgeted version. We prove that the unbudgeted problem is solvable in polynomial time (in both its maximization and minimization version). For the budgeted version, on the other hand, in contrast to a number of other opinion or influence maximization formalisms, the objective function is neither submodular nor supermodular. Computationally, we find that it is NP-hard. In light of this, we introduce a greedy heuristic for the budgeted setting, which we show has desirable performance on various real and synthetic datasets.

There are various open directions suggested by this framework. The first is whether we can provide any approximation algorithms for the budgeted setting. Through simulations, we also note that while not giving optimal results, local search algorithms yield satisfactory results in finding target-sets in the unbudgeted setting. Exploring this connection might yield insights into the opinion formation process. For the budgeted problem, rather than a budget on the number of agents, we can also consider having a total budget B for the total change in susceptibility for all agents. That is, given resistance values α_i , we want to set $\tilde{\alpha}_i$, such that $\sum_i |\alpha_i - \tilde{\alpha}_i| \leq B$. Finally, from an experimental point of view, it is worth designing scalable, well performing heuristics

ACKNOWLEDGEMENTS

Rediet Abebe was supported by a Google scholarship, a Facebook scholarship, and a Simons Investigator Award, and Jon Kleinberg was supported in part by a Simons Investigator Award, an ARO MURI grant, a Google Research Grant, and a Facebook Faculty Research Grant. We thank De et al. [13] for sharing the Twitter dataset.

Charalampos Tsourakakis would like to thank his newborn son Eftychios for the happiness he brought to his family.

REFERENCES

- [1] Rediet Abebe, Lada Adamic, and Jon Kleinberg. 2018. Mitigating Overexposure in Viral Marketing. In *Proc. of the 32nd AAAI Conf. on Artificial Intelligence*. AAAI-18.
- [2] Daron Acemoglu and Asuman Ozdaglar. 2011. Opinion dynamics and learning in social networks. *Dynamic Games and Applications* 1, 1 (2011), 3–49.
- [3] Joseph P. Allen, Maryfrances R. Porter, and F. Christy McFarland. 2006. Leaders and followers in adolescent close friendships: Susceptibility to peer influence as a predictor of risky behavior, friendship instability, and depression. *Development and psychopathology* 18, 1 (2006), 155–172.
- [4] Solomon E. Asch. 1955. Opinions and social pressure. *Scientific American* 193, 5 (1955), 31–35.
- [5] Jonah Berger and Chip Heath. 2007. Where consumers diverge from others: Identity signaling and product domains. *Journal of Consumer Research* 34, 2 (2007), 121–134.
- [6] David Bindel, Jon Kleinberg, and Sigal Oren. 2015. How bad is forming your own opinion? *Games and Economic Behavior* 92 (2015), 248–265.
- [7] Robert B. Cialdini. 2001. The science of persuasion. *Scientific American* 284, 2 (2001), 76–81.
- [8] Robert B. Cialdini. 2009. *Influence: Science and practice*. Vol. 4. Pearson education Boston, MA.
- [9] Ayn E. Crowley and Wayne D. Hoyer. 1994. An integrative framework for understanding two-sided persuasion. *Journal of Consumer research* 20, 4 (1994), 561–574.
- [10] Abhimanyu Das, Sreenivas Gollapudi, Arindham Khan, and Renato Paes Leme. 2014. Role of conformity in opinion dynamics in social networks. In *Proceedings of the second ACM conference on Online social networks*. ACM, 25–36.
- [11] Abhimanyu Das, Sreenivas Gollapudi, and Kamesh Munagala. 2014. Modeling opinion dynamics in social networks. In *Proceedings of the 7th ACM international conference on Web search and data mining*. ACM, 403–412.
- [12] Abhimanyu Das, Sreenivas Gollapudi, Rina Panigrahy, and Mahyar Salek. 2013. Debiasing social wisdom. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 500–508.
- [13] Abir De, Sourangshu Bhattacharya, Parantapa Bhattacharya, Niloy Ganguly, and Soumen Chakrabarti. 2014. Learning a linear influence model from transient opinion dynamics. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM, 401–410.
- [14] Morris H. DeGroot. 1974. Reaching a consensus. *J. Amer. Statist. Assoc.* 69, 345 (1974), 118–121.
- [15] TE Dielman, Deborah D. Kloska, Sharon L. Leech, John E. Schulenberg, and Jean T. Shope. 1992. Susceptibility to peer pressure as an explanatory variable for the differential effectiveness of an alcohol misuse prevention program in elementary schools. *Journal of School Health* 62, 6 (1992), 233–237.
- [16] Pedro Domingos and Matt Richardson. 2001. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 57–66.
- [17] William N. Evans, Wallace E. Oates, and Robert M. Schwab. 1992. Measuring peer group effects: A study of teenage behavior. *Journal of Political Economy* 100, 5 (1992), 966–991.
- [18] Brian J. Fogg. 2002. Persuasive technology: using computers to change what we think and do. *Ubiquity* 2002, December (2002), 5.
- [19] Marieke L. Fransen, Edith G. Smit, and Peeter W. Verlegh. 2015. Strategies and motives for resistance to persuasion: an integrative framework. (2015), 1201 pages.
- [20] Noah E. Friedkin and Eugene C. Johnsen. 1999. Social influence networks and opinion change. *Advances in Group Processes* 16 (1999), 1–29.
- [21] Javad Ghaderi and R. Srikant. 2013. Opinion dynamics in social networks: A local interaction game with stubborn agents. In *American Control Conference (ACC)*, 2013. IEEE, 1982–1987.
- [22] Javad Ghaderi and R. Srikant. 2013. Opinion dynamics in social networks: A local interaction game with stubborn agents. In *American Control Conference (ACC)*, 2013. IEEE, 1982–1987.
- [23] Aristides Gionis, Evimaria Terzi, and Panayiotis Tsaparas. 2013. Opinion maximization in social networks. In *Proceedings of the 2013 SIAM International Conference on Data Mining*. SIAM, 387–395.
- [24] Aristides Gionis, Evimaria Terzi, and Panayiotis Tsaparas. 2013. Opinion maximization in social networks. In *Proceedings of the 2013 SIAM International Conference on Data Mining*. SIAM, 387–395.
- [25] Yansong Hu and Christophe Van den Bulte. 2014. Nonmonotonic status effects in new product adoption. *Marketing Science* 33, 4 (2014), 509–533.
- [26] Wijnand IJsselstein, Yvonne De Kort, Cees Midden, Berry Eggen, and Elise Van Den Hoven. 2006. *Persuasive technology for human well-being: setting the scene*. 1–5.
- [27] Yogesh V. Joshi, David J. Reibstein, and Z. John Zhang. 2015. Turf wars: Product line strategies in competitive markets. *Marketing Science* 35, 1 (2015), 128–141.
- [28] Maurits Kaptein, Panos Markopoulos, Boris de Ruyter, and Emile Aarts. 2009. Can you be persuaded? individual differences in susceptibility to persuasion. In *IFIP Conference on Human-Computer Interaction*. Springer, 115–118.
- [29] David Kempe, Jon Kleinberg, and Eva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 137–146.
- [30] Donald Ervin Knuth. 1993. *The Stanford GraphBase: a platform for combinatorial computing*. Vol. 37. Addison-Wesley Reading.
- [31] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2007. Graph evolution: Densefication and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 1, 1 (2007), 2.
- [32] William J. McGuire. 1964. Inducing Resistance to Persuasion: Some Contemporary Approaches. *Advanced in Experimental Social Psychology* 1 (1964), 191 – 229.

- [33] Mehdi Moussaïd, Juliane E Kämmer, Pantelis P Analytis, and Hansjörg Neth. 2013. Social influence and the collective dynamics of opinion formation. *PLoS one* 8, 11 (2013), e78433.
- [34] Kevin M Murphy and Andrei Shleifer. 2004. Persuasion in politics. *American Economic Review* 94, 2 (2004), 435–439.
- [35] Cameron Musco, Christopher Musco, and Charalampos E Tsourakakis. 2017. Minimizing Polarization and Disagreement in Social Networks. *arXiv preprint arXiv:1712.09948* (2017).
- [36] Jacob R Neihsel and Sarah Niebler. 2015. On the limits of persuasion: Campaign ads and the structure of voters’ interpersonal discussion networks. *Political Communication* 32, 3 (2015), 434–452.
- [37] Mark EJ Newman. 2006. Finding community structure in networks using the eigenvectors of matrices. *Physical review E* 74, 3 (2006), 036104.
- [38] Matthew Richardson and Pedro Domingos. 2002. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 61–70.
- [39] Samuel F Sampson. 1969. *Crisis in a cloister*. Ph.D. Dissertation. Ph. D. Thesis. Cornell University, Ithaca.
- [40] Laurence Steinberg and Kathryn C Monahan. 2007. Age differences in resistance to peer influence. *Developmental psychology* 43, 6 (2007), 1531.
- [41] Zakary L Tormala and Richard E Petty. 2002. What doesn’t kill me makes me stronger: The effects of resisting persuasion on attitude certainty. *Journal of personality and social psychology* 83, 6 (2002), 1298.
- [42] Zakary L Tormala and Richard E Petty. 2004. Source credibility and attitude certainty: A metacognitive analysis of resistance to persuasion. *Journal of Consumer Psychology* 14, 4 (2004), 427–442.
- [43] Melanie Wakefield, Brian Flay, Mark Nichter, and Gary Giovino. 2003. Effects of anti-smoking advertising on youth smoking: a review. *Journal of health communication* 8, 3 (2003), 229–247.
- [44] Julie A Wall, Thomas G Power, and Consuelo Arbona. 1993. Susceptibility to antisocial peer pressure and its relation to acculturation in Mexican-American adolescents. *Journal of Adolescent Research* 8, 4 (1993), 403–418.
- [45] Hao Yin, Austin R Benson, Jure Leskovec, and David F Gleich. 2017. Local higher-order graph clustering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 555–564.
- [46] Wayne W Zachary. 1977. An information flow model for conflict and fission in small groups. *Journal of anthropological research* 33, 4 (1977), 452–473.