

# Persuading Voters: It’s Easy to Whisper, It’s Hard to Speak Loud

Matteo Castiglioni, Andrea Celli, Nicola Gatti

Politecnico di Milano, Piazza Leonardo da Vinci 32, I-20133, Milan, Italy  
name.surname@polimi.it

## Abstract

We focus on the following natural question: *is it possible to influence the outcome of a voting process through the strategic provision of information to voters who update their beliefs rationally?* We investigate whether it is computationally tractable to design a signaling scheme maximizing the probability with which the sender’s preferred candidate is elected. We resort to the model recently introduced by Arieli and Babichenko (2019) (*i.e.*, without inter-agent externalities), and focus on, as illustrative examples,  $k$ -voting rules and plurality voting. There is a sharp contrast between the case in which *private* signals are allowed and the more restrictive setting in which only *public* signals are allowed. In the former, we show that an optimal signaling scheme can be computed efficiently both under a  $k$ -voting rule and plurality voting. In establishing these results, we provide two contributions applicable to general settings beyond voting. Specifically, we extend a well-known result by Dughmi and Xu (2017) to more general settings and prove that, when the sender’s utility function is anonymous, computing an optimal signaling scheme is fixed-parameter tractable in the number of receivers’ actions. In the public signaling case, we show that the sender’s optimal expected return cannot be approximated to within *any* factor under a  $k$ -voting rule. This negative result easily extends to plurality voting and problems where utility functions are anonymous.

## Introduction

Information is the foundation of any democratic election, as it allows voters for better choices. In many settings, uninformed voters have to rely on inquiries of third party entities to make their decision. For example, in most trials, jurors are not given the possibility of choosing which tests to perform during the investigation or which questions are asked to witnesses. They have to rely on the prosecutor’s investigation and questions. The same happens in elections, in which voters gather information from third-party sources. With the advent of modern media environments, malicious actors have unprecedented opportunities to garble this information and influence the outcome of the election via misinformation

and fake news (Allcott and Gentzkow 2017). Reaching voters with targeted messages has never been easier. Hence, we pose the question: *can a malicious actor influence the outcome of a voting process only by the provision of information to voters who update their beliefs rationally?*

We describe the problem through the *Bayesian persuasion* framework by Kamenica and Gentzkow (2011). At its core, the model involves an informed *sender* trying to influence the behavior of self-interested *receivers* through the provision of payoff-relevant information. Kamenica and Gentzkow (2011) study how a single sender, having access to some private information, can design a signaling scheme to persuade a single receiver to select a favorable action. The model assumes the sender’s commitment, which is realistic in many settings (Kamenica and Gentzkow 2011; Dughmi 2017). One argument to that effect is that reputation and credibility may be a key factor for the long-term utility of the sender (Rayo and Segal 2010).

A number of recent works study *social influence* as a means of election control (Sina et al. 2015; Faliszewski et al. 2018; Wilder and Vorobeychik 2018; 2019). The crucial difference in our model is that voters are strategic players, who update their beliefs rationally. This property forces the sender to carefully craft the signaling scheme to preserve persuasiveness (*i.e.*, incentive compatibility). Other mechanisms for election interference that have been studied are *bribery* (Faliszewski et al. 2009; Erdélyi, Reger, and Yang 2017), and *adding/deleting voters/candidates* (Loreggia et al. 2015; Faliszewski, Hemaspaandra, and Hemaspaandra 2011; Liu et al. 2009; Chen et al. 2017). Both these mechanisms differ from ours in the provision of tangible incentives (the former) or the modification of the election setting (the latter). In our model, the malicious actor (*i.e.*, the sender), can influence the outcome of the election only by deciding *who gets to know what*.

We extend the fundamental model of Arieli and Babichenko (2019) to describe general voting problems. Specifically, our model comprises of multiple receivers, an arbitrary number of actions and states of nature, and no inter-agent externalities.<sup>1</sup> We adopt the sender’s perspective, and

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>Without inter-agent externalities, the utility of a voter depends only on her appreciation of the candidate she selected, and on the

study how to compute optimal *private* (i.e., different voters may receive different information) and *public* (i.e., all voters observe the same signal) communication schemes. We observe a sharp contrast between the two settings: all results for the private setting are positive (i.e., polynomial time tractability), while coordinating voters via public signaling is largely intractable.

**Original contributions.** We focus on two commonly adopted voting rules: *k*-voting rules and *plurality voting*. First, we provide an efficient implementation of the optimal private signaling scheme under a *k*-voting rule. Then, by generalizing a result by Dughmi and Xu (2017)—later revisited by Xu (2019)—we describe a necessary and sufficient condition for the efficient computation of private signaling schemes for a general class of sender’s objective functions. This condition is employed to show that private Bayesian persuasion is fixed-parameter tractable in the number of receivers’ actions when sender’s utility is anonymous, and to show that an optimal private signaling scheme under plurality voting can be found in polynomial time. As for public signaling, we provide a new inapproximability result on the problem of computing an optimal public signaling scheme under a *k*-voting rule, showing that it cannot be approximated to within *any* factor of the input size. This result extends to anonymous sender’s utility functions and plurality voting.

## Related Works

The classic model of Bayesian persuasion is due to Kamenica and Gentzkow (2011). Later, (Bergemann and Morris 2016a; 2016b; 2019) highlighted the connection between optimal information disclosure and the *best* Bayes correlated equilibrium from the sender’s perspective. A number of works deal with the multiple receivers generalization of the model, e.g., (Schnakenberg 2015; Taneva 2015; Wang 2013). Among these works, those by Bardhi and Guo (2018), Alonso and Câmara (2016) and Chan et al. (2019) are closely related to ours, representing the first attempts of applying the Bayesian persuasion framework to voting problems. In particular, Bardhi and Guo (2018) and Chan et al. (2019) focus on problems with binary actions and state spaces. The former studies private Bayesian persuasion for unanimity voting, while the latter analyzes private and public persuasion with *k*-majority voting rules. Alonso and Câmara (2016) employ a novel geometric tool to characterize an optimal public signaling scheme in a voting framework and characterize voters’ preferences over electoral rules. However, the aforementioned works only provide the economic groundings of Bayesian persuasion in (simple) voting settings. Their characterization does not include any computational result.

Dughmi and Xu (2016) are the first to analyze Bayesian persuasion from a computational perspective, in the single-receiver case. Arieli and Babichenko (2019) introduce the fundamental model of persuasion with no inter-agent externalities.<sup>2</sup> In the case of binary actions and state spaces, Arieli

state of nature. It does not depend on actions taken by other voters.

<sup>2</sup>Problems with externalities are largely intractable even in very

and Babichenko (2019) provide an explicit characterization of the optimal private signaling scheme when the sender’s utility function is either supermodular, anonymous submodular, or supermajority. Moreover, Arieli and Babichenko (2019) also provide necessary and sufficient conditions for the existence of a public signaling scheme matching the performance of the best private signaling scheme. Other related works focusing on the no externality setting with receivers’ binary action spaces are the following: Babichenko and Barman (2016) describe a tight  $(1 - 1/e)$ -approximate private signaling scheme for monotone submodular sender’s utility functions and show that an optimal private scheme for anonymous utility functions can be found in polynomial time. Dughmi and Xu (2017) generalize the model to the case of many states of nature, while assuming sender’s utility to be a monotone set function. In this setting, the supermodular and anonymous cases can be efficiently solved. When the sender’s utility is submodular, Dughmi and Xu (2017) show that an  $(1 - 1/e)$ -approximation to the optimal revenue can be obtained by sending conditionally independent private signals to receivers. Moreover, the authors show that it is NP-hard to approximate the sender’s value provided by the optimal public scheme, within any constant factor. Finally, Xu (2019) focuses on the complexity of public signaling when there are no inter-agent externalities and the action spaces are binary. Finding an optimal public signal is shown to be *fixed-parameter tractable* under some non-degeneracy assumptions. The author describes a PTAS with a bi-criteria guarantee for (monotone) submodular sender’s objectives.

## Model

Our model is a generalization of the fundamental special case introduced by Arieli and Babichenko (2019). It comprises a *sender* and a finite set  $R$  of *receivers* (voters) that must choose one alternative from a set  $C = \{c_0, \dots, c_\ell\}$  of candidates (i.e.,  $C$  is the set of voters’s available actions). Each voter must choose a candidate from  $C$ . Each voter’s utility depends only on her own action and the (random) state of nature, but not on the actions of other voters. In particular, we write  $u_r : C \times \Theta \rightarrow \mathbb{R}$ , where  $\Theta = \{\theta_i\}_{i=1}^n$  is the finite space of states of nature. The value of  $u_r(c, \theta)$  is a measure of voter  $r$ ’s appreciation of candidate  $c$  when the state of nature is  $\theta$ . A profile of votes (i.e., one candidate for each voter) is denoted by  $\mathbf{c} \in \times_{r \in R} C$ . In general settings, beyond voting, we denote the sender’s utility when the state of nature is  $\theta$  with  $f_\theta : \times_{r \in R} C \rightarrow \mathbb{R}$  (here  $C$  may be an arbitrary space of actions). Furthermore, we say that  $f$  is *anonymous* if its value depends only on  $\theta$  and on the number of players selecting each action. In the specific context of voting, the sender’s objective is to maximize the winning probability of  $c_0$  (according to some voting rules). In this setting, instead of using  $f$ , we denote the sender’s utility function by  $W : \times_{r \in R} C \rightarrow \{0, 1\}$ , where  $W(\cdot) = 1$

simple settings; see, e.g., (Bhaskar et al. 2016; Rubinstein 2015; Dughmi 2018). The no-externality assumption removes the equilibrium selection and (partially) the computation concerns arising in more general setting, and allows one to focus on the problem of coordinating receivers’ actions.




Figure 1: Interaction between the sender and a receiver.

if  $c_0$  wins, and  $W(\cdot) = 0$  otherwise. The state of nature influences the receivers' preferences but it does not affect the sender's payoff, which only depends on the final votes.<sup>3</sup>

As it is customary in Bayesian persuasion (see, e.g., Kamenica (2018)), we assume  $\theta$  is drawn from a common prior  $\mu \in \text{int}(\Delta(\Theta))$ , which is explicitly known to both sender and receivers.<sup>4</sup> Their interaction goes as follows (see Figure 1): the sender commits to a publicly known *signaling scheme*  $\phi$  that maps states of nature to *signals* for the voters. The *signal set* of receiver  $r$  is denoted by  $S_r$ , while  $s_r \in S_r$  is a single signal for  $r$ . A profile of signals is denoted by  $\mathbf{s} = (s_1, \dots, s_{|R|}) \in S$ , where  $S = \times_{r \in R} S_r$ . The sender observes the realized state of nature  $\theta \sim \mu$ , and exploits the knowledge of  $\theta$  to compute  $\mathbf{s} \in S$ , drawn according to  $\phi$  under  $\theta$ . Each voter  $r$  observes  $s_r$  and rationally updates her posterior over  $\Theta$ . Then, each voter  $r$  selects a candidate  $c_r^*$  maximizing  $r$ 's expected reward.

A general signaling scheme  $\phi : \Theta \rightarrow \Delta(S)$  has a private communication channel for each receiver. We refer to these as *private signaling schemes*. The probability with which the sender recommends  $\mathbf{s} \in S$  after observing  $\theta$  is  $\phi(\theta, \mathbf{s})$ . Therefore, it holds  $\sum_{\mathbf{s} \in S} \phi(\theta, \mathbf{s}) = 1$  for each  $\theta$ . Let  $S_{-r} = \times_{j \neq r} S_j$ . Then,  $\phi_r(\theta, s_r) = \sum_{\mathbf{s}_{-r} \in S_{-r}} \phi(\theta, (s_r, \mathbf{s}_{-r}))$  is the *marginal probability* with which  $s_r \in S_r$  is recommended to  $r$  when  $\theta$  is observed. Furthermore, we call  $\phi_r$  a *marginal signaling scheme*. It is immediate to observe that the same set  $\{\phi_r\}_{r \in R}$  may be obtained from more than one distribution over  $S$ . A special setting is the one of *public signaling schemes*, where the same signal is received by all receivers through a public communication channel (equivalently, all voters have to receive the same private message). With an overload of notation, we denote a public signaling scheme by  $\phi : \Theta \rightarrow S$ , where  $S$ , denotes the set of public signals (the meaning of  $\phi$  will be clear from the context).

A revelation-principle style argument shows that, in the case of no inter-agent externalities, there exists an optimal (private or public) signaling scheme which is both *direct* and *persuasive* (Kamenica and Gentzkow 2011; Dughmi and Xu 2017).<sup>5</sup> A signaling scheme is direct when signals can be mapped to actions of the receivers, and interpreted as action recommendations, i.e., in the voting setting,  $S_r = C$

<sup>3</sup>The sender's utility function is state-independent in many settings, e.g., voting (Alonso and Câmara 2016), and marketing (Candogan 2019; Babichenko and Barman 2016).

<sup>4</sup> $\text{int}(X)$  is the *interior* of set  $X$ , and  $\Delta(X)$  is the set of all probability distributions on  $X$ .

<sup>5</sup>An optimal signaling scheme always exists when the sender's expected utility is upper semicontinuous, which is always the case when receivers break ties in favor of the sender (Kamenica and Gentzkow 2011).

for each  $r \in R$ . We say that a signaling scheme is *persuasive* if following recommendations is an equilibrium of the underlying *Bayesian game* (Bergemann and Morris 2016a; 2016b). Analogously, a *marginal signaling scheme*  $\phi_r$  is persuasive if  $r$  has no incentives in deviating from its recommendations. When not specified, we assume to work with direct signaling schemes. Moreover, when the sender's signaling scheme  $\phi$  is direct and persuasive we write  $W(\phi)$  to denote the sender's expected utility. Finally, function  $\delta : S \times C \rightarrow \mathbb{N}$  is s.t.  $\delta(\mathbf{s}, c)$  is the number of voters that are recommended  $c$  by  $\mathbf{s}$ .

To further clarify the notion of Bayesian persuasion in voting problems, we provide the following simple example.

**Example 1.** *There are three voters  $R = \{1, 2, 3\}$  that have to choose one among two candidates  $C = \{c_0, c_1\}$ . The sender (e.g., a politician or a lobbyist) is interested in having  $c_0$  elected, and can observe the realized state of nature, drawn uniformly from  $\Theta = \{A, B, C\}$ . The state of nature describes the position of  $c_0$  on a matter of particular interest to the voters. Moreover, voters perceive  $c_1$  in a slightly negative way, independently of the state of nature. Table 1 describes the utility of the three voters.*

		State A	State B	State C	
		$c_0$	$c_1$	$c_0$	$c_1$
Voters	1	+1	-1/4	-1	-1/4
	2	-1	-1/4	+1	-1/4
	3	-1	-1/4	-1	-1/4

Table 1: Payoffs from voting different candidates.

		not A	Signals not B	not C
States	A	0	1/2	1/2
	B	1/2	0	1/2
	C	1/2	1/2	0

Table 2: Optimal signaling scheme.

*We consider a simple-majority voting rule. Without any form of signaling (or with a completely uninformative signal), all voters would choose  $c_1$  because it provides an expected utility of  $-1/4$ , against  $-1/3$ . This would clearly be the worst possible outcome for the sender, i.e.,  $W((c_1, c_1, c_1)) = 0$ . The sender would still get 0 utility with a fully informative (private) signal, since two out of three receivers would pick  $c_1$ . However, the sender can design a public signaling scheme granting herself utility 1 for each state of nature.<sup>6</sup> An optimal signaling scheme should convince two out of the three receivers to vote for  $c_0$ . Table 2 describes one such scheme with generic signals. Suppose the observed state is A, and that the signal is not B (sampled uniformly from {not B, not C}). Then, the posterior distribution over the states of nature is  $(1/2, 0, 1/2)$ . Therefore, both player 1 and 3 would vote for  $c_0$ , since their*

<sup>6</sup>Since a public signaling scheme is a special case of private signaling, the same result could be achieved in the latter scenario.

expected utility would be 0 against  $-1/4$ . The sender's payoff would be  $W((c_0, c_1, c_0)) = 1$ , and the same happens for any  $\theta \in \Theta$ . An equivalent direct signaling scheme would publicly reveal a tuple of candidates' suggestions. For example, not **A** would become  $(c_1, c_0, c_0)$ .

We consider two commonly adopted voting rules: *k*-voting rule and plurality voting rule (see, e.g., Brandt et al. (2016)). In an election with a *k*-voting rule each voter chooses a candidate after observing the sender's signal. Candidate  $c_i$  is elected if it receives at least  $k$  votes, where  $k \in [|R|]$  is the established electoral rule.<sup>7</sup> The problem of designing the optimal sender's persuasive signaling scheme under a *k*-voting rule is denoted by K-V. In an election with plurality voting rule the winner is determined as the candidate with a plurality (greatest number) of votes. The problem of finding an optimal persuasive signaling scheme for the sender with plurality voting is denoted by PL-V. In both settings, we focus on maximizing the winning probability of the sender. The problem can be written as the optimization problem:  $\max_{\phi} \sum_{\theta \in \Theta, \mathbf{s} \in S} \mu(\theta) \phi(\theta, \mathbf{s}) W(\mathbf{s})$ , subject to  $\phi$  being persuasive for each voter. An approximate solution for this problem is a signaling scheme that is persuasive, but guarantees to the sender a sub-optimal expected utility.

### Private Signaling with *k*-Voting Rules

In this section, we show that a solution to K-V (i.e., finding an optimal persuasive signaling scheme under a *k*-voting rule) can be found in polynomial time when the sender can employ a private signaling scheme.

First, we show that the sender can restrict the choice of a signaling scheme to the set of the schemes  $\phi$  whose marginal signaling schemes are Pareto efficient on the set  $\{\phi_r(\theta, c_0)\}_{\theta \in \Theta, r \in R}$  (Lemma 1), and recommend with positive probability either  $c_0$  or the candidate giving  $r$  the best utility under  $\theta$  (Lemma 2).

**Lemma 1.** *Given a signaling scheme  $\phi'$  and a set of persuasive marginal signaling schemes  $\{\phi_r\}_{r \in R}$ , if  $\phi_r(\theta, c_0) \geq \phi'_r(\theta, c_0)$  for each  $r \in R$  and  $\theta \in \Theta$ , there exists a persuasive signaling scheme  $\phi$  such that  $W(\phi) \geq W(\phi')$ .*

*Proof.* Be given a signaling scheme  $\phi'$  and a set of persuasive marginal signaling schemes  $\{\phi_r\}_{r \in R}$  s.t.  $\phi_r(\theta, c_0) \geq \phi'_r(\theta, c_0)$  for each  $r \in R$ ,  $\theta \in \Theta$ . Intuitively, we show that it is possible to move probability mass to  $c_0$  while guaranteeing persuasiveness with the following procedure. Let  $\phi^0 = \phi'$ , and  $R = \{r_1, \dots, r_{|R|}\}$ . We iterate over  $j \in [|R|]$ , and update the signaling scheme with the following procedure. Let  $A_{r_j}$  be an arbitrary mapping from  $[|S_{-r_j}|]$  to  $S_{-r_j}$ , which serves as an arbitrary ordering of elements in  $S_{-r_j}$  (i.e.,  $A_{r_j}(i)$  returns the  $i$ -th element of  $S_{-r_j}$  in such ordering). Moreover,  $\forall \theta \in \Theta$ , we define  $\Delta_{r_j}^0(\theta) = \phi_{r_j}(\theta, c_0) - \phi'_{r_j}(\theta, c_0)$ . For each  $j \in [|R|]$ , we iterate over  $i \in [|S_{-r_j}|]$ , and perform the following updates:  $\mathbf{s}_{-r_j} = A_{r_j}(i)$ ,

<sup>7</sup>We denote by  $[n]$  the set  $\{1, \dots, n\}$ .

$$\phi^j(\theta, (c_0, \mathbf{s}_{-r_j})) = \min \left\{ \phi^{j-1}(\theta, (c_0, \mathbf{s}_{-r_j})) + \Delta_{r_j}^{i-1}(\theta), \sum_{c \in C} \phi^{j-1}(\theta, (c, \mathbf{s}_{-r_j})) \right\}, \quad (1)$$

and

$$\Delta_{r_j}^i(\theta) = \Delta_{r_j}^{i-1}(\theta) - \phi^j(\theta, (c_0, \mathbf{s}_{-r_j})) + \phi^{j-1}(\theta, (c_0, \mathbf{s}_{-r_j})),$$

where  $\phi^j(\theta, (c, \mathbf{s}_{-r_j}))$  is the probability of recommending  $c$  to  $r_j$  and  $\mathbf{s}_{-r_j}$  to the other receivers, under  $\theta$  (at iteration  $j$ ). Finally, for each  $\mathbf{s}_{-r_j}$ , and  $c \neq c_0$ , set:

$$\phi_{r_j}^j(\theta, c) = \frac{\phi_{r_j}^j(\theta, c) \left( \sum_{c' \in C} \phi^{j-1}(\theta, (c', \mathbf{s}_{-r_j})) - \phi^j(\theta, (c_0, \mathbf{s}_{-r_j})) \right)}{\sum_{c' \in C \setminus \{c_0\}} \phi_{r_j}(\theta, c')},$$

the numerator is well-defined because of the minimization in Equation 1. After having enumerated all the receivers, we obtain  $\phi^{|R|}$ . We show that  $\phi = \phi^{|R|}$  is precisely the desired signaling scheme. First, we show that, at each iteration  $r$ ,  $\phi^r$  is well formed. For each iteration  $j$ , and pair  $(\theta, \mathbf{s}_{-r_j})$ , we show that  $\sum_{c \in C} \phi^j(\theta, (c, \mathbf{s}_{-r_j})) = \sum_{c \in C} \phi^{j-1}(\theta, (c, \mathbf{s}_{-r_j}))$ . We have:  $\sum_{c \in C} \phi^j(\theta, (c, \mathbf{s}_{-r_j})) = \phi^j(\theta, (c_0, \mathbf{s}_{-r_j})) + \sum_{c \in C \setminus \{c_0\}} \phi^j(\theta, (c, \mathbf{s}_{-r_j}))$ . Then, by expanding  $\phi^j(\theta, (c, \mathbf{s}_{-r_j}))$  via the update rule, we obtain:

$$\sum_{c \in C} \phi^j(\theta, (c, \mathbf{s}_{-r_j})) = \phi^j(\theta, (c_0, \mathbf{s}_{-r_j})) + \sum_{c \in C} \phi^{j-1}(\theta, (c, \mathbf{s}_{-r_j})) - \phi^j(\theta, (c_0, \mathbf{s}_{-r_j})),$$

which is precisely  $\sum_{c \in C} \phi^{j-1}(\theta, (c, \mathbf{s}_{-r_j}))$ . This implies that  $\sum_{\mathbf{s} \in S} \phi^j(\theta, \mathbf{s}) = 1$ , and that receiver  $r_j$ 's marginal probabilities are modified only at iteration  $r$ . Now, we show that receiver  $r_j$ 's marginals are updated correctly. We distinguish the following two cases.

i) It is easy to see that, for candidate  $c_0$ ,

$$\begin{aligned} \sum_{\mathbf{s}_{-r_j} \in S_{-r_j}} \phi^j(\theta, (c_0, \mathbf{s}_{-r_j})) &= \\ &= \Delta_r^0(\theta) + \sum_{\mathbf{s}_{-r_j} \in S_{-r_j}} \phi^{j-1}(\theta, (c_0, \mathbf{s}_{-r_j})) = \phi_{r_j}(\theta, c_0). \end{aligned}$$

ii) For each  $c \neq c_0$ , we have:  $\sum_{\mathbf{s}_{-r_j} \in S_{-r_j}} \phi^j(\theta, (c, \mathbf{s}_{-r_j})) =$

$$\begin{aligned} &= \sum_{\mathbf{s}_{-r_j} \in S_{-r_j}} \frac{\phi_{r_j}(\theta, c) \left( \sum_{c' \in C} \phi^{j-1}(\theta, (c', \mathbf{s}_{-r_j})) - \phi^j(\theta, (c_0, \mathbf{s}_{-r_j})) \right)}{\sum_{c' \in C \setminus \{c_0\}} \phi_{r_j}(\theta, c')} \\ &= \frac{\phi_{r_j}(\theta, c) \left( \sum_{\mathbf{s} \in S} \phi^{j-1}(\theta, \mathbf{s}) - \sum_{\mathbf{s}_{-r_j} \in S_{-r_j}} \phi^j(\theta, (c_0, \mathbf{s}_{-r_j})) \right)}{\sum_{c' \in C \setminus \{c_0\}} \phi_{r_j}(\theta, c')} = \\ &= \frac{\phi_{r_j}(\theta, c)(1 - \phi_{r_j}(\theta, c_0))}{\sum_{c' \in C \setminus \{c_0\}} \phi_{r_j}(\theta, c')} = \phi_{r_j}(\theta, c). \end{aligned}$$

Since  $\{\phi_r\}_{r \in R}$  are persuasive, also the new signaling scheme  $\phi$  is persuasive. Finally, we show that the new signaling scheme does not decrease sender's expected utility. Let  $S^* = \{s \in S \mid \delta(s, c_0) \geq k\}$  be the set of joint signals recommending  $c_0$  to more than  $k$  voters (under a  $k$ -voting rule). Then,  $W(\phi) = \sum_{\theta \in \Theta} \mu(\theta) \sum_{s \in S^*} \phi(\theta, s)$ . It is enough to show that, for each iteration  $j$ , for each  $\theta \in \Theta$ , and, for each  $s_{-r_j} \in S_{-r_j}$ , it holds

$$\sum_{c \in C} (\phi^j(\theta, (c, s_{-r_j})) - \phi^{j-1}(\theta, (c, s_{-r_j}))) \mathbb{1}_{(c, s_{-r_j}) \in S^*} \geq 0.$$

We distinguish three cases. i) When  $\delta(s_{-r_j}, c_0) < k - 1$ , a change in  $r_j$ 's marginal probabilities does not affect the sender's winning probability, term  $\mathbb{1}_{(c, s_{-r_j}) \in S^*}$  being always 0. ii) When  $\delta(s_{-r_j}, c_0) = k - 1$ ,  $\mathbb{1}_{(c, s_{-r_j}) \in S^*} = 1$  only if  $c = c_0$ , and  $\phi^j(\theta, (c_0, s_{-r_j})) \geq \phi^{j-1}(\theta, (c_0, s_{-r_j}))$ . iii) When  $\delta(s_{-r_j}, c_0) > k - 1$ ,  $\mathbb{1}_{(c, s_{-r_j}) \in S^*}$  is always 1, and we already know that  $\sum_{c \in C} (\phi^j(\theta, (c, s_{-r_j})) - \phi^{j-1}(\theta, (c, s_{-r_j}))) = 0$ .  $\square$

**Lemma 2.** *There always exists a solution  $\phi$  to K-V in which, for all  $r \in R$  and  $\theta \in \Theta$ ,  $\phi_r(\theta, c) > 0$  if and only if one of the following two conditions is satisfied: i)  $c = c_0$  ii)  $c \in \arg \max_{c' \in C} u_r(\theta, c')$ .*

*Proof.* Given a persuasive signaling scheme  $\phi'$ , we show that it is possible to build a collection  $\{\phi_r\}_{r \in R}$  with the property above, s.t.  $\phi_r(\theta, c_0) \geq \phi'_r(\theta, c_0)$  for each  $r \in R$ ,  $\theta \in \Theta$ . This, together with Lemma 1, proves our result. We build  $\phi$  iteratively. For each pair  $(\theta, r)$ , select  $c^* \in \arg \max_{c \in C} u_r(\theta, c)$ , and set  $\phi_r(\theta, c^*) = 1 - \phi'_r(\theta, c_0)$ ,  $\phi_r(\theta, c_0) = \phi'_r(\theta, c_0)$ , and  $\phi_r(\theta, c) = 0$  for each other  $c \in C \setminus \{c_0, c^*\}$ . It is immediate to see that, for each  $\theta$  and  $r$ ,  $\sum_{c \in C} \phi_r(\theta, c) = 1$ . Next, we show that each  $\phi_r$  is persuasive, i.e.,  $\sum_{\theta \in \Theta} \mu(\theta) \phi_r(\theta, c) (u_r(\theta, c) - u_r(\theta, c')) \geq 0$  for each  $r \in R$ , and  $c, c' \in C$ . If  $c = c_0$ , we have  $\phi_r(\theta, c_0) > \phi'_r(\theta, c_0)$  only if  $c_0 \in \arg \max_{c \in C} u_r(\theta, c)$ , which means  $(u_r(\theta, c_0) - u_r(\theta, c')) \geq 0$ , in the remaining cases we have  $\phi_r(\theta, c_0) = \phi'_r(\theta, c_0)$ . If  $c \neq c_0$ ,  $c \in \arg \max_{c' \in C} u_r(\theta, c')$   $\forall \theta \in \Theta$  with  $\phi(\theta, c) > 0$ , which makes the incentive constraint satisfied.  $\square$

By exploiting Lemma 2, we show that an optimal persuasive signaling scheme under a  $k$ -voting rule can be computed in polynomial time via the following linear program (LP). Let  $\beta_\theta \in \mathbb{R}$  be the probability with which  $k$  voters vote for  $c_0$  with state  $\theta$ . Then, we can compute an optimal solution to K-V as follows (the proof is provided below):

$$\max_{\beta \in [0, 1]^{|\Theta|}, z \in \mathbb{R}^{|\Theta| \times k \times |R|}} \sum_{\theta \in \Theta} \mu(\theta) \beta_\theta \quad (2a)$$

$$\begin{aligned} & t, q \in \mathbb{R}^{|\Theta| \times k} \\ & \phi_r(\cdot, c_0) \in [0, 1]^{|R \times \Theta|} \\ & \text{s.t. } \sum_{\theta \in \Theta} \mu(\theta) \phi_r(\theta, c_0) (u_r(\theta, c_0) - u_r(\theta, c)) \geq 0 \end{aligned} \quad (2b)$$

$$\begin{aligned} & \forall r \in R, \forall c \in C \setminus \{c_0\} \\ & \beta_\theta \leq \frac{1}{k - m} q_{\theta, m} \forall \theta \in \Theta, \forall m \in \{0, \dots, k - 1\} \end{aligned} \quad (2c)$$

$$q_{\theta, m} \leq (|R| - m) t_{\theta, m} + \sum_{r \in R} z_{\theta, r, m} \quad (2d)$$

$$\begin{aligned} & \forall \theta \in \Theta, \forall m \in \{0, \dots, k - 1\} \\ & \phi_r(\theta, c_0) \geq t_{\theta, m} + z_{\theta, m, r} \\ & \forall r \in R, \forall \theta \in \Theta, \forall m \in \{0, \dots, k - 1\}. \end{aligned} \quad (2e)$$

This formulation allows us to state the following:

**Theorem 3.** *It is possible to compute an optimal persuasive private signaling scheme for K-V in poly( $n, \ell, |R|$ ) time.*

*Proof.* Formulation 2 has a polynomial number of variables and constraints. Then, proving Theorem 3 amounts to show that a solution to Formulation 2 is also a solution to K-V.

Let  $c_{\theta, r}^* = \arg \max_{c \in C} u_r(\theta, c)$ , for each  $\theta$  and  $r$ . First, by Lemma 2, the space of available signals can be restricted to those in which, for each  $r$  and  $\theta$ , only  $\phi_r(\theta, c_0)$  and  $\phi_r(\theta, c_{\theta, r}^*)$  are  $> 0$ , and  $\phi_r(\theta, c_{\theta, r}^*) = 1 - \phi_r(\theta, c_0)$ . Constraints (2b) are the incentive constraints for action  $c_0$ . For any  $c \neq c_0$ , the incentive constraints are satisfied by construction. Objective (2a) is given by the sum over all  $\theta \in \Theta$  of the prior of state  $\theta$ , multiplied by the probability of having at least  $k$  vote for  $c_0$  given  $\theta$ . We need to show the correctness of  $\beta_\theta$ . For each state of nature the maximum probability with which at least  $k$  receivers play  $c_0$  is given by:

$$\beta_\theta = \min \left\{ \min_{m \in \{0, \dots, k-1\}} \frac{1}{k - m} q_{\theta, m}, 1 \right\},$$

where  $q_{\theta, m}$  is the sum of the lowest  $|R| - m$  elements in the set  $\{\phi_r(\theta, c_0)\}_{r \in R}$ ; for further details, see (Arieli and Babichenko 2019, Lemma 3). This is enforced via Constraints (2c). Constraints (2d) and (2e) ensure  $q_{\theta, m}$ 's consistency, and are derived from the dual of a simple LP of this kind:  $\min_{y \in \mathbb{R}^n} x^\top y$  s.t.  $\mathbb{1}^\top y = w$  and  $0 \leq y \leq 1$  (where  $x \in \mathbb{R}^n$  is the vector from which we want to extract the sum of the smallest  $w$  entries). This concludes the proof.  $\square$

## A Condition for Efficient Private Signaling

In the following, we provide a necessary and sufficient condition for the poly-time computation of persuasive private signaling schemes under a general class of sender's objective functions. In the next section, this result will be exploited when dealing with anonymous utility functions and plurality voting. We allow for general sender's utility functions of type  $f_\theta : \times_{r \in R} C \rightarrow \mathbb{R}$ , which generalizes previous results by Dughmi and Xu (2017) where the receivers' action space has to be binary. Given a collection of set functions  $\mathcal{F}$ ,  $P(\mathcal{F})$  denotes the class of persuasion instances in which, for each  $\theta \in \Theta$ ,  $f_\theta \in \mathcal{F}$ . We can state the following (the proof can be found in the extended version of the paper (Castiglioni, Celli, and Gatti 2019)).

**Theorem 4.** *Let  $\mathcal{F}$  be any collection of set functions including  $f_0(\cdot) = 0$ . Given any instance in  $P(\mathcal{F})$ , there exists a polynomial-time algorithm for computing an optimal persuasive private signaling scheme if and only if there is a polynomial-time algorithm that computes*

$$\max_{c \in \times_{r \in R} C} f(c) + \sum_{r \in R} w_r(c^r), \quad (3)$$

for any  $f \in \mathcal{F}$ , and any weights  $w_r(c^r) \in \mathbb{R}$ , where  $c^r$  is the action chosen by  $r$  in  $c$ .

The crucial difference with the result by Dughmi and Xu (2017) is that they consider set functions depending only on the set of players choosing the target action, among the two available. Theorem 4 generalizes this setting as it allows for functions taking as input any action profile  $c$ . This is crucial in settings like plurality voting, where the sender is not only interested in votes favorable to  $c_0$ , but also in the distribution of the other preferences. Dughmi and Xu (2017)'s result cannot be applied to such settings.

### Further Positive Results for Private Signaling

Despite Theorem 4, in the case of general utility functions the problem of determining an optimal persuasive private signaling scheme is still largely intractable. An intuition behind that is that there may be an exponential (in  $|C|$ ) number of values of  $f$  (e.g., in the case of anonymous utility functions, there are  $\binom{|R|+|C|-1}{|R|}$  values of  $f$ ). In order to identify tractable classes of the problem, we need to make some further assumptions on  $\mathcal{F}$ .

**Anonymous Utility Functions.** A reasonable (in the context of voting) restriction is to *anonymous utility functions* (see, e.g., Arieli and Babichenko (2019)). Previous results on the computational complexity of private signaling with anonymous utility functions focus on the case of binary actions, which is shown to be tractable (Babichenko and Barman 2016; Arieli and Babichenko 2019; Dughmi and Xu 2017). We generalize these results to a generic number of states of nature and receiver's actions with the following result (the proof is provided in Castiglioni, Celli, and Gatti (2019)).

**Theorem 5.** *Private Bayesian persuasion with anonymous sender's utility functions is fixed-parameter tractable in the number of receivers' actions.*

Theorem 5 implies that, for any anonymous voting rule, private Bayesian persuasion is fixed-parameter tractable in the number of candidates.

**Plurality Voting.** By further restricting our attention to specific voting rules, we can see the consequences of Theorem 4 to an even better extent. A simple and widespread voting rule is *plurality voting*.<sup>8</sup> In this setting  $W(\mathbf{s}) = 1$  if and only if  $\delta(\mathbf{s}, c_0) > \delta(\mathbf{s}, c)$  for any  $c \neq c_0$ , and  $W(\mathbf{s}) = 0$  otherwise. We can state the following:

**Theorem 6.** *PL-V with private signaling can be solved in  $\text{poly}(n, \ell, |R|)$  time.*

*Proof.* We exploit Theorem 4, and show that the maximization Problem (3) can be solved efficiently. With an overload of notation, generic actions profiles are represented via signals. Then, the maximization problem reads:  $\max_{\mathbf{s} \in S} W(\mathbf{s}) + \sum_{r \in R} w_r(s_r)$ . We split the maximization problem in two steps. First, we consider the maximization over non-winning action profiles, i.e., signals in the set  $\bar{S} = \{\mathbf{s} \in S \mid \exists c \neq c_0 \text{ s.t. } \delta(\mathbf{s}, c) > \delta(\mathbf{s}, c_0)\}$ . An upper bound to the optimal value of the maximization problem

<sup>8</sup>See, e.g., its (discussed) adoption in direct presidential elections in a number of states (Blais, Massicotte, and Dobrzynska 1997).

restricted to  $\bar{S}$  is given by  $\max_{\mathbf{s} \in S} \sum_r w_r(s_r)$ . The latter problem can be solved independently for each receiver  $r$ , by choosing  $c$  maximizing  $w_r(c)$ . Once the relaxed problem has been solved, the objective function of the separation problem is adjusted by checking whether  $c_0$  is winning or not. The resulting value is then compared with the value from the following step.

We consider the maximization over winning action profiles, i.e., signals in  $S^* = S \setminus \bar{S}$ . For any  $\mathbf{s} \in S^*$ ,  $W(\mathbf{s}) = 1$ . Then, we have to maximize the same objective of the previous case with the following additional constraints:  $\delta(\mathbf{s}, c_0) > \delta(\mathbf{s}, c)$ , for all  $c \neq c_0$ . To determine an optimal solution to this problem, we enumerate over  $k \in \{\lceil \frac{|R|-1}{\ell} \rceil + 1, \dots, |R|\}$ , i.e., the number of votes that make  $c_0$  a potential winner of the election. Then, for each value of  $k$ , we consider action profiles such that  $\delta(\mathbf{s}, c_0) = k$ , and  $\delta(\mathbf{s}, c) < k$ , for all  $c \neq c_0$  (i.e., winning signals where  $c_0$  receives exactly  $k$  votes). An optimal solution for a fixed  $k$  can be determined with this LP:

$$\begin{aligned} \max_{\chi \in \mathbb{R}_+^{|R \times C|}} \quad & \sum_{(r,c) \in R \times C} \chi_r(c) w_r(c) \\ \text{s.t.} \quad & \sum_{r \in R} \chi_r(c_0) = k \\ & \sum_{r \in R} \chi_r(c) \leq k - 1 \forall c \in C \setminus \{c_0\} \\ & \sum_{c \in C} \chi_r(c) = 1 \forall r \in R. \end{aligned}$$

We look for an integer solution of the problem, which always exists and can be found in polynomial time (see, e.g., (Orlin 1997)). This is because the formulation is an instance of the *maximum cost flow problem*, which is, in its turn, a variation of the *minimum cost flow problem*. Once an integer solution has been found, an optimal action profile of the original maximization problem is the one obtained by recommending to each  $r$  the candidate  $c$  s.t.  $\chi_r(c) = 1$ .  $\square$

### Public Signaling

In contrast with the results for private signaling problems, we show that public persuasion in the context of voting is largely intractable.

We reduce from MAXIMUM  $k$ -SUBSET INTERSECTION (MSI) (Clifford and Popa 2011).

**Definition 1 (MSI).** *An instance of MAXIMUM  $k$ -SUBSET INTERSECTION is a tuple  $(\mathcal{E}, A_1, \dots, A_m, k, q)$ , where  $\mathcal{E} = \{e_1, \dots, e_n\}$  is a finite set of elements, each  $A_i$ ,  $i \in [m]$ , is a subset of  $\mathcal{E}$ , and  $k, q$  are positive integers. It is a "yes"-instance if there exist exactly  $k$  sets  $A_{i_1}, \dots, A_{i_k}$  such that  $|\cap_{j \in [k]} A_{i_j}| \geq q$ , and a "no"-instance otherwise.*

MSI has been recently shown to be NP-hard (Xavier 2012; Elkind et al. 2015). Now, we prove the following negative result:

**Theorem 7.**  *$K$ -V with public signaling, even with two candidates, cannot be approximated in polynomial time to within any factor, unless  $P=NP$ .*

*Proof.* Given an instance of MSI, we build a public signaling problem with the following features.

**Mapping.** It has a voter  $r_i$  for each  $A_i$ ,  $i \in [m]$ , and  $m$  voters  $r_{e,j}$ ,  $j \in [m]$ , for each  $e \in \mathcal{E}$ . There is a state of nature  $\theta_e$  for each  $e \in \mathcal{E}$ , and  $\mu(\theta_e) = 1/n$  for each  $\theta_e$ . Finally,  $C = \{c_0, c_1\}$ . Receivers have the following utility functions: for each  $r_i$ ,  $i \in [m]$ ,

$$u_{r_i}(\theta_e, c) = \begin{cases} 1 & \text{if } e \in A_i, c = c_0 \\ -n^2 & \text{if } e \notin A_i, c = c_0 \\ 0 & \text{if } c = c_1 \end{cases},$$

for each  $r_{e,j}$ ,  $e \in \mathcal{E}$ , and  $j \in [m]$ ,

$$u_{r_{e,j}}(\theta_{e'}, c) = \begin{cases} 1 & \text{if } e = e', c = c_0 \\ -\frac{1}{q-1} & \text{if } e \neq e', c = c_0 \\ 0 & \text{if } c = c_1 \end{cases}.$$

The sender needs at least  $k + mq$  votes (for  $c_0$ ) in order to win the election (i.e., we are considering a  $(k + mq)$ -voting rule). We prove our theorem by showing that  $c_0$  has a strictly positive probability of winning the election if and only if the corresponding MSI instance is satisfiable.

**If.** Suppose there exists a set  $A^* = \{A_{i_1}, \dots, A_{i_k}\}$  satisfying the MSI instance, and let  $I = \bigcap_{j \in [k]} A_{i_j}$ . Define a signaling scheme  $\phi$  with two signals ( $\gamma_0$  and  $\gamma_1$ ) such that, for each  $e \in I$ ,  $\phi(\theta_e, \gamma_0) = 1$ , and, for each  $e \notin I$ ,  $\phi(\theta_e, \gamma_1) = 1$ , and it is equal to 0 otherwise. We show that such a signaling scheme guarantees a strictly positive winning probability for the sender. First, we show that, when the realized state of nature  $\theta_e$  is such that  $e \in I$  (i.e., the sender recommends  $\gamma_0$ ), at least  $k + mq$  receivers vote for  $c_0$ . Each receiver  $r_i$  such that  $A_i \in A^*$  will choose  $c_0$  when recommended  $\gamma_0$ . Specifically,  $\sum_{\theta_e} \frac{1}{n} \phi(\theta_e, \gamma_0) u_{r_i}(\theta_e, c_0) = \frac{q}{n}$ , while  $\sum_{\theta_e} \frac{1}{n} \phi(\theta_e, \gamma_0) u_{r_i}(\theta_e, c_1) = 0$ . Receivers  $r_{e,j}$  with  $e \in I$  will vote for  $c_0$  after observing  $\gamma_0$ . This is because, for each  $e \in I$  and  $j \in [m]$ ,  $r_{e,j}$  has expected utility  $\frac{1}{n} \phi(\theta_e, \gamma_0) - \sum_{\theta_{e'}: e' \neq e} \frac{1}{n} \frac{1}{q-1} \phi(\theta_{e'}, c_0) = 0$  for voting  $c_0$ , and expected utility 0 for voting  $c_1$ . Then, when the realized state of nature is  $\theta_e$  with  $e \in I$ , there are at least  $k + mq$  receivers voting for  $c_0$ . Therefore, the sender's winning probability is at least  $\frac{k}{n}$  (i.e., the probability of observing  $\theta_e$  with  $e \in I$  under a uniform prior).

**Only if.** Suppose, by contradiction, that MSI is not satisfiable, and that the sender's winning probability under the optimal signaling scheme is not null. This implies the existence of a signal  $\gamma_0$  such that, when recommended, a set of receivers  $R^*$  votes for  $c_0$ , and  $|R^*| \geq k + mq$ . Then, there exist at least  $q$  states  $\theta_e$  in which all voters  $r_{e,j}$ ,  $j \in [m]$ , vote for  $c_0$ . Each receiver  $r_{e,j}$ , having observed  $\gamma_0$ , votes for  $c_0$  only if  $\phi(\theta_e, \gamma_0) - \frac{1}{q-1} \sum_{\theta_{e'}: e' \neq e} \phi(\theta_{e'}, \gamma_0) \geq 0$ . This implies that  $\phi(\theta_e, \gamma_0) - \sum_{\theta_{e'}} \phi(\theta_{e'}, \gamma_0) + \phi(\theta_e, \gamma_0) \geq 0$  and  $\phi(\theta_e, \gamma_0) \geq \sum_{\theta_{e'} \in \Theta} \phi(\theta_{e'}, \gamma_0)/q$ . Then, there are exactly  $q$  states  $\theta_e$  in which  $\gamma_0$  is played with probability  $\sum_{\theta_{e'} \in \Theta} \phi(\theta_{e'}, \gamma_0)/q$ , while  $\gamma_0$  is never played in the remaining states. As a consequence,  $R^*$  includes exactly  $mq$  voters  $r_{e,j}$ , and at least  $k$  voters  $r_i$ .

Each voter  $r_i \in R^*$ , after observing  $\gamma_0$ , choose candidate  $c_0$ . Therefore,  $\sum_{\theta_e \in \Theta} \mu(\theta_e) \phi(\theta_e, \gamma_0) (u_{r_i}(\theta_e, c_0) -$

$u_{r_i}(\theta_e, c_1)) \geq 0$ . We obtain  $\sum_{e \in A_i} \phi(\theta_e, \gamma_0) - n^2 \sum_{e \notin A_i} \phi(\theta_e, \gamma_0) \geq 0$ . Then,

$$\sum_{e \in A_i} \phi(\theta_e, \gamma_0) - n^2 \sum_{e \in \mathcal{E}} \phi(\theta_e, \gamma_0) + n^2 \sum_{e \in A_i} \phi(\theta_e, \gamma_0) \geq 0.$$

Let  $\xi(A_i) = \sum_{e \in A_i} \phi(\theta_e, \gamma_0)$ . We have  $\xi(A_i) \geq \frac{n^2}{n^2+1} \sum_{e \in \mathcal{E}} \phi(\theta_e, \gamma_0)$  for each  $i \in [m]$  such that  $r_i \in S^*$ .

Let  $\mathcal{E}^*$  be the set of elements  $e$  such that  $r_{e,j} \in S^*$ , for all  $j \in [m]$ . In this case, since MSI is not satisfiable, there exists a pair  $(r_i, e) \in R \times \mathcal{E}^*$  such that  $r_i \in R^*$  and  $e \notin A_i$  (otherwise  $\{A_i\}_{i: r_i \in R^*}$  would be a feasible solution with intersection  $\mathcal{E}^*$ ). We observed that, in each  $\theta_e$  with  $e \in \mathcal{E}^*$ ,  $\gamma_0$  is recommended with probability  $\sum_{e \in \mathcal{E}} \phi(\theta_e, \gamma_0)/q$ . Then,  $\xi(A_i) = \sum_{e \in A_i} \phi(\theta_e, \gamma_0) \leq \frac{q-1}{q} \sum_{e \in \mathcal{E}} \phi(\theta_e, \gamma_0)$ . This leads to a contradiction since

$$\frac{q-1}{q} \sum_{e \in \mathcal{E}} \phi(\theta_e, \gamma_0) \geq \frac{n^2}{n^2+1} \sum_{e \in \mathcal{E}} \phi(\theta_e, \gamma_0)$$

has no solutions (since  $q$  and  $n$  are positive integers and  $q \leq n$ ). This concludes our proof.  $\square$

Theorem 7 implies that the public signaling problem is intractable also with more general sender's utility functions. It is immediate to see that the same negative result holds for anonymous utility functions (a  $k$ -voting rule induces a sender's anonymous utility function), and we prove that the same hardness result holds also for plurality voting (see Castiglioni, Celli, and Gatti (2019)).

**Corollary 1.** *PL-V with public signaling, even with two candidates, cannot be approximated in polynomial time to within any factor, unless  $P=NP$ .*

## Discussion and Future Research

This paper studies how a malicious actor may influence the outcome of a voting process by the strategic provision of information to voters that update their beliefs rationally. We focus on the case with no inter-agent externalities, and allow for an arbitrary number of candidates and states of nature, thus generalizing the model of Arieli and Babichenko (2019). We draw a sharp contrast between the tractability of the problem of computing an optimal persuasive signaling scheme in the private and public case, respectively. In the former setting, we show that under various voting rules (i.e.,  $k$ -voting rules and plurality voting) the problem can be solved efficiently. In doing so, we provide a generalization of the fundamental necessary and sufficient condition first described by Dughmi and Xu (2017). In the public signaling case, we propose a new inapproximability result which strongly improves previously known results, showing that in this setting the problem is unlikely to be tractable.

In the future, we are interested in identifying special classes of instances in which optimal persuasive public signaling schemes can be found efficiently and in using our results for the analysis of other voting rules. Moreover, we are interested in the extremely challenging case in which multiple competing senders face the problem of manipulating the same election.

## Acknowledgments

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

## References

- Allcott, H., and Gentzkow, M. 2017. Social media and fake news in the 2016 election. *J ECON PERSPECT* 31(2):211–36.
- Alonso, R., and Câmara, O. 2016. Persuading voters. *AM ECON REV* 106(11):3590–3605.
- Arieli, I., and Babichenko, Y. 2019. Private bayesian persuasion. *J ECON THEORY* 182:185–217.
- Babichenko, Y., and Barman, S. 2016. Computational aspects of private bayesian persuasion. *arXiv preprint arXiv:1603.01444*.
- Bardhi, A., and Guo, Y. 2018. Modes of persuasion toward unanimous consent. *THEOR ECON* 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):586–91.
- Bergemann, D., and Morris, S. 2019. Information design: A unified perspective. *J ECON LIT* 57(1):44–95.
- Bhaskar, U.; Cheng, Y.; Ko, Y. K.; and Swamy, C. 2016. Hardness results for signaling in bayesian zero-sum and network routing games. In *ACM EC*, 479–496.
- Blais, A.; Massicotte, L.; and Dobrzynska, A. 1997. Direct presidential elections: a world summary. *ELECT STUD* 16(4):441–455.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D. 2016. *Handbook of computational social choice*. Cambridge University Press.
- Candogan, O. 2019. Persuasion in networks: Public signals and k-cores. *Available at SSRN*.
- Castiglioni, M.; Celli, A.; and Gatti, N. 2019. Persuading voters: It’s easy to whisper, it’s hard to speak loud. *arXiv preprint arXiv:1908.10620*.
- 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.
- Clifford, R., and Popa, A. 2011. Maximum subset intersection. *Information Processing Letters* 111(7):323–325.
- Dughmi, S., and Xu, H. 2016. Algorithmic bayesian persuasion. In *ACM STOC*, 412–425.
- Dughmi, S., and Xu, H. 2017. Algorithmic persuasion with no externalities. In *ACM EC*, 351–368.
- Dughmi, S. 2017. Algorithmic information structure design: a survey. *ACM SIGEC EX* 15(2):2–24.
- Dughmi, S. 2018. On the hardness of designing public signals. *GAME ECON BEHAV.*
- Elkind, E.; Markakis, E.; Obraztsova, S.; and Skowron, P. 2015. Equilibria of plurality voting: Lazy and truth-biased voters. In *SAGT*, 110–122. Springer.
- Erdélyi, G.; Reger, C.; and Yang, Y. 2017. The complexity of bribery and control in group identification. In *AAMAS*, 1142–1150.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2009. Llull and copeland voting computationally resist bribery and constructive control. *J ARTIF INTELL RES* 35:275–341.
- 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. *J ARTIF INTELL RES* 40:305–351.
- Kamenica, E., and Gentzkow, M. 2011. Bayesian persuasion. *AM ECON REV* 101(6):2590–2615.
- Kamenica, E. 2018. Bayesian persuasion and information design. *ANNU REV ECON* 11.
- 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 *AAMAS*, 1737–1738.
- Orlin, J. B. 1997. A polynomial time primal network simplex algorithm for minimum cost flows. *MATH PROGRAM* 78(2):109–129.
- Rayo, L., and Segal, I. 2010. Optimal information disclosure. *J POLIT ECON* 118(5):949–987.
- Rubinstein, A. 2015. Honest signaling in zero-sum games is hard, and lying is even harder. *arXiv preprint arXiv:1510.04991*.
- Schnakenberg, K. E. 2015. Expert advice to a voting body. *J ECON THEORY* 160:102–113.
- Sina, S.; Hazon, N.; Hassidim, A.; and Kraus, S. 2015. Adapting the social network to affect elections. In *AAMAS*, 705–713.
- Taneva, I. A. 2015. Information design.
- Wang, Y. 2013. Bayesian persuasion with multiple receivers. *Available at SSRN* 2625399.
- Wilder, B., and Vorobeychik, Y. 2018. Controlling elections through social influence. In *AAMAS*, 265–273.
- Wilder, B., and Vorobeychik, Y. 2019. Defending elections against malicious spread of misinformation. In *AAAI*, volume 33, 2213–2220.
- Xavier, E. C. 2012. A note on a maximum k-subset intersection problem. *Information Processing Letters* 112(12):471–472.
- Xu, H. 2019. On the tractability of public persuasion with no externalities. *CoRR* abs/1906.07359.