

Persuading Voters in District-based Elections

Matteo Castiglioni, Nicola Gatti

Politecnico di Milano, Piazza Leonardo da Vinci 32, I-20133, Milan, Italy
matteo.castiglioni@polimi.it, nicola.gatti@polimi.it

Abstract

We focus on the scenario in which an agent can exploit his information advantage to manipulate the outcome of an election. In particular, we study *district-based* elections with two candidates, in which the winner of the election is the candidate that wins in the majority of the districts. District-based elections are adopted worldwide (e.g., UK and USA) and are a natural extension of widely studied voting mechanisms (e.g., k -voting and plurality voting). We resort to the Bayesian persuasion framework, where the manipulator (*sender*) strategically discloses information to the voters (*receivers*) that update their beliefs rationally. We study both *private* signaling, in which the sender can use a private communication channel per receiver, and *public* signaling, in which the sender can use a single communication channel for all the receivers. Furthermore, for the first time, we introduce *semi-public* signaling in which the sender can use a single communication channel per district. We show that there is a sharp distinction between private and (semi-)public signaling. In particular, optimal private signaling schemes can provide an arbitrarily better probability of victory than (semi-)public ones and can be computed efficiently, while optimal (semi-)public signaling schemes cannot be approximated to within any factor in polynomial time unless $P = NP$. However, we show that reasonable relaxations allow the design of multi-criteria PTASs for optimal (semi-)public signaling schemes. In doing so, we introduce a novel property, namely *comparative stability*, and we design a bi-criteria PTAS for public signaling in general Bayesian persuasion problems beyond elections when the sender’s utility function is state-dependent.

Introduction

The fairness and efficiency of democratic elections largely depend on the news provided by the media. Indeed, often, citizens are called to express opinions on complex choices they do not know deeply enough to express informed judgments. Therefore, multiple and reliable sources of information providing fair and in-depth coverage of the public debate are crucial to guarantee the democratic process. However, most of the information shaping the voters’ opinions is not disclosed to inform in a disinterested way but instead aims to direct voters’ political orientation, thus *persuading* them to prefer one specific candidate over another.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

As recently showed by Allcott and Gentzkow (2017) and Guess, Nyhan, and Reifler (2018) for the 2016 US presidential election, the spread of fake news has become a major public concern for democracy. The problem of assessing the extent to which it is possible to manipulate an election has received considerable attention under the general framework of election control and has been investigated according to several perspectives, such as control by bribery (Faliszewski et al. 2009; Erdélyi, Reger, and Yang 2020) or by adding and deleting candidates and voters (Loreggia et al. 2015; Faliszewski, Hemaspaandra, and Hemaspaandra 2011; Liu et al. 2009; Chen et al. 2017). More recently, Sina et al. (2015), Faliszewski et al. (2018), Wilder and Vorobeychik (2018), Wilder and Vorobeychik (2019), and Castiglioni, Ferraioli, and Gatti (2020) studied social influence as a means of election control. In this paper, we pose the following question: *can an informed agent use his information advantage to influence an election’s outcome by the partial disclosure of information to rational voters?*

According to the classical Bayesian persuasion framework by Kamenica and Gentzkow (2011), the above problem can be formulated as a game with asymmetric information, where a sender can influence the behavior of the receiver(s) through the strategic provision of payoff-relevant information. In particular, the sender can strategically reveal information by means of a signaling scheme that determines “who knows what” about the parameters that govern the payoff functions. Alonso and Câmara (2016), Chan et al. (2019), and Bardhi and Guo (2018) provide the seminal attempts to apply the Bayesian persuasion framework to voting. More recently, Castiglioni, Celli, and Gatti (2020a) and Castiglioni, Celli, and Gatti (2020b) investigated its computational issues. All the aforementioned works focus on k -voting or plurality-voting elections. Differently, in this paper, we study for the first time how Bayesian persuasion can be adopted in more challenging settings such as district-based elections with two candidates, in which the winner of the election is the candidate winning in the majority of the districts. We focus on the setting with *no inter-agent externalities* where each receiver’s utility depends only on his action and the realized state of nature, but not on the other receivers’ actions. This assumption is well-motivated, as voting for the most preferred candidate is a weakly dominant strategy in two-candidate elections.

Two forms of signals are customarily investigated in the literature. With *private* signals, the sender can target different information to different receivers. Instead, with *public* signals, the sender can only communicate the same information to every receiver. Even if private persuasion may be more beneficial for the sender, sometimes, as in election settings where there are too many receivers, privately communicating to each receiver may be impracticable. At the same time, public communication is much easier to implement, *e.g.*, through TVs or newspapers. We introduce a new form of signaling, called *semi-public*, to model situations between private and public settings in district-based elections, where all the receivers of the same district observe the same signal, but the sender can target different information to different districts. Indeed, the voters are often reached by local communication shared with the voters in the same location, *e.g.*, local newspapers, electoral posters and rallies.

Original Contributions We study the efficiency and complexity of signaling in district-based elections with two candidates. First, we compare private, public, and semi-public signaling schemes in terms of efficiency when used to manipulate elections, showing that private signaling schemes perform arbitrarily better than (semi-)public schemes. Then, we show that optimal private signaling schemes can be computed efficiently, while the direct use of the results provided by Castiglioni, Celli, and Gatti (2020a) shows that the problem is inapproximable with (semi-)public signaling. However, we prove that multi-criteria Polynomial-Time Approximation Schemes (PTASs) for public and semi-public signaling schemes are possible when some relaxations are made. In particular, in the case of semi-public persuasion, we allow ϵ -persuasiveness and lower the number of districts to control needed to win the election by an arbitrary constant factor w.r.t. the majority. Instead, in the case of public persuasion, we also need to lower the number of votes needed to win in a district by an arbitrary constant factor w.r.t. the majority. In doing so, we introduce a novel property, namely *comparative stability*, and we design a bi-criteria PTAS to compute public signaling schemes in general Bayesian persuasion problems beyond district-based elections. Our result extends that by Xu (2020), allowing state-dependent sender’s utility functions and generalizing from stable to comparative stable sender’s utility functions.

Related Works The seminal model of Bayesian persuasion with a single receiver is introduced by Kamenica and Gentzkow (2011). This model is extended, allowing multiple receivers, by Bergemann and Morris (2016a), Bergemann and Morris (2016b), Wang (2013), and Taneva (2019). Furthermore, Alonso and Câmara (2016), Bardhi and Guo (2018) and Chan et al. (2019) provide the first attempts of applying the Bayesian persuasion framework to voting. In particular, Bardhi and Guo (2018) and Chan et al. (2019) study unanimity voting and k -voting rules, respectively, in settings with binary actions and state spaces. Instead, Alonso and Câmara (2016) employ a novel geometric tool to characterize an optimal public signaling scheme in voting.

In addition to the works mentioned above, which provide the economic groundings of Bayesian persuasion in (simple) voting settings, other works study election problems from a computational perspective. In particular, Arieli and Babichenko (2019) study the problem of private Bayesian persuasion with no inter-agent externalities. In the case of binary state spaces and k -voting rule, they provide a characterization of the optimal private signaling schemes.¹ Cheng et al. (2015) study the same k -voting problem with public persuasion, providing a polynomial-time approximation algorithm for a relaxed version of the problem in which the number of votes needed to win the election is reduced by an arbitrary constant factor and ϵ -persuasiveness is adopted. Castiglioni, Celli, and Gatti (2020a) extend the previous models to settings with an arbitrary number of states of nature and candidates. In particular, they prove that a private signaling scheme for k -voting can be computed in polynomial time, while the optimal public signaling scheme is NP-hard to approximate within any factor. Castiglioni, Celli, and Gatti (2020b) strengthen this hardness result, showing that finding a public signaling scheme that is approximately optimal and ϵ -persuasive requires quasi-polynomial time, assuming the exponential time hypothesis.

Our work is also closely related to Bayesian persuasion in general settings beyond elections and the relation between stability of the sender’s utility function and the computation of approximations. In particular, Cheng et al. (2015) introduce the notion of stability and show that this is a sufficient property to compute approximately optimal ϵ -persuasive signaling schemes in polynomial time. Xu (2020) extends this framework to incorporate α -approximable sender’s utility functions and shows that approximately optimal and ϵ -persuasive signaling schemes can be computed in polynomial time when the sender’s utility function is stable and independent of the state of nature.

Problem Formulation

In this section, we introduce the two frameworks we use in our work: district-based elections and Bayesian persuasion.

District-based Elections There is a set of candidates $C = \{c_0, c_1\}$ and a set of voters $R = \{r_1, \dots, r_{|R|}\}$ divided in a set D of districts. The set of voters of district $d \in D$ is denoted with R^d . Each voter casts a vote for one of the two candidates. Once the voters expressed their preferences, the election process proceeds in two steps. For the sake of simplicity, we study the basic case in which both steps follow a *majority-voting* rule.² The election works as follows.

1. For each $d \in D$, the votes expressed by all $r \in R^d$ are locally aggregated, and the candidate with the majority of the votes is elected as the winner of the district.
2. The outcomes of all the districts are aggregated, and the candidate that is the winner in the majority of the districts is chosen as the winner of the district-based election.

¹In k -voting, a candidate wins if he collects at least k votes.

²In majority voting, the candidate with the most votes wins.

We assume that the manipulator prefers c_0 to be the winner of the election. Let $\mathbf{c} \in \mathbf{C}$ be a tuple composed by the votes of all the voters, where $\mathbf{C} = C^{|R|}$. Similarly, \mathbf{c}^d is the tuple of the votes of the voters in district d . The manipulator's utility $\mathcal{W} : \mathbf{C} \rightarrow \{0, 1\}$ is defined as the composition of a collection of functions $W^d : C^{|R^d|} \rightarrow C$, each representing the majority voting run in district d , and the function $\overline{W} : C^{|D|} \rightarrow \{0, 1\}$, representing the majority voting that aggregates the outcomes of all the districts. We define $K_D = \lceil |D|/2 \rceil$ and, for each district d , $K_d = \lceil |R^d|/2 \rceil$. Then, \mathcal{W} is defined as $\mathcal{W}(\mathbf{c}) = \overline{W}(W^1(\mathbf{c}^1), \dots, W^D(\mathbf{c}^D))$, where $W^d(\mathbf{c}^d)$ assumes value c_0 if at least K_d of the voters in district d vote for candidate c_0 , and \overline{W} assumes value 1 if and only if c_0 wins in at least K_D districts.

We introduce some relaxations for the majority-voting rules W^d and \overline{W} . In the first relaxation, we allow the number of votes that the target candidate c_0 needs to win in each district d to be smaller than K_d . We denote with W_δ^d the resulting majority voting rule. Formally, $W_\delta^d : C^{|R^d|} \rightarrow C$ assumes value c_0 if at least $\lceil (1-\delta)K_d \rceil$ voters in district d vote for c_0 and c_1 otherwise. The manipulator's utility function of this first relaxed problem, denoted with \mathcal{W}_δ , is defined as $\mathcal{W}_\delta = \overline{W}(W_\delta^1(\mathbf{c}^1), \dots, W_\delta^D(\mathbf{c}^D))$. In the second, stronger relaxation, we also allow the number of districts that the target candidate c_0 needs to control to win the election to be smaller than K_D . We denote with \overline{W}_δ the resulting majority voting rule aggregating the outcomes of the districts. Formally, $\overline{W}_\delta : C^{|D|} \rightarrow \{0, 1\}$ assumes value 1 when c_0 wins in at least $\lceil (1-\delta)K_D \rceil$ districts. The manipulator's utility function of this second relaxed problem, denoted with $\mathcal{W}_{\delta\delta}$, is defined as $\mathcal{W}_{\delta\delta}(\mathbf{c}) = \overline{W}_\delta(W_\delta^1(\mathbf{c}^1), \dots, W_\delta^D(\mathbf{c}^D))$.

Bayesian Persuasion Framework Our model includes a sender (the manipulator) and a set R of receivers (voters) that must choose an action (a candidate) from the set $C = \{c_0, c_1\}$. Each voter r 's utility u_r depends only on his own action and a state of nature $\theta \in \Theta$ drawn from a prior distribution $\mu \in \Delta_\Theta$, where Δ_Θ is the set of probability distributions supported on Θ . In particular, we define $u_r : \Theta \times C \rightarrow [0, 1]$, where $u_r(\theta, c)$ expresses how much receiver r appreciates candidate c when the state of nature is θ . We use $u_r(\theta) = u_r(\theta, c_0) - u_r(\theta, c_1)$ to denote how much voter r prefers candidate c_0 over c_1 , in state of nature θ . In general Bayesian persuasion problems, the sender's utility, usually denoted with f_θ , depends on the state of nature θ and maps the receivers' action profiles to values in $[0, 1]$. In our setting, f_θ does not depend on θ and is set equal to $\mathcal{W}, \mathcal{W}_\delta, \mathcal{W}_{\delta\delta}$ depending on the specific problem we tackle.

The interaction among the sender and the receivers goes as follows (see Fig. 1). The sender commits to a randomized publicly known signaling scheme ϕ that maps states of nature to signals for the receivers. The signal set of a receiver r is denoted with S_r , while $s_r \in S_r$ is a signal for receiver r . The set of possible signals is then $\mathcal{S} = \times_{r \in R} S_r$, while a profile of signals is denoted with $\mathbf{s} = (s_1, \dots, s_{|R|})$. The sender observes the state of nature sampled from μ and computes $\mathbf{s} \in \mathcal{S}$ according to ϕ . After observing the signal s_r , each




Figure 1: Interaction between the sender and a receiver.

receiver r performs a Bayesian update, and infers a posterior belief $\mathbf{p}^r \in \mathcal{P}$ (where $\mathcal{P} = \Delta_\Theta$) as follows: the realized state of nature is θ with probability $p_\theta^r = \frac{\mu_\theta \phi(\theta, s_r)}{\sum_{\theta' \in \Theta} \mu_{\theta'} \phi(\theta', s_r)}$. Then, each receiver plays an action maximizing his expected utility according to posterior \mathbf{p}^r .

We introduce three forms of signaling schemes. A private signaling scheme exploits a private communication channel toward each receiver. Sometimes this assumption is not realistic, and the sender has only a single communication channel observed by all the receivers, *i.e.*, $s_r = s'_r$ for all $r, r' \in R$. We call these signaling schemes public. Finally, we introduce a novel form of communication that suits our election model, where the sender has a communication channel toward each district d , and all the receivers in the same district receive the same signal, *i.e.*, $s_r = s'_r$ for all $r, r' \in R^d$. We call these signaling schemes semi-public.

In all these settings, a revelation-principle style argument shows that there always exists a signaling scheme that is *direct* and *persuasive*. More precisely, a signaling scheme is direct if the signals are action recommendations, while it is persuasive if each receiver has the interest to follow the recommendations. Thus, a direct signaling scheme is a mapping $\phi : \Theta \rightarrow \Delta_C$, and $\phi(\theta, \mathbf{c})$ is the probability whereby the sender recommends \mathbf{c} in state θ . In order for the signaling scheme to be persuasive, the receivers must have an incentive to follow the recommendation. This is customarily assured by forcing constraints on ϕ depending on the specific form of signaling. In particular, the incentive constraints associated with a receiver r in a district d are:

- $\sum_{\theta, \mathbf{c}: c_r = c} \phi(\theta, \mathbf{c})(u_r(\theta, c) - u_r(\theta, c')) \geq 0 \quad \forall c, c' \in C$ (private signaling);
- $\sum_{\theta} \phi(\theta, \mathbf{c})(u_r(\theta, c_r) - u_r(\theta, c')) \geq 0 \quad \forall \mathbf{c} \in \mathbf{C}, c' \in C$ (public signaling);
- $\sum_{\theta, \mathbf{c}: \mathbf{c}^d = \bar{\mathbf{c}}} \phi(\theta, \mathbf{c})(u_r(\theta, \bar{c}_r) - u_r(\theta, c')) \geq 0 \quad \forall \bar{\mathbf{c}} \in C^{|R^d|}, c' \in C$ (semi-public signaling).

Similarly, a direct signaling scheme is ϵ -persuasive if the incentive constraints are violated by at most ϵ .

Finally, we state the optimization problems we study in this paper. PRIVATE-DBE is the problem of designing a private signaling scheme maximizing the probability of having candidate c_0 elected in district-based elections. PUBLIC-DBE and SEMIPUBLIC-DBE refer to the same problem with public and semi-public signaling, respectively.

An Example of Inefficiency of (Semi-)Public Persuasion

To clarify better the Bayesian persuasion framework, we provide an example of its application to majority voting

without districts. This example is also useful to show that the restriction to (semi-)public signaling can decrease the sender's utility by an arbitrarily large factor.

Example 1. Consider a (non-relaxed) majority-voting election with seven voters $R = \{r_1, r_2, r_3, r_4, r_5, r_6, r_7\}$ and two candidates $C = \{c_0, c_1\}$. The objective of the sender is to maximize the probability with which candidate c_0 is elected. Therefore, he needs to persuade at least half of the voters (i.e., $\lceil |R|/2 \rceil = 4$) to make candidate c_0 be the winner. There are three states of nature, namely, $\Theta = \{\theta_A, \theta_B, \theta_C\}$, and each state is equally probable. Tab. 1 provides the parameters $u_r(\theta)$ of the voters, defined as $u_r(\theta) = u_r(\theta, c_0) - u_r(\theta, c_1)$ and capturing the net payoff of voter r from having candidate c_0 elected, in state of nature θ .

	State θ_A	State θ_B	State θ_C	
Voters	r_1, r_2	+1/2	-1	-1
	r_3, r_4	-1	+1/2	-1
	r_5, r_6	-1	-1	+1/2
	r_7	+1/2	+1/2	+1/2

Table 1: Payoffs of the voters in Example 1.

The sender can design a direct and persuasive private signaling scheme such that at least four voters prefer candidate c_0 over c_1 for every signal profile s . Hence, this scheme ensures that candidate c_0 is elected with a probability of 1. Specifically, in each state θ the scheme recommends candidate c_0 to every voter r with utility $u_r(\theta) \geq 0$ and to one voter among those with $u_r(\theta) < 0$ chosen randomly with uniform probability. It is easy to see that this private signaling scheme satisfies the incentive constraints. Consider, for example, voter r_1 . The marginal probabilities with which he is recommended to vote for candidate c_0 are: $\phi_1(\theta_A, c_0) = 1$, $\phi_1(\theta_B, c_0) = 1/4$ and $\phi_1(\theta_C, c_0) = 1/4$. Therefore, when he receives the recommendation to vote for c_0 , he has a posterior distribution \mathbf{p} with $p_{\theta_A} = \frac{\mu_{\theta_A} \cdot \phi_1(\theta_A, c_0)}{\sum_{\theta \in \Theta} \mu_{\theta} \cdot \phi_1(\theta, c_0)} = \frac{1/3}{1/3+1/3 \cdot 1/4+1/3 \cdot 1/4} = 2/3$

and $p_{\theta_B} = p_{\theta_C} = 1/6$. Thus, the voter has expected utility $u(\theta_A)p_{\theta_A} + u(\theta_B)p_{\theta_B} + u(\theta_C)p_{\theta_C} = 0$ and will follow the recommendation. Similarly, we can show that the incentive constraints associated with the other voters are satisfied.

We switch to public signals and we show that we cannot design a public signaling scheme that guarantees candidate c_0 to be elected with positive probability. Any public signaling scheme making candidate c_0 win the election with positive probability must assign a strictly positive probability to at least one signal that makes at least four voters prefer candidate c_0 over c_1 . We show that we cannot design such a public signal. In particular, we show that there is no posterior $\mathbf{p} \in \mathcal{P}$ that provides an expected utility larger than or equal to zero to at least four voters.³ Since receiver r_7 prefers candidate c_0 in every state of nature, he votes for

³Recall that in a public signaling scheme, all the receivers observe the same signal, perform the same update of the belief, and have the same posterior belief.

c_0 independently from the posterior induced by the signal. Therefore, it is sufficient to persuade three voters among the first six. Suppose that voters r_1 and r_2 vote for c_0 . This implies that $p_{\theta_A}/2 - p_{\theta_B} - p_{\theta_C} = p_{\theta_A}/2 - (1 - p_{\theta_A}) \geq 0$ and $p_{\theta_A} \geq 2/3$. Suppose, by contradiction, that also voters r_3 and r_4 vote for c_0 . This requires that $-p_{\theta_A} + p_{\theta_B}/2 - p_{\theta_C} \geq 0$ and $p_{\theta_B} \geq 2/3$, reaching a contradiction with $\mathbf{p} \in \mathcal{P}$. It is easy to see that, by the symmetry of the instance, all the other sets of four voters cannot vote for c_0 at the same time.

From the previous example, we can state the following:

Proposition 1. There is an instance of majority-voting election in which the optimal private signaling scheme guarantees that candidate c_0 wins the election with a probability of 1, while the optimal public signaling scheme cannot guarantee a winning probability strictly larger than 0.

This inefficiency result can be easily generalized to the case of public and semi-public signaling scheme in district-based elections. Indeed, with only a single district, semi-public signals correspond to public signals and a district-based election reduces to a simple majority-voting election as the one presented above.

Private Persuasion in District-based Elections

In this section, we show that an optimal private signaling scheme for district-based elections can be found in polynomial time. Our result is built upon the previous works by Arieli and Babichenko (2019) and Castiglioni, Celli, and Gatti (2020a) on k -voting. Let $a_{d,\theta}$ be the probability with which K_d voters vote for c_0 in district d when the state of nature is θ . Similarly, let α_θ be the probability that c_0 wins in at least K_D districts with state of nature θ . Finally, given a direct private signaling scheme ϕ , we denote with $\phi_r(\theta, c) = \sum_{\mathbf{c}:c_r=c} \phi(\theta, \mathbf{c})$ the marginal probabilities of ϕ whereby c is recommended to r with state of nature θ . We can compute an optimal private signaling scheme by LP (1) (all the proofs are in the Supplemental Material).

Theorem 1. LP (1) computes an optimal solution of PRIVATE-DBE in polynomial time.

Proof sketch. Constraints (1b) force the marginal probabilities $\phi_r(\theta, c_0)$ of the signaling scheme ϕ to satisfy the incentive constraints. For each state of nature θ , the maximum probability $a_{d,\theta}$ with which at least K_d receivers in R^d vote for c_0 given marginal probabilities $\phi_r(\theta, c_0)$ is:

$$a_{d,\theta} = \min \left\{ \min_{m \in \{0, \dots, K_d-1\}} \frac{1}{K_d - m} v_{\theta, m}; 1 \right\},$$

where $v_{\theta, m}$ is the sum of the lowest $|R^d| - m$ elements in the set $\{\phi_r(\theta, c)\}_{r \in R^d}$; for further details, see Arieli and Babichenko (2019). The above equation is enforced via Constraints (1f). Constraints (1g) and (1h) ensure that the values of $v_{\theta, m}$ are consistent with the values of the other variables. The computation of the maximum probability α_θ with which at least K_D districts elect c_0 given probabilities $a_{d,\theta}$ is similar to the computation of $a_{d,\theta}$ given $\phi_r(\theta, c_0)$. This is enforced by Constraints (1c), (1d), and (1e). Finally, Objective (1a) maximizes the sum over all $\theta \in \Theta$ of the prior probability multiplied by α_θ , i.e., the probability that c_0 wins when the state of nature is θ . \square

$$\begin{aligned}
& \max_{\substack{\alpha \in [0,1]^{|\Theta|}, a \in [0,1]^{|D| \times |\Theta|} \\ i, l \in \mathbb{R}^{|\Theta| \times K_D}, o \in \mathbb{R}^{|D| \times |\Theta| \times K_D} \\ t_{d,\theta,m}, v_{d,\theta,m} \in \mathbb{R} \forall d \in D, \theta \in \Theta, m \in \{1, \dots, K_d\} \\ z_{d,\theta,r,m} \in \mathbb{R} \forall d \in D, \theta \in \Theta, r \in R, m \in \{1, \dots, K_d\} \\ \phi_r(\cdot, c_0) \in [0,1]^{|\Theta|} \forall r \in R}} \sum_{\theta \in \Theta} \mu_\theta \alpha_\theta \quad (1a) \\
\text{s.t. } & \sum_{\theta \in \Theta} \mu_\theta \phi_r(\theta, c_0) u_r(\theta) \geq 0 \quad \forall r \in R \quad (1b) \\
& \alpha_\theta \leq \frac{1}{K_D - m} i_{\theta,m} \quad (1c) \\
& \quad \forall \theta \in \Theta, \forall m \in \{0, \dots, K_D - 1\} \\
& i_{\theta,m} \leq (|D| - m) l_{\theta,m} + \sum_{d \in D} o_{d,\theta,m} \quad (1d) \\
& \quad \forall \theta \in \Theta, \forall m \in \{0, \dots, K_D - 1\} \\
& a_{d,\theta} \geq l_{\theta,m} + o_{d,\theta,m} \quad (1e) \\
& \quad \forall d \in D, \forall \theta \in \Theta, \forall m \in \{0, \dots, K_D - 1\} \\
& a_{d,\theta} \leq \frac{1}{K_d - m} v_{d,\theta,m} \quad (1f) \\
& \quad \forall d \in D, \forall \theta \in \Theta, \forall m \in \{0, \dots, K_d - 1\} \\
& v_{d,\theta,m} \leq (|R^d| - m) t_{d,\theta,m} + \sum_{r \in R^d} z_{d,\theta,r,m} \quad (1g) \\
& \quad \forall d \in D, \forall \theta \in \Theta, \forall m \in \{0, \dots, K_d - 1\} \\
& \phi_r(\theta, c_0) \geq t_{d,\theta,m} + z_{d,\theta,r,m} \quad (1h) \\
& \quad \forall d \in D, \forall r \in R^d, \forall \theta \in \Theta, \forall m \in \{0, \dots, K_d - 1\}
\end{aligned}$$

Public and Semi-public Persuasion in District-based Elections

We turn our attention to the design of optimal public and semi-public signaling schemes. There is a sharp distinction between the nature of these problems and that one of private signaling. Indeed, in addition to being inefficient w.r.t. private signals (see Proposition 1), optimal (semi-)public signaling schemes are also inapproximable. The hardness follows from previous results with public signaling. Specifically, Castiglioni, Celli, and Gatti (2020a) prove that it is NP-hard to approximate the optimal public signaling scheme within any factor in elections with majority voting. The extension of this hardness result to public and semi-public signaling in district-based elections is direct as a district-based election reduces to majority-voting when there is only a single district. Thus, we focus on possible relaxations that make the problem computationally tractable.

Motivated by the fact that voters are somewhat biased to follow the sender's recommendations, several works relax the incentive constraints allowing the receivers to vote for the target candidate even if other candidates give them a slightly better expected utility (ϵ -persuasiveness). Recently, Castiglioni, Celli, and Gatti (2020b) prove that even allowing this relaxation the problem of designing an approximate public signaling scheme remains intractable with majority voting. Therefore, we focus on other different relaxations. In particular, Cheng et al. (2015) employ two forms of relaxation, adopting ϵ -persuasiveness and lowering the

number of votes needed to win the election by an arbitrary constant factor. With these two relaxations, they prove that an approximate public signaling scheme with majority-voting can be computed efficiently. We prove that, adapting these two relaxations to our settings, both PUBLIC-DBE and SEMIPUBLIC-DBE admit a multi-criteria PTAS. As a preliminary step, we prove some results on the relation between the notion of stability and the design of approximately optimal signaling schemes that are of general interest in Bayesian persuasion beyond elections.

Comparative Stability and Public Signaling Schemes

We refer to the notion of stability of a function introduced by Xu (2020). In particular, a function is said stable if, for every action profile, the introduction of small perturbations leads to small changes in the value of the function. Here, we extend the notion of stability to pairs of functions, and we call it comparative. Our extension is such that comparative stability corresponds to (simple) stability in the degenerate case in which the two functions of the pair are the same. Furthermore, if function g satisfies the comparative stability property w.r.t. function h , we also say that g is β -stable compared with h . Initially, we introduce the notion of perturbation by the concept of α -noisy distribution.

Definition 1. Let $\mathbf{c} \in \mathbf{C}$ be an action profile and \mathbf{y} be a probability distribution supported on $\Delta_{\mathbf{C}}$. For any $\alpha \in (0, 1]$, we say that \mathbf{y} is an α -noisy distribution around \mathbf{c} if for all $i \in \{1, \dots, n\}$: $\Pr_{\tilde{\mathbf{y}} \sim \mathbf{y}}[\tilde{y}_i \neq c_i] \leq \alpha$.

Hence, an α -noisy distribution bounds the *marginal probability* of any single element of $\{1, \dots, n\}$ to be corrupted. However, no assumption is made on how the corruptions of the elements correlate with each other. Now, we define our notion of comparative stability.

Definition 2. Given two functions $g, h : \mathbf{C} \rightarrow [0, 1]$ and a real number $\beta \geq 0$, we say that g is β -stable compared with h if and only if the following holds for all action profiles $\mathbf{c} \in \mathbf{C}$, $\alpha \in (0, 1]$, and α -noisy distributions \mathbf{y} around \mathbf{c} :

$$\mathbb{E}_{\tilde{\mathbf{y}} \sim \mathbf{y}} [g(\tilde{\mathbf{y}})] \geq h(\mathbf{c})(1 - \alpha\beta).$$

Intuitively, if g satisfies the comparative stability property w.r.t. h , then, for every action profile, the value of h in that action profile is close to the value of g in the corresponding perturbed action profile.

We exploit the notion of comparative stability to design an efficient algorithm that computes approximate public signaling schemes. More precisely, we study a generic multi-agent Bayesian persuasion problem, where the sender faces a set of receivers R , and each receiver needs to choose an action between a couple of alternatives. Let g, h be two sets of arbitrary functions depending on the state of nature θ and denoted with $g_\theta : \mathbf{C} \rightarrow [0, 1]$ and $h_\theta : \mathbf{C} \rightarrow [0, 1]$, respectively. According to Definition 2, we say that g is β -stable compared with h if g_θ is β -stable with respect to h_θ for all the states of nature $\theta \in \Theta$.

For the sake of clarity, in the following, we use indirect signaling schemes, and we express a signaling scheme as a weighted set of posteriors to which the receivers respond at best. Now, we describe the optimal behavior of the receivers.

Definition 3 (Receivers' behavior with persuasiveness). Given a set of functions $\{f_\theta\}_{\theta \in \Theta}$ such that $f_\theta : \mathbf{C} \rightarrow [0, 1]$, the receivers' optimal behavior $\mathbf{b}^{\mathbf{P}} \in \mathbf{C}$ with persuasiveness given posterior $\mathbf{p} \in \mathcal{P}$ is as follows. Let:

- $A = \{r \in R : \sum_{\theta} p_{\theta} u_r(\theta) > 0\}$ the set of receivers whose unique best response is action c_0 ,
- $B = \{r \in R : \sum_{\theta} p_{\theta} u_r(\theta) < 0\}$ the set of receivers whose unique best response is action c_1 ,
- $E = \{r \in R : \sum_{\theta} p_{\theta} u_r(\theta) = 0\}$ the set of receivers who are indifferent between action c_0 and c_1 .

Then, we have:

$$\mathbf{b}^{\mathbf{P}} = \arg \max_{\mathbf{c} \in \mathbf{C}: c_r = c_0 \forall r \in A, c_r = c_1 \forall r \in B} \sum_{\theta} p_{\theta} f_{\theta}(\mathbf{c}).$$

Similarly, we define the notion of ϵ -best response.

Definition 4 (Receivers' behavior with ϵ -persuasiveness). Given a set of functions $\{f_\theta\}_{\theta \in \Theta}$ such that $f_\theta : \mathbf{C} \rightarrow [0, 1]$, the receivers' optimal behavior $\mathbf{b}^{\mathbf{P}, \epsilon} \in \mathbf{C}$ with ϵ -persuasiveness given posterior $\mathbf{p} \in \mathcal{P}$ is as follows. Let:

- $A_{\epsilon} = \{r \in R : \sum_{\theta} p_{\theta} u_r(\theta) > \epsilon\}$ the set of receivers whose unique best response is action c_0 ,
- $B_{\epsilon} = \{r \in R : \sum_{\theta} p_{\theta} u_r(\theta) < -\epsilon\}$ the set of receivers whose unique best response is action c_1 ,
- $E_{\epsilon} = \{r \in R : \sum_{\theta} p_{\theta} u_r(\theta) \in [-\epsilon, \epsilon]\}$ the set of receivers who are indifferent between action c_0 and c_1 .

Then, we have:

$$\mathbf{b}^{\mathbf{P}, \epsilon} = \arg \max_{\mathbf{c} \in \mathbf{C}: c_r = c_0 \forall r \in A_{\epsilon}, c_r = c_1 \forall r \in B_{\epsilon}} \sum_{\theta} p_{\theta} f_{\theta}(\mathbf{c}).$$

Now, we show that computing a direct public signaling scheme is equivalent to derive a Bayes plausible distribution of posteriors $\gamma \in \Delta_{\mathcal{P}}$ that maximizes the sender's utility. Let $\text{supp}(\gamma)$ denote the set of posteriors induced with strictly positive probability. Similarly, let $\text{supp}(\phi)$ denote the set of posteriors induced by ϕ with strictly positive probability. Finding a public signaling scheme is equivalent to finding a probability distribution $\gamma \in \Delta_{\mathcal{P}}$ on the set of posteriors \mathcal{P} such that $\sum_{p \in \text{supp}(\gamma)} \gamma_p p_{\theta} = \mu_{\theta}$ for every $\theta \in \Theta$. Given a well-defined distribution over posteriors γ , we can recover a direct signaling schemes ϕ that induces such a probability distribution by setting $\phi_{\theta}(\mathbf{c}) = \sum_{\mathbf{p} \in \text{supp}(\gamma): \mathbf{c} = \mathbf{b}^{\mathbf{P}}} \gamma_p p_{\theta}$. For this reason, in the following, we represent signaling schemes as probability distributions on the posteriors. We introduce some further notation. For every $\mathbf{p} \in \mathcal{P}$ and set of functions $f = \{f_{\theta}\}_{\theta \in \Theta}$, we define the sender's expected utility with persuasiveness as $f(\mathbf{p}) = \sum_{\theta} p_{\theta} f_{\theta}(\mathbf{b}^{\mathbf{P}})$, and with ϵ -persuasiveness as $f_{\epsilon}(\mathbf{p}) = \sum_{\theta} p_{\theta} f_{\theta}(\mathbf{b}^{\mathbf{P}, \epsilon})$. Finally, we define q -uniform probability distributions as follows.

Definition 5. A probability distribution $\mathbf{x} \in \Delta_X$ is q -uniform if and only if it is the average of a multiset of q basis vectors in $|X|$ -dimensional space.

Therefore, we say that a probability distribution $\mathbf{p} \in \mathcal{P}$ is q -uniform if each of its entry p_{θ} is a multiple of $1/q$. Moreover, we use the notation $\mathcal{Q} \subset \Delta_{\Theta}$ to denote the set of all q -uniform distributions over Θ . Our first result shows that we can decompose each posterior in a convex combination $\gamma \in \Delta_{\mathcal{Q}}$ of q -uniform posteriors (with q constant), such that

$\sum_{\mathbf{p} \in \mathcal{Q}} \gamma_p g_{\epsilon}(\mathbf{p})$ closely approximates $h(\mathbf{p}^*)$. This is a generalization of the result by Xu (2020) to state-dependent utility functions (and couples of functions), and it is crucial to prove the following results.

Lemma 1. Let $\beta, \epsilon > 0, \eta \in (0, 1]$ and set $q = 32 \log \left(\frac{4}{\eta \min\{1, 1/\beta\}} \right) / \epsilon^2$. Then, given a posterior $\mathbf{p}^* \in \mathcal{P}$ and two sets of functions g, h with g β -stable compared with h , there exists a $\gamma \in \Delta_{\mathcal{Q}}$ with $\sum_{\mathbf{p} \in \mathcal{Q}} \gamma_p \mathbf{p} = \mathbf{p}^*$ and

$$\sum_{\mathbf{p} \in \mathcal{Q}} \gamma_p \sum_{\theta} p_{\theta} g_{\theta}(\mathbf{b}^{\mathbf{P}, \epsilon}) \geq (1 - \eta) \sum_{\theta} p_{\theta}^* h_{\theta}(\mathbf{b}^{\mathbf{P}^*}). \quad (2)$$

Now, we can prove the main result of this section. Consider a couple of sets of functions g, h where g is β -stable compared with h . With abuse of notation, we define $g(\phi)$ and $h(\phi)$ as the functions which evaluate the expected sender's utility of a public signaling scheme ϕ with g and h , respectively. We can resort to Lemma 1 to state the following result. The proof is based on solving a linear program that works only with q -uniform posteriors.

Theorem 2. Let $\beta, \epsilon > 0$ and $\eta \in (0, 1]$. Consider two arbitrary state-dependent sets of functions g, h such that $g_{\theta} : \mathbf{C} \rightarrow [0, 1]$ is β -stable compared with $h_{\theta} : \mathbf{C} \rightarrow [0, 1]$ for all $\theta \in \Theta$. Then there exists a poly $\left(|R| |\Theta|^{\log \left(\frac{1}{\eta \min\{1, 1/\beta\}} \right) / \epsilon^2} \right)$ time algorithm that returns an ϵ -persuasive public signaling scheme ϕ_{ϵ} such that:

$$g(\phi_{\epsilon}) \geq (1 - \eta) \max_{\phi \in \Phi} h(\phi),$$

where Φ is the set of persuasive signaling schemes.

By setting $h = g$, we obtain a generalization of the result by Xu (2020) to state-dependent functions.

Comparative Stability of Voting Functions We apply this novel concept of stability to voting problems. Our first result proves that the two relaxed majority-voting functions previously introduced satisfy the comparative stability property. This result is similar to that by Cheng et al. (2015). However, we use multiplicative factors (in place of additive factors) and prove a slightly stronger result than stability. In particular, we prove that the decrease in utility is small even if only the perturbations from action c_0 to c_1 are bounded.

Lemma 2. W_{δ} is $1/\delta$ -stable compared with W . Moreover, for all $\mathbf{c} \in \mathbf{C}$, $r \in R$, $\alpha \in (0, 1]$, and $\mathbf{y} \in \Delta_{\mathbf{C}}$ such that $\Pr_{\mathbf{y}}(\tilde{y}_r = c_1 \wedge c_r = c_0) \leq \alpha$, it holds:

$$\mathbb{E}_{\tilde{\mathbf{y}} \sim \mathbf{y}} [W_{\delta}(\tilde{\mathbf{y}})] \geq W(\mathbf{c}) \left(1 - \frac{\alpha}{\delta} \right).$$

We can use the result above to prove that $W_{\delta\delta}$ satisfies the property of comparative stability with respect to W . Intuitively, the result follows from the observation that W is the composition of two majority-voting steps.

Lemma 3. $W_{\delta\delta}$ is $\frac{1}{\delta^2}$ -stable with respect to W .

Finally, we derive a stronger decomposition lemma for majority-voting. Specifically, Lemma 1 shows that the decrease in the expected sender's utility when decomposing a posterior in q -uniform posteriors can be bounded. However, in generic settings, the sender's expected utility in a given state of nature can change arbitrarily. This is not the case in

majority voting, where, instead, this decrease is bounded. In particular, we can show the following, that is crucial when addressing the SEMIPUBLIC-DBE problem.

Lemma 4. *Let $\epsilon > 0, \eta \in (0, 1]$ and set $q = 32 \log\left(\frac{4}{\eta\delta}\right) / \epsilon^2$. Then, given a posterior $\mathbf{p}^* \in \mathcal{P}$, there exists a $\gamma \in \Delta_{\mathcal{Q}}$ with $\sum_{\mathbf{p} \in \mathcal{Q}} \gamma_{\mathbf{p}} \mathbf{p} = \mathbf{p}^*$ and*

$$\sum_{\mathbf{p} \in \mathcal{Q}} \gamma_{\mathbf{p}} p_{\theta} W_{\delta}(\mathbf{b}^{\mathbf{p}, \epsilon}) \geq (1 - \eta) p_{\theta}^* W(\mathbf{b}^{\mathbf{p}^*}) \quad \forall \theta \in \Theta.$$

Computing Public and Semi-public Signaling Schemes in District-based Elections We present two multi-criteria PTASs for the SEMIPUBLIC-DBE and PUBLIC-DBE problems, respectively, when our relaxations are adopted. First, we focus on the problem of designing public signaling schemes. We assume ϵ -persuasive signaling schemes, and we replace function \mathcal{W} with \mathcal{W}_{δ} (this corresponds to relaxing both the majority voting within every single district and the majority voting aggregating the outcomes of all the districts). Let $\mathcal{W}(\phi)$ and $\mathcal{W}_{\delta}(\phi)$ denote the functions returning the sender's expected utility provided by a public signaling scheme ϕ with voting rules \mathcal{W} and \mathcal{W}_{δ} , respectively. We show that it is possible to compute efficiently an ϵ -persuasive public signaling scheme ϕ_{ϵ} that approximates the optimal persuasive signaling scheme with an approximation factor arbitrarily close to 1. Since the relaxed function \mathcal{W}_{δ} is $1/\delta^2$ -stable compared to the non-relaxed function \mathcal{W} by Theorem 3, we can immediately apply Theorem 2 to these functions and then derive the following.

Corollary 1. *Let $\epsilon > 0, \delta \in (0, 1)$ and $\eta \in (0, 1]$, then there exists a poly $\left(|R| |\Theta|^{\log\left(\frac{1}{\eta\delta^2}\right) / \epsilon^2}\right)$ time algorithm that returns an ϵ -persuasive public signaling scheme ϕ_{ϵ} such that:*

$$\mathcal{W}_{\delta}(\phi_{\epsilon}) \geq (1 - \eta) \max_{\phi \in \Phi} \mathcal{W}(\phi), \quad (3)$$

where Φ is the set of persuasive signaling schemes.

Then, we focus on the SEMIPUBLIC-DBE problem. As highlighted above, to overcome the intractability result, also in this setting, it is necessary to relax the problem. Specifically, we use ϵ -persuasive signaling schemes and we replace function \mathcal{W} with \mathcal{W}_{δ} (this corresponds to relaxing the majority voting aggregating the outcomes of all the districts). We show that it is possible to compute efficiently an ϵ -persuasive semi-public signaling scheme ϕ_{ϵ} that approximates the optimal persuasive signaling scheme with an approximation factor arbitrarily close to 1. Computing a semi-public signaling scheme ϕ amounts to determining a collection $\{\phi_d\}_{d \in D}$ of $|D|$ public signaling schemes, one for each district, and correlate them. The crucial point concerns the computation of good marginal probabilities of the signaling scheme. Indeed, their aggregation is equivalent to computing a private signaling scheme in majority-voting elections, and this can be done efficiently (see LP (1) and Theorem 1). The main idea of our proof is that there are approximately optimal marginal probabilities of the signaling scheme that use only q -uniform posteriors (with q constant). Let α_{θ} be the probability that c_0 wins in at least K_D districts with state of nature θ , $\alpha_{d, \theta}^{\delta}$ be the probability that candidate c_0 receives at least $\lceil (1 - \delta) K_d \rceil$

votes in district d with state of nature θ , and γ^d be a probability distribution over posteriors for the receivers in district d . Finally, let $\mathbb{I}[\mathcal{E}]$ denote the indicator function for the event \mathcal{E} . Then, the following formulation computes an approximately optimal signaling scheme in polynomial time.

$$\max_{\substack{\alpha \in [0, 1]^{|\Theta|}, a^{\delta} \in [0, 1]^{|D| \times |\Theta|} \\ i, l \in \mathbb{R}^{|\Theta| \times K_D}, o \in \mathbb{R}^{|D| \times |\Theta| \times K_D} \\ \gamma^d \in \Delta_{\mathcal{Q}} \forall d \in D}} \sum_{\theta \in \Theta} \mu_{\theta} \alpha_{\theta} \quad (4a)$$

$$\text{s.t. } \alpha_{\theta} \leq \frac{1}{K_D - m} i_{\theta, m} \quad \forall \theta \in \Theta, \forall m \in \{0, \dots, K_D - 1\} \quad (4b)$$

$$i_{\theta, m} \leq (|D| - m) l_{\theta, m} + \sum_{d \in D} o_{d, \theta, m} \quad \forall \theta \in \Theta, \forall m \in \{0, \dots, K_D - 1\} \quad (4c)$$

$$\alpha_{d, \theta}^{\delta} \geq l_{\theta, m} + o_{d, \theta, m} \quad \forall d \in D, \forall \theta \in \Theta, \forall m \in \{0, \dots, K_D - 1\} \quad (4d)$$

$$\alpha_{d, \theta}^{\delta} \leq \sum_{\mathbf{p} \in \mathcal{Q}} \frac{\gamma_{\mathbf{p}}^d p_{\theta}}{\mu_{\theta}} \mathbb{I}(W_{\delta}^d(\mathbf{b}^{\mathbf{p}, \epsilon}) = c_0) \quad \forall d \in D, \forall \theta \in \Theta \quad (4e)$$

$$\sum_{\mathbf{p} \in \mathcal{Q}} \gamma_{\mathbf{p}}^d p_{\theta} = \mu_{\theta} \quad \forall d \in D, \forall \theta \in \Theta \quad (4f)$$

Theorem 3. *Let $\epsilon > 0, \delta \in (0, 1)$ and $\eta \in (0, 1]$, then there exists a poly $\left(|R| |\Theta|^{\log\left(\frac{1}{\eta\delta}\right) / \epsilon^2}\right)$ time algorithm that outputs an ϵ -persuasive semi-public signaling scheme ϕ_{ϵ} such that:*

$$\mathcal{W}_{\delta}(\phi_{\epsilon}) \geq (1 - \eta) \max_{\phi \in \Phi} \mathcal{W}(\phi), \quad (5)$$

where Φ is the set of persuasive signaling schemes.

Conclusions and Future Works

In this paper, we study how a manipulator can exploit his information advantage to manipulate a district-based election through the strategic provision of information to rational voters. We show that private signaling schemes can be computed efficiently while computing optimal (semi-)public signaling schemes is intractable. However, we show that reasonable relaxations allow the design of multi-criteria PTASs for (semi-)public persuasion. An interpretation of these relaxations is that (semi-)public signaling is often tractable, except when the target candidate wins in at least half of the districts, but it is impossible to make slightly more than half of them elect such a candidate. In most cases, the receivers are slightly biased to follow the sender recommendations, and the manipulator's preferred candidate can either win or not win by at least a small, but not negligible, margin. With these assumptions, our algorithm approximates arbitrarily well the optimal signaling scheme in polynomial time.

In the future, we will study classes of instances in which optimal (semi-)public signaling schemes can be computed efficiently. We are also interested in settings in which the sender is uncertain about the voters' preferences.

Acknowledgments

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Market”.

References

- Allcott, H.; and Gentzkow, M. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives* 31(2): 211–236.
- Alonso, R.; and Câmara, O. 2016. Persuading voters. *American Economic Review* 106(11): 3590–3605.
- Arieli, I.; and Babichenko, Y. 2019. Private bayesian persuasion. *J ECON THEORY* 182: 185–217.
- Bardhi, A.; and Guo, Y. 2018. Modes of persuasion toward unanimous consent. *Theoretical Economics* 13(3): 1111–1149.
- Bergemann, D.; and Morris, S. 2016a. Bayes correlated equilibrium and the comparison of information structures in games. *THEOR ECON* 11(2): 487–522.
- Bergemann, D.; and Morris, S. 2016b. Information design, Bayesian persuasion, and Bayes correlated equilibrium. *AM ECON REV* 106(5): 586–591.
- Castiglioni, M.; Celli, A.; and Gatti, N. 2020a. Persuading Voters: It’s Easy to Whisper, It’s Hard to Speak Loud. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 1870–1877. AAAI Press. URL <https://aaai.org/ojs/index.php/AAAI/article/view/5555>.
- Castiglioni, M.; Celli, A.; and Gatti, N. 2020b. Public Bayesian Persuasion: Being Almost Optimal and Almost Persuasive. *ArXiv abs/2002.05156*.
- Castiglioni, M.; Ferraioli, D.; and Gatti, N. 2020. Election Control in Social Networks via Edge Addition or Removal. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 1878–1885. AAAI Press. URL <https://aaai.org/ojs/index.php/AAAI/article/view/5556>.
- Chan, J.; Gupta, S.; Li, F.; and Wang, Y. 2019. Pivotal persuasion. *Journal of Economic Theory* 180: 178–202.
- Chen, J.; Faliszewski, P.; Niedermeier, R.; and Talmon, N. 2017. Elections with few voters: candidate control can be easy. *J ARTIF INTELL RES* 60: 937–1002.
- Cheng, Y.; Cheung, H. Y.; Dughmi, S.; Emamjomeh-Zadeh, E.; Han, L.; and Teng, S. 2015. Mixture Selection, Mechanism Design, and Signaling. In *FOCS*, 1426–1445.
- Erdélyi, G.; Reger, C.; and Yang, Y. 2020. The complexity of bribery and control in group identification. *Autonomous Agents and Multi-Agent Systems* 34(1): 8.
- Faliszewski, P.; Gonen, R.; Koutecký, M.; and Talmon, N. 2018. Opinion Diffusion and Campaigning on Society Graphs. In *IJCAI*, 219–225.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2011. Multimode control attacks on elections. *Journal of Artificial Intelligence Research* 40: 305–351.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2009. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research* 35: 275–341.
- Guess, A.; Nyhan, B.; and Reifler, J. 2018. Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign. *European Research Council* 9.
- Kamenica, E.; and Gentzkow, M. 2011. Bayesian persuasion. *AM ECON REV* 101(6): 2590–2615.
- Liu, H.; Feng, H.; Zhu, D.; and Luan, J. 2009. Parameterized computational complexity of control problems in voting systems. *THEOR COMPUT SCI* 410(27-29): 2746–2753.
- Loreggia, A.; Narodytska, N.; Rossi, F.; Venable, K. B.; and Walsh, T. 2015. Controlling elections by replacing candidates or votes. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 1737–1738.
- Sina, S.; Hazon, N.; Hassidim, A.; and Kraus, S. 2015. Adapting the social network to affect elections. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 705–713.
- Taneva, I. 2019. Information Design. *American Economic Journal: Microeconomics* 11(4): 151–85. doi:10.1257/mic.20170351. URL <https://www.aeaweb.org/articles?id=10.1257/mic.20170351>.
- Wang, Y. 2013. Bayesian persuasion with multiple receivers. Available at SSRN 2625399 .
- Wilder, B.; and Vorobeychik, Y. 2018. Controlling elections through social influence. In *Proceedings of the 17th international conference on autonomous agents and multi-agent systems*, 265–273. International Foundation for Autonomous Agents and Multiagent Systems.
- Wilder, B.; and Vorobeychik, Y. 2019. Defending elections against malicious spread of misinformation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2213–2220.
- Xu, H. 2020. On the Tractability of Public Persuasion with No Externalities. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2708–2727. SIAM.