

Algorithmic Aspects of Private Bayesian Persuasion*

Yakov Babichenko¹ and Siddharth Barman²

- 1 The William Davidson Faculty of Industrial Engineering and Management, Technion - Israel Institute of Technology, Haifa, Israel
yakovbab@tx.technion.ac.il
- 2 Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India
barman@csa.iisc.ernet.in

Abstract

We consider a multi-receivers Bayesian persuasion model where an informed sender tries to persuade a group of receivers to take a certain action. The state of nature is known to the sender, but it is unknown to the receivers. The sender is allowed to commit to a signaling policy where she sends a private signal to every receiver. This work studies the computation aspects of finding a signaling policy that maximizes the sender's revenue.

We show that if the sender's utility is a submodular function of the set of receivers that take the desired action, then we can efficiently find a signaling policy whose revenue is at least $(1 - 1/e)$ times the optimal. We also prove that approximating the sender's optimal revenue by a factor better than $(1 - 1/e)$ is NP-hard and, hence, the developed approximation guarantee is essentially tight. When the sender's utility is a function of the number of receivers that take the desired action (i.e., the utility function is anonymous), we show that an optimal signaling policy can be computed in polynomial time. Our results are based on an interesting connection between the Bayesian persuasion problem and the evaluation of the concave closure of a set function.

1998 ACM Subject Classification F.2.0 Analysis of Algorithms and Problem Complexity

Keywords and phrases Economics of Information, Bayesian Persuasion, Signaling, Concave Closure

Digital Object Identifier 10.4230/LIPIcs.ITCS.2017.34

1 Introduction

Sender-receiver models have been extensively studied in economics to understand the role of information in strategic settings; see, e.g., [5, 2, 12, 18]. Since these models study how information, or lack thereof, affects strategic decisions, they have found applications in various domains such as voting [2, 27], regulation policies [21, 28], marketing [8, 3], and auctions [7]. In these models, there is a sender who is more informed than the receiver, and the receiver has to take an action that affects both the sender's and the receiver's utility. An important objective behind studying these models is to quantify the *informational advantage* of the sender. In particular, the goal is to understand the optimal policy using which the sender can transmit (partial) information – or, equivalently send signals – to persuade the receiver into taking an action that is beneficial for the sender. Therefore, research work in

* The first author was supported by Israel Science Foundation, grant # 2021296. The second author was supported in part by Ramanujan Fellowship (SB/S2/RJN-128/2015).



this field is aimed at characterizing signaling policies that yield the sender the maximum possible revenue.

A fundamental sender-receiver model considered in the literature is the *Bayesian persuasion* model, wherein the sender is allowed to *commit* to a signaling (information-revelation) policy *before* she receives the information.¹ In this model, the utilities of the sender and the receiver (obtained from different sender/receiver actions) depend on the state of nature, which is drawn from a prior distribution. This distribution is known to the receiver and the sender. But, only the sender knows the realized state of nature (i.e., the draw) and, hence, has an informational advantage. In this Bayesian model – building upon the classical work by Aumann and Maschler [5] – Kamenica and Gentzkow [18] study the case in which there is exactly one receiver and they obtain a clean characterization of the sender’s optimal signaling policy and her optimal revenue. However, this clean characterization fails to hold when there are multiple receivers; this is true even if the receivers have no payoff externalities.

In fact, except for a few specific instances [29, 11], little is known about optimal signaling policies in settings wherein the sender has to transmit to multiple receivers. To address this limitation, in a recent work Arieli and Babichenko [4] examine the following multiple-receiver setting, which will also be the focus of our work: there are two states of nature, a single sender, and n receivers with binary ($\{0,1\}$) actions. Like the standard Bayesian setting, in this model the state of nature is drawn from a prior distribution that is known to the receivers and the sender; but, only the sender has access to the realized state. Each receiver’s utility depends on her action and the state of nature (which she does not know). Furthermore, in this model the sender is allowed to send *private* signals to the receivers, and her utility is a monotonic set function of the *profile* of actions of the receivers, i.e., the sender’s utility is a function which depends upon the set of receivers that, say, play action one.

Our goal is to examine the algorithmic aspects of this model. In particular, we focus on the computation of an (approximately) optimal policy for the sender. When the sender’s utility is a submodular function, we show that a signaling policy whose revenue is at least $(1 - \frac{1}{e} - \epsilon)$ times the sender’s optimal revenue can be found in polynomial time; here $\epsilon > 0$ is an arbitrarily small constant. In addition, we establish that there does not exist a polynomial-time algorithm which approximates the sender’s optimal revenue by a factor better than $(1 - \frac{1}{e} + \epsilon)$, unless $P = NP$. Hence, the obtained approximation ratio is almost tight. For the case in which the sender’s utility is anonymous – i.e., depends only on the number of receivers that play action one and not on their identities – we develop a polynomial algorithm for finding the optimal policy.

Our results are based on an interesting connection between the sender’s optimal policy and the *concave closure* of a set function. We observe that computation/approximation, of an optimal signaling policy is equivalent to computation/approximation of the concave closure. Our signaling results follow from the following analogous results for the concave closure: a tight $(1 - \frac{1}{e})$ approximation bound for the concave closure of submodular functions and a polynomial-time algorithm for the concave closure of anonymous functions. Since concave closure is a fundamental object of study in discrete convex analysis (see, e.g., [26]) our results for the evaluation of the concave closure might be of independent interest.

Although our focus is on a basic setting (in particular, on a setting in which the receivers have no payoff externalities,² they have binary actions, and the state of nature is binary), the considered model does capture several interesting – albeit stylized – scenarios.

¹ The Bayesian persuasion model complements the classical *cheap talk* model [12] which assumes that sender decides which signal to send after she receives the information.

² That is, each receiver’s utility depends on her action and the state of nature, and not on the action of other receivers.

For example, the model can be used to represent a marketer (a sender) who is trying to persuade consumers (the receivers) to buy a certain product. Here, it is natural to assume that the marketer has more information about the quality of the product than the consumers. Also, different potential consumers may have different utilities for adopting the product. Note that each consumer has two possible (binary) actions, either to buy the product or not. In addition, we can consider the product (state of nature) to be in one of two possible states: the product is either high-quality or low-quality. Furthermore, the marketer's production cost may not be linear. Since it is reasonable to assume that the marketer's utility depends on the number of consumers that buy the product, this example is exactly captured by our result on anonymous utilities. In particular, using the algorithm developed for anonymous utilities we can efficiently compute an optimal signaling policy for the marketer.

Along these lines, the considered model also captures a viral marketing scenario where the receivers are the "opinion leaders", say, in a social network. After persuading a subset of "opinion leaders" to adopt a product, the information about the product will be spread through the network according to some diffusion process. As demonstrated in the notable work for Kempe et al. [23], many diffusion processes satisfy the submodularity property. Hence, our result on submodular utilities can be applied in such settings.

Another example is that of a lobbyist who is trying to persuade politicians to support a certain proposal. The proposal will pass if the number of supporters is above a specified threshold. Under the assumption that politicians vote sincerely³ (i.e., they vote in favor of the alternative that maximizes their utility, given the information they possess about the proposal), this example is captured by the anonymous supermajority utility case.

Techniques. As mentioned above, our proofs are based on an interesting connection between the Bayesian persuasion model and the *concave closure of a set function* [25, 13, 30]. Specifically, in Section 2 (Lemma 1) we show that computation (or approximation) of sender's optimal revenue with utility $V : 2^{[n]} \rightarrow \mathbb{R}_+$ is (computationally) equivalent to evaluation (or approximation) of the concave closure of the function V , here $[n]$ is the set of receivers. Concave closure has been studied in the context of submodular maximization, see, e.g., [10] and [30]. In particular, prior work has shown that even though the concave closure provides a tight relaxation for constrained maximization of submodular functions, it is NP-hard to compute. Hence, instead of focusing on the computation of the concave closure, approximation results for submodular maximization typically rely on finding the *multilinear relaxation*. It turns out that in the context of the Bayesian persuasion problem the concave closure is not just a technical tool, but a core object. Therefore, a key focus of the paper is on efficiently approximating concave closures.

Specifically, we develop a $(1 - \frac{1}{e})$ -approximation algorithm for computing the concave closure of monotone submodular functions. Our tight approximation result rests on an approximation preserving reduction between computing the concave closure and the problem of maximizing a monotone submodular function subject to a matroid constraint. We obtain such a reduction by a careful rounding of the problem parameter to a discrete grid. Since submodular maximization under matroid constraints admits a $(1 - \frac{1}{e})$ approximation (see [10]) the desired result follows. For the hardness result, we use tools from [24] and [17] which were developed to establish the hardness of approximating the maximum social welfare in combinatorial auctions.

³ Note that this simplifying assumption assumes that politicians adopt a simple behavioral rule rather than a sophisticated equilibrium behavior. In equilibrium, each politician should condition his vote on the event that his vote will be pivotal, which leads the politicians to a completely different behavior, see e.g., [6].

We establish the result for anonymous utility functions (which are not necessarily submodular/concave) by developing a polynomial-time algorithm for computing the concave closure of anonymous functions. This result builds upon a lemma from [4], which characterizes the maximum “mass” that can be assigned to subsets of size $k \leq n$ under given marginal constraints; see Lemma 3 below for details. In addition to Lemma 3, to obtain the result we show a non-trivial property that in this case the maximal assignment is “monotonous across all ks ,” see details in Lemma 4. This additional property allows us to formulate the original concave closure problem as an LP with a *polynomial* number of variables.

1.1 Additional Related Work

The current literature on Bayesian persuasion starts with the result of [18] who – building upon the classical work by [5] – analyze the case of a single sender and a single receiver. Several extensions of this model – including ones that consider multiple receivers – have been studied in recent years, see, e.g., [1, 19, 20]. The setting wherein the sender is only allowed to send a public signal to the receivers is considered in [2, 27]. Furthermore, the complementary setting in which the sender is allowed to send private signals to the receivers has been studied in [4, 29, 31]. Our result is most closely related to the work of [4] where the optimal policy and the optimal revenue are characterized for supermodular utilities, utilities that are both anonymous and submodular, and also for supermajority utilities. In particular, we build upon the work of [4] with a computation perspective. We show that for *every* anonymous utility the optimal revenue can be computed in polynomial time. This result generalizes the claim in [4] where the same result (along with a closed-form expression for the sender’s revenue) was established for anonymous utilities that are also submodular. We provide a tight approximation result for submodular utilities. Our inapproximability result for submodular functions (Theorem 9) indicates that it is unlikely that there exists a closed-form expression for the sender’s revenue when her utility function is submodular.

A number of recent results in the computer science community have examined algorithmic questions surrounding the above mentioned models and signaling in general [15, 14, 16, 22, 9]. In particular, an interesting paper by Dughmi and Xu [15] studies the complexity of Bayesian persuasion in the single-receiver model of Kamenica and Gentzkow [18]. They consider the case in which the receiver has n actions and there are $\exp(n)$ states of nature. Dughmi and Xu [15] show that when the payoff profiles are i.i.d. distributed (for all receiver’s actions) the problem can be solved in polynomial time. The same is not true if the payoff profiles are independently, but not identically, distributed – in this case the problem becomes $\#P$ -hard. Finally, they also show that the general problem (with arbitrary payoff profiles) can be approximately solved efficiently in a query model, if we assume that the receiver follows the recommended action by the sender in all cases where no ε -better action (according to her belief) exists.⁴

2 Notations and Preliminaries

We consider a Bayesian persuasion model with a sender and n receivers, $[n] = \{1, 2, \dots, n\}$. Write $\Omega = \{\omega_0, \omega_1\}$ to denote the two possible states of nature. Each receiver $i \in [n]$ has two actions, $\{0, 1\}$, and a utility function, u_i , that depends on the state of nature and its own action, $u_i : \Omega \times \{0, 1\} \rightarrow \mathbb{R}$. All receivers share a common prior distribution, where

⁴ This notion is called ε -incentive compatibility.

$0 < \gamma < 1$ is the probability of state ω_1 , and $1 - \gamma$ of state ω_0 . Note that even though the receivers' utilities are dependent on the realized state of nature they are a priori unaware of it. Throughout, we will use $\Delta(A)$ to denote the set of probability distributions over set A .

It is shown in [4] that, without loss of generality, we can assume that $u_i(\omega_0, 0) > u_i(\omega_0, 1)$ and $u_i(\omega_1, 1) > u_i(\omega_1, 0)$ for all receivers $i \in [n]$. In particular, it is shown in [4] that we can efficiently reduce an instance with arbitrary utility functions to an instance in which the receivers prefer to play 1 when the state of nature is ω_1 and prefer to play 0 when the state is ω_0 . Hence, throughout the paper we will work with this assumption on receivers' utilities.

As mentioned earlier, the sender's utility, V , depends on the set of receivers that play action 1, $V : \{0, 1\}^n \rightarrow \mathbb{R}$. With a slight abuse of notation, for a subset $S \subseteq [n]$, we will use $V(S)$ to denote $V(1_S)$, where 1_S is the characteristic vector of subset S . Throughout we will assume that the sender's utility monotonically increases with the set of receivers that play action 1: $V(S) \leq V(T)$ for every $S \subseteq T$.

Note that we have restricted our attention to the case wherein the sender's utility does not depend on Ω . More generally, the sender's utility can be defined to be a function of the state of nature as well, i.e., we can have $V : \Omega \times \{0, 1\}^n \rightarrow \mathbb{R}$. It is shown in [4] that such a general case can always be efficiently reduced to a setting in which $V(\omega_0, S) = V(\omega_1, S)$ for each S . Hence, in this paper we will focus on utility functions, V , that are state independent.

Recall that the utility function V is said to be submodular if it satisfies the decreasing marginal property. That is, for every subset $S \subset T \subseteq [n]$ and each $i \in [n] \setminus T$, we have $V(S \cup \{i\}) - V(S) \geq V(T \cup \{i\}) - V(T)$.

As is typical in Bayesian persuasion models, we assume that only the sender knows the realized state. The receivers remain unaware of it. Furthermore, following the model of Kamenica and Gentzkow [18] we allow the sender to commit in advance to an information-revelation/signaling policy. In this work, however, we allow the sender to reveal the information to every receiver privately. This translates to a state dependent signaling distribution. Formally a policy of the informed sender consists of n finite sets $\{\Theta_i\}_{i=1, \dots, n}$, where Θ_i is the private signal set of receiver i , and a mapping $F : \Omega \rightarrow \Delta(\Theta_1 \times \dots \times \Theta_n)$. Write $\Theta := \Theta_1 \times \dots \times \Theta_n$. The sender can commit to a policy F that is known to the receivers prior to stage in which the state ω is realized.

The sequence of the interaction between the sender and the receivers is as follows. First, the sender commits to a signaling policy F . Then, a state $\omega \in \Omega$ is realized in accordance with the prior $(\gamma, 1 - \gamma)$. After that a profile of signals $\theta = (\theta_1, \dots, \theta_n)$ is generated according to the distribution $F(\omega)$. Every receiver i observes her private signal realization $\theta_i \in \Theta_i$ and forms a posterior $\mathbb{P}_F(\omega_1 | \theta_i) = p(\theta_i)$.

With the posterior in hand, receiver i selects an action that maximizes her expected utility. In other words, receiver i plays action 1 if and only if

$$p(\theta_i)u_i(\omega_1, 1) + (1 - p(\theta_i))u_i(\omega_0, 1) \geq p(\theta_i)u_i(\omega_1, 0) + (1 - p(\theta_i))u_i(\omega_0, 0).$$

We assume that in case of indifference, receivers plays action 1. Let $g_i(\theta_i) \in \{0, 1\}$ denote receiver i 's best-reply action when she observes the signal θ_i . Also, write $g(\theta)$ to be the action profile of the receivers when the realized vector of signals is θ . We will use $F_1 \in \Delta(\Theta)$ to denote the signal distribution conditional on state ω_1 and $F_0 \in \Delta(\Theta)$ to denote the signal distribution conditional on state ω_0 .

Let $s(F)$ be the sender's utility from the policy (Θ, F) :

$$s(F) := \gamma \mathbb{E}_{\theta \sim F_1}[V(g(\theta))] + (1 - \gamma) \mathbb{E}_{\theta \sim F_0}[V(g(\theta))]. \quad (1)$$

A signaling policy (Θ, F) is said to be *optimal* if it maximizes sender's utility among all possible signal sets Θ and all possible signals $F : \Omega \rightarrow \Delta(\Theta)$.

We begin by stating a result from [4] (see Lemma 1 in [4]) that shows the existence of an optimal policy with the following useful properties:

- For every receiver i , the private signal set Θ_i is equal to $\{0, 1\}$ and receiver i 's best reply $g_i(\theta_i) = \theta_i$. In other words, when signal θ_i is recommended by the sender to receiver i it is profitable (after receiver i performs a Bayesian update of her belief on the state of the world) for her to follow the recommendation. In [18], such policies are called *straightforward*.
- In the optimal policy $F_1(1, 1, \dots, 1) = 1$, i.e., when state ω_1 is realized the sender recommends everyone to adopt the product. Recall that F_1 is a distribution over the set Θ , which (by the previous property) is $\{0, 1\}^n$ for the optimal policy under consideration.
- When the realized state is ω_0 , the sender recommends to receiver i to adopt the product with probability of at most $a_i := \min \left\{ \frac{\gamma}{1-\gamma} \frac{u_i(\omega_{1,1}) - u_i(\omega_{1,0})}{u_i(\omega_{0,0}) - u_i(\omega_{0,1})}, 1 \right\}$. Write marginal $F_0(\theta_i = 1) := \sum_{\theta \in \{0,1\}^n: \theta_i=1} F_0(\theta)$. We succinctly express this condition as $F_0(\theta_i = 1) \leq a_i$. The number a_i can be interpreted as the maximal probability that the sender can “lie” to the receiver, and will be called the *persuasion level of player i* .

Under such an optimal policy, the sender's utility is given by

$$s(F) = \gamma V([n]) + (1 - \gamma) \mathbb{E}_{\theta \sim F_0} V(\theta) \quad (2)$$

Overall, in light of these properties the problem of determining an optimal policy (over general signal sets Θ and mappings $F : \Omega \rightarrow \Delta(\Theta)$) reduces to the following well-structured maximization problem:

$$\text{maximize } \mathbb{E}_{\theta \sim F_0} V(\theta) \quad \text{subject to } F_0(\theta_i = 1) \leq a_i \quad \forall i \in [n]. \quad (3)$$

Note that the prior γ and the utility $V([n])$ are fixed parameters. Hence, an optimal solution of (3) gives us an optimal solution of (2).

For each subset $S \subset [n]$, with characteristic vector 1_S , write μ_S to be the probability that exactly the receivers in S will receive the recommendation to adopt the product, i.e., $\mu_S := F_0(1_S)$. For a given persuasion levels profile $a := (a_1, \dots, a_n)$, the maximization problem (3) can be written as

$$\begin{aligned} V^+(a) := \max & \sum_{S \subseteq [n]} \mu_S V(S) \\ \text{s.t.} & \sum_{S \subseteq [n]} \mu_S 1_S \leq a \\ & \sum_{S \subseteq [n]} \mu_S = 1 \\ & \mu_S \geq 0. \end{aligned} \quad (4)$$

This is exactly the definition of the concave closure of the set function V evaluated at a given vector a .

Throughout, we will use $V^+(a)$ to denote the concave closure of the sender's utility function V at $a \in [0, 1]^n$. We will refer to solving (approximating) the optimization problem (4) as computing (approximating) the concave closure. Interestingly, the concave closure has been studied in the optimization literature. In particular, it is considered as a technical tool for submodular maximization; see, e.g., [10]. Note that computing the concave closure corresponds to solving a linear programming with polynomial number of constraints, but with an exponential number of variables (the variables are μ_S for every $S \subset [n]$).

In some cases we will be interested in approximating the optimal revenue of the sender and, hence, we introduce here the following lemma that states that computing (approximating) the concave closure is computationally equivalent to computing (approximating) the sender's optimal revenue. Note that, for a given parameter $\alpha \in (0, 1]$, an α approximation of the concave closure corresponds to a distribution $\{\mu_S\}_{S \subseteq [n]}$ that satisfies the feasibility constraints of the optimization problem (4) and obtains an objective function value, $\sum_{S \subseteq [n]} \mu_S V(S)$, that is at least α times the optimal. The next lemma states that there exists an approximation-preserving, polynomial-time reduction between computing the concave closure and finding the optimal revenue of the sender. Specifically, the lemma establishes that computing the concave closure of the sender's utility function lies at the core of determining a revenue-maximizing policy for the sender.

► **Lemma 1.** *Given a persuasion profile $a \in [0, 1]^n$ and utility function V along with an α approximation of the concave closure $V^+(a)$. We can find, in polynomial time, a policy for the sender (with utility function V and persuasion profile a) that obtains revenue at least α times the optimal.*

Furthermore, for every $\varepsilon > 0$, there exists a polynomial-time reduction from the problem of α approximating a sender's revenue (with utility function V and persuasion profile a) to the problem of computing an $(\alpha + \varepsilon)$ approximation of the concave closure $V^+(a)$.

Proof. The forward direction is direct: by equation (2), an α approximation of $\max \mathbb{E}_{\theta \sim F_0} V(\theta)$ is also an α approximation of $\gamma V(N) + (1 - \gamma) \mathbb{E}_{\theta \sim F_0} V(\theta)$.

For the other direction, given function V and persuasion level profile $a = (a_i)_{i \in [n]}$, we can set the prior γ to be very small (e.g., $\gamma = \frac{\varepsilon(V(N))^2}{1 - \alpha}$ suffices) and we set receiver i utilities to be

$$u_i(\omega_0, 0) = 1, \quad u_i(\omega_0, 1) = u_i(\omega_1, 0) = 0, \quad u_i(\omega_1, 1) = a_i \frac{1 - \gamma}{\gamma}.$$

Such a choice guarantees that indeed $a_i = \min \left\{ \frac{\gamma}{1 - \gamma} \frac{u_i(\omega_1, 1) - u_i(\omega_0, 1)}{u_i(\omega_0, 0) - u_i(\omega_0, 1)}, 1 \right\}$. It follows that for such instances an α approximation of the sender's revenue implies $(\alpha + \varepsilon)$ approximation of the concave closure of V . ◀

In subsequent sections we establish algorithmic and hardness results for the problem of finding the optimal policy (and revenue) of the sender. We do so by using the above mentioned lemma and, in particular, addressing the computation of the concave closure.

3 Anonymous Utility

This section considers the case wherein the sender's utility function is anonymous i.e., it satisfies $V(S) = f(|S|)$ for some monotonically increasing function $f : [n] \rightarrow \mathbb{R}$. Our main result for anonymous utilities is as follows.

► **Theorem 2.** *There exists a polynomial algorithm for computing the maximum revenue and an optimal signaling policy for a sender that has a monotone, anonymous utility function.*

3.1 Proof of Theorem 2

We show that the concave closure of anonymous function can be computed in polynomial time.

We use \mathcal{S}_k to denote all the size- k subsets of $[n]$, $\mathcal{S}_k := \{S \subseteq [n] \mid |S| = k\}$. We denote by $\text{marg}(\mu)_i := \sum_{S \subseteq [n]: i \in S} \mu(S)$ the marginal probability of the i th coordinate to be equal 1. Note that the constraints of the concave closure $V^+(a)$ can be written as $\text{marg}(\mu)_i \leq a_i$ for every $i \in [n]$.

The following lemma from [4] characterizes the maximum probability mass that can be assigned to subsets of size k under the constraints imposed by the profile $b = (b_1, \dots, b_n)$.

► **Lemma 3 ([4]).** *Let $1 \geq b_1 \geq b_2 \geq \dots \geq b_n \geq 0$ be a monotonic sequence. The optimal value of the following maximization problem*

$$\begin{aligned} \max \quad & \sum_{S \in \mathcal{S}_k} \nu(S) \\ \text{s.t.} \quad & \sum_{S: i \in S} \nu(S) \leq b_i \quad \forall i \in [n] \\ & \nu_S \geq 0 \end{aligned} \tag{5}$$

where ν is a positive measure (not necessarily a probability measure), is equal to

$$\beta_k(b_1, \dots, b_n) = \min_{0 \leq m < k} \frac{1}{k - m} (b_{m+1} + \dots + b_n).$$

Moreover, a measure ν that maximizes (5) can be computed in polynomial time.

The key idea is to use this lemma to solve the LP corresponding to the concave closure – which has exponential (in n) number of variables – by another LP that has a polynomial number of variables. First, we assume, without loss of generality, that the point a (where we want to evaluate the concave closure) satisfies $a_1 \geq a_2 \geq \dots \geq a_n$. We split the original problem into n problems of finding a measure μ_k over \mathcal{S}_k for every $k = 1, \dots, n$ (the final measure is defined by $\mu = \mu_1 + \dots + \mu_n$). The new maximization problem has n^2 variables $(a_i^j)_{i,j \in [n]}$, where (a_1^j, \dots, a_n^j) represents the marginal constrain vector on subsets of size k . We denote by α_k the measure that is assigned to subsets of size k , then the original maximization problem can be translated to the following

$$\begin{aligned} \max \quad & \alpha_1 f(1) + \alpha_2 f(2) + \dots + \alpha_n f(n) \\ \text{s.t.} \quad & \sum_{k \in [n]} \alpha_k = 1, \quad 0 \leq \alpha_k \leq \beta_k(a_1^k, a_2^k, \dots, a_n^k) \text{ for } k \in [n], \text{ and } \sum_{j \in [n]} a_i^j \leq a_i. \end{aligned} \tag{6}$$

where the first constraint corresponds to $\sum_{S \subseteq [n]} \mu_S = 1$, the second follows from Lemma 3, and the last constraint uses the fact that marginals preserve additivity, and thus correspond to $\sum_{S \subseteq [n]} \mu_S 1_S \leq a$.

Note that the only nonlinear constraints in (6) are $\alpha_k \leq \beta_k(a_1^k, a_2^k, \dots, a_n^k)$. Interesting, these constraints are “almost linear” in the following sense: If $a_1^k \geq a_2^k \geq \dots \geq a_n^k$ then the constraint $\alpha_k \leq \beta_k(a_1^k, a_2^k, \dots, a_n^k)$ can be written as

$$\begin{cases} \alpha_k \leq \frac{1}{k} (a_1 + a_2 + \dots + a_n) \\ \alpha_k \leq \frac{1}{k-1} (a_2 + a_3 + \dots + a_n) \\ \vdots \\ \alpha_k \leq \frac{1}{1} (a_k + a_{k+1} + \dots + a_n) \end{cases}$$

So the only obstacle is that the marginal constrain vector, in principle, is not guaranteed to satisfy the monotonicity constraint $a_1^k \geq a_2^k \geq \dots \geq a_n^k$ (for all $k \in [n]$). The following Lemma 4 proves that, in fact, there always exists an optimal solution that satisfies this monotonicity

constraint (for all $k \in [n]$). Therefore we can impose this constraint in the optimization problem (6), and then it becomes an LP maximization with $\text{poly}(n)$ variables (and $\text{poly}(n)$ constraints).

We also note that the proof of Lemma 3 in [4] is constructive, and computationally efficient. Thus, to compute an optimal policy (not only optimal revenue) after we have computed the values of $(a_i^j)_{i,j \in [n]}$ that maximize (6) we can use the constructive algorithm of the proof of Lemma 3 for all $k \in [n]$. This completes the proof of Theorem 2.

► **Lemma 4.** *Assume that the point a satisfies $a_1 \geq a_2 \geq \dots \geq a_n$. For $(\mu_S)_{S \subset [n]}$ we denote by $a_i^k = \sum_{S \in \mathcal{S}_k: i \in S} \mu(S)$. There exists μ that maximizes (6) and satisfies in addition $a_1^k \geq a_2^k \geq \dots \geq a_n^k$ for all $k \in [n]$.*

Proof. The proof builds upon ideas that were used in the proof of Lemma 3 in [4].

Let ν be a distribution that satisfies the constraints $\text{marg}(\nu)_i \leq a_i$, and let $\alpha_k = \nu(\mathcal{S}_k)$ be the weight of ν on subsets of size k . It is sufficient to construct another distribution μ that satisfies $\mu(\mathcal{S}_k) = \alpha_k$ (and thus μ has the same revenue as ν), and in addition $a_1^k \geq a_2^k \geq \dots \geq a_n^k$.

The construction is done in n steps, where the steps $k = n, n - 1, \dots, 1$ are done in an decreasing order. At step k we assign a measure of α_k to subsets of size k , and we denote the assigned measure by μ_k . Each step k is done in finite number of stages. Here we describe the assignment of measure at stage $k.m$.

During the construction we “assign mass” and thus, we “spend marginal constraints.” We take track of the remaining marginal constraints vector. At the beginning, we set the constraints vector $(a_1^{n,0}, \dots, a_n^{n,0}) = (a_1, \dots, a_n)$ to be the original constraints.

During the process we preserve the monotonicity of the marginal constraints vector and therefore we can denote the marginal constraints vector at stage $k.m$ by

$$(a_1^{k,m}, \dots, a_n^{k,m}) = (b_1, \dots, b_j, \underbrace{c, c, \dots, c}_{l-j \text{ times}}, b_{l+1}, \dots, b_n)$$

where $b_j > c > b_{l+1}$ and $j < k \leq l$. Note that if $a_k^{k,j} = a_{k+1}^{k,j} = \dots = a_n^{k,j}$ then $l = n$ and for simplicity of notation we denote $b_{n+1} = 0$. Note that if $a_1^{k,j} = a_2^{k,j} = \dots = a_k^{k,j}$ then $j = 0$, and for simplicity of notation we denote $b_0 > b_1$.

At stage $k.m$, the idea is to distribute mass equally over the subsets S of size k that satisfy $[j] \subseteq S \subseteq [l]$ (we have $\binom{l-j}{k-j}$ such sets). If we do so, after we have distributed x units of mass the remaining marginal constraints vector will be

$$b(x) = (b_1 - x, \dots, b_j - x, c - \frac{k-j}{l-j}x, \dots, c - \frac{k-j}{l-j}x, b_{l+1}, \dots, b_n) \tag{7}$$

because every element $i = j + 1, j + 2, \dots, l$ appears in exactly $\frac{k-j}{l-j}$ fraction of the above subsets. Step $k.m$ terminates at the moment when one of the following three happens:

- (1) The total mass that has been assigned during step k reaches α_k . In such a case we proceed to step $k - 1$.
- (2) The j th coordinate becomes equal to the $(j + 1)$ th coordinate. In such a case we proceed to stage $k.(m + 1)$.
- (3) The l th coordinate becomes equal to the $(l + 1)$ th coordinate. In such a case we proceed to stage $k.(m + 1)$.

We denote by $\alpha_{k.m}$ the amount of mass that has been assigned during step $k.m$. We denote by $b(\alpha_{k.m})$ the marginal constraints vector after step $k.m$, where $b(\cdot)$ is defined in

equation (7). This marginal constraints serves as the marginal constraint vector for the next step (in case (1) happens) or the next stage (in case (2) or (3) happens).

We argue the following two statements, which will complete the proof.

1. The described process succeeds to complete all the n steps.
2. The described process at each step k assigns mass in a way that $a_1^k \geq a_2^k \geq \dots \geq a_n^k$.

Statement (2) follows from the fact that at each stage $k.m$ the marginals of the assigned mass is of the form $(\underbrace{x, \dots, x}_{j_m \text{ times}}, \underbrace{cx, \dots, cx}_{l_m - j_m \text{ times}}, 0, \dots, 0)$ for $x = \alpha_{k.m}$ and $c < 1$. Moreover, during step k the coordinate j_m is monotonically decreasing, and the coordinate l_m is monotonically increasing. Therefore, the sum of those vectors, which is equal to the vector $(a_1^k, a_2^k, \dots, a_n^k)$ is monotonically increasing.

Assume by way of contradiction that statement (1) is false. The above process cannot assign the required measure only if we are at step k and the marginal constraints vector becomes $(d_1, d_2, \dots, d_m, 0, \dots, 0)$ for $m < k$. In such a case indeed the process cannot proceed, because it will turn the $m + 1$ coordinate of the marginal constraint vector negative. We denote by α'_k the measure at step k that has been assigned up to the moment of termination.

We argue that this is impossible from the fact that $\alpha_n, \alpha_{n-1}, \dots, \alpha_k$ are feasible weights for some distribution ν . The idea is that the described above process has minimal marginals on the elements $m + 1, \dots, n$, thus if this process cannot proceed neither could some other distribution ν . Formally, we denote $\nu = \nu_1 + \dots + \nu_n$, where ν_j is a measure over \mathcal{S}_j . Note that $|\nu_j| = \alpha_j$. We denote $(d_i^j)_{i \in [n]}$ the marginals of ν_j . We argue that $\sum_{i=m+1}^n d_i^j \geq (j - m)\alpha_j$, because every subset of size j contains at least $j - m$ elements from the set $\{m + 1, \dots, n\}$. Therefore we have

$$a_{m+1} + \dots + a_n \geq \sum_{j=k}^n \sum_{i=m+1}^n d_i^j \geq \sum_{j=k}^n (j - m)\alpha_j \quad (8)$$

On the other hand, the constructed measure μ_n with marginals $(a_i^j)_{i \in [n]}$ satisfies $\sum_{i=m+1}^n a_i^j = (j - m)\alpha_j$, because this process assigns positive probability *only* to subsets that contain $\{1, \dots, m\}$ (because $m < k \leq j$ and $a_m^j > a_{m+1}^j$). Since the process spent all the marginal constraints a_{m+1}, \dots, a_n we have

$$a_{m+1} + \dots + a_n = \sum_{j=k}^n \sum_{i=m+1}^n a_i^j = (k - m)\alpha'_k + \sum_{j=k+1}^n (j - m)\alpha_j < \sum_{j=k+1}^n (j - m)\alpha_j \quad (9)$$

Inequalities (8) and (9) yield a contradiction. ◀

4 Submodular Utilities

This section considers private Bayesian persuasion settings in which the sender's utility function is submodular. In particular, we develop a tight $(1 - 1/e)$ approximation of the optimal signaling policy when the sender's utility is a monotone submodular function.

It is relevant to note that our algorithmic results require only query access to the submodular function, i.e., our results hold as long as we can access to $V(S)$, for any subset $S \subseteq [n]$. This, in particular, implies that we can address submodular functions that admit a *succinct* representation.

We begin by noting that finding the concave closure of a submodular function is NP-hard: Given a succinct, monotone, submodular function $f : 2^{[n]} \rightarrow \mathbb{R}$ and a vector $a \in [0, 1]^n$, it is NP-hard to compute the concave closure $f^+(a)$; see, e.g. [30] and [13].

It is relevant to note that while the concave closure of submodular functions are known to be computationally hard, approximation algorithms and inapproximability results for them have not been directly addressed in prior work.

4.1 Approximation Algorithm for Submodular Utilities

This section provides a $(1 - 1/e)$ -approximation algorithm for computing the concave closure of a monotone, submodular function V . We obtain the $(1 - \frac{1}{e})$ approximation by reducing the computation of the concave closure to the problem of maximizing a submodular function subject to a matroid constraint. The key implication of this approximation result is the following theorem.

► **Theorem 5.** *If in a private Bayesian persuasion problem the utility of the sender, V , is a monotone, submodular function. Then, in polynomial time, we can compute a signaling policy that achieves a revenue of at least $(1 - 1/e - \varepsilon)$ times the optimal; here, ε is an arbitrarily small constant.*

4.1.1 Proof of Theorem 5

We show that the concave closure of submodular function V at any given vector $a = (a_1, a_2, \dots, a_n) \in [0, 1]^n$ can be approximated to within a factor of $(1 - \frac{1}{e} - \varepsilon)$, for an arbitrarily small $\varepsilon > 0$, then Theorem 5 follows from Lemma 1.

We split the marginal values a_i into two sets: $\{a_i : a_i \geq \frac{1}{n^2}\}$ are the *high values* and $\{a_i : a_i < \frac{1}{n^2}\}$ are the *low values*. Without loss of generality we assume that a_1, \dots, a_m are the high values and a_{m+1}, \dots, a_n are the low values, for $m \leq n$.

Every distribution μ over subsets of $[n]$ induces a distribution $\nu = \nu(\mu)$ over subsets of $[m]$ in the following natural way: the probability mass μ_S on $S \subseteq [n]$ is moved to the set $S \cap [m]$, formally for each subset $T \subseteq [m]$ define $\nu_T := \sum_{S \subseteq [n]: S \cap [m] = T} \mu_S$. The following lemma holds for distribution ν .

► **Lemma 6.** *For every distribution μ that satisfies the marginal constraints (i.e., $\sum_{S \subseteq [n]: S \ni i} \mu_S \leq a_i$) we have*

$$\sum_{S \subseteq [n]} \mu_S V(S) \leq \sum_{T \subseteq [m]} \nu_T V(T) + \sum_{i=m+1}^n a_i V(\{i\}).$$

Here $a_1, \dots, a_m \geq \frac{1}{n^2}$ and $a_m, a_{m+1}, \dots, a_n \leq \frac{1}{n^2}$.

Proof.

$$\begin{aligned} \sum_{S \subseteq [n]} \mu_S V(S) &\leq \sum_{S \subseteq [n]} \mu_S \left[V(S \cap [m]) + \sum_{i \in S, i > m} V(\{i\}) \right] \\ &= \sum_{T \subseteq [m]} \nu_T V(T) + \sum_S \sum_{i \in S, i > m} \mu_S V(\{i\}) \\ &= \sum_{T \subseteq [m]} \nu_T V(T) + \sum_{i > m} \sum_{S: i \in S} \mu_S V(\{i\}) \\ &\leq \sum_{T \subseteq [m]} \nu_T V(T) + \sum_{i > m} a_i V(\{i\}). \end{aligned}$$

Here, the first inequality follows from subadditivity of V . The second equation follows from the definition of $\nu = \nu(\mu)$. The third equation is obtained by changing the order of summation and the last inequality follows from the fact that μ satisfies the marginal constraints. ◀

We can consider the optimization problem corresponding to the concave closure restricted to the set $[m]$:

$$\begin{aligned} V_m^+(a) &:= \max \sum_{T \subseteq [m]} \nu_T V(T) \\ \text{s.t.} \quad &\sum_{T \subseteq [m]: T \ni i} \nu_T \leq a_i \quad \forall i \in [m] \\ &\nu \text{ is a probability measure.} \end{aligned} \tag{10}$$

Given a distribution $\bar{\nu}$ that α -approximates problem (10), we define distribution $\bar{\mu}$ over $[n]$ as follows: For each subset $T \subseteq [m]$, set $\bar{\mu}_T := (1 - \frac{1}{n})\bar{\nu}$. In addition, for every $i > m$ set $\bar{\mu}_{\{i\}} := a_i$. Finally, to ensure that $\bar{\mu}$ is a probability measure we assign a probability mass of $c = \frac{1}{n} - a_{m+1} - \dots - a_n > 0$ to the empty set, i.e., $\mu_\emptyset := c$.

► **Lemma 7.** *If distribution $\bar{\nu}$ α -approximates problem (10), then $\bar{\mu}$ provides a $(1 - \frac{1}{n})\alpha$ -approximation of the original concave closure problem (4).*

Proof. Recall that $V^+(a)$ denotes the concave closure of function V evaluated at a and, similarly, $V_m^+(a)$ is optimal value of (10).

$$\begin{aligned} \sum_{S \subseteq [n]} \bar{\mu}_S V(S) &= (1 - \frac{1}{n}) \sum_{T \subseteq [m]} \bar{\nu}_T V(T) + \sum_{i > m} a_i V(\{i\}) \\ &\geq (1 - \frac{1}{n}) \alpha V_m^+(a) + \sum_{i > m} a_i V(\{i\}) \\ &\geq (1 - \frac{1}{n}) \alpha [V_m^+(a) + \sum_{i > m} a_i V(\{i\})] \\ &\geq (1 - \frac{1}{n}) \alpha V^+(a). \end{aligned}$$

Here the first equation is implied by the definition of $\bar{\mu}$. The second inequality follows from the fact that $\bar{\nu}$ α -approximates the concave closure (on the set $[m]$). The third inequality is trivial and the last one follows from Lemma 6. ◀

Lemma 7 reduces the original concave closure problem to the problem of computing the concave closure over $[m]$ where (unlike the original problem) we know that $a_i \geq \frac{1}{n^2}$ for each $i \in [m]$. In the remainder of the proof, we consider the later problem. The idea is to translate this problem into a discrete one. A natural way do to so is by rounding the underlying terms to integer multiples of a parameter $\delta := \frac{1}{n^4(n+1)}$ and then working with the multiples, instead of the fractional terms.

Since (10) is a linear program (over variables $\{\nu_T\}_{T \subseteq [m]}$) with at most $n + 1$ non-trivial constraints, without loss of generality we can restrict attention to solutions that have support size of at most $n + 1$.

As mentioned previously, we set a grid of size $\delta := \frac{1}{n^4(n+1)}$, and we consider the maximization problem of V_m^+ where we restrict the probabilities $\{\nu_S\}$ to be integer multiples of δ .

$$\begin{aligned}
\max \quad & \sum_{T \subseteq [m]} \nu_T V(T) \\
\text{s.t.} \quad & \sum_{T \subseteq [m]: T \ni i} \nu_T \leq a_i \quad \forall i \in [m] \\
& \nu \text{ is a probability measure.} \\
& \nu_T \in \{0, \delta, 2\delta, \dots, 1\}.
\end{aligned} \tag{11}$$

► **Lemma 8.** *If distribution $\hat{\nu}$ be an α -approximate solution of optimization problem (11) with support size at most $n + 1$. Then, $\hat{\nu}$ is a $(1 - \frac{1}{n\alpha})\alpha$ -approximate solution of the concave closure $V_m^+(a)$ as well.*

Proof. We prove that if we restrict our attention to probabilities in the set $\{0, \delta, 2\delta, \dots, 1\}$, then we incur at most a multiplicative loss of $(1 - \frac{1}{n\alpha})$. Given a distribution ν with support size at most $n + 1$ we round down the probabilities to integer multiples of δ (and put all the remaining probability mass on the empty set), we denote the resulting distribution by ν' . Formally $\nu'_T = \ell\delta$ where $k = \max\{j \in \mathbb{Z} : j\delta \leq \nu_T\}$. Note that

$$\begin{aligned}
& \sum_T \nu_T V(T) - \sum_T \nu'_T V(T) = \sum_T (\nu_T - \nu'_T) V(T) \leq \sum_T \delta V(T) \\
& \leq \sum_{T \in \text{supp}(\nu)} \frac{1}{n^4(n+1)} V([m]) \leq \frac{1}{n^4} V([m]) \\
& \leq \frac{1}{n^4} \sum_{i \in [m]} V(\{i\}) \leq \frac{1}{n^2} \sum_{i \in [m]} a_i V(\{i\}),
\end{aligned} \tag{12}$$

where the first equality is trivial. The first and the second inequality is a consequence of the rounding and the value of δ . The third inequality follows from the fact that the support size is at most $n + 1$. The subadditivity of V gives us the fourth inequality and the last inequality follows from the fact that $a_i \geq \frac{1}{n^2}$.

Note also that $V_m^+(a) \geq \frac{1}{n} \sum_{i \in [m]} a_i V(\{i\})$ because one feasible option is to put a mass of $\frac{a_i}{n}$ on the singleton $\{i\}$, and the remaining probability mass to put on the empty set. This is indeed feasible because $\sum_i \frac{a_i}{n} \leq \sum_i \frac{1}{n} \leq 1$.

Finally let $\bar{\nu}$ be an α approximation for $V_m^+(a)$, and let $\bar{\nu}'$ be the corresponding rounding. Then

$$\begin{aligned}
\frac{\sum_T \bar{\nu}'_T V(T)}{\sum_T \bar{\nu}_T V(T)} &= 1 - \frac{\sum_T \bar{\nu}_T V(T) - \sum_T \bar{\nu}'_T V(T)}{\sum_T \bar{\nu}_T V(T)} \\
&\geq 1 - \frac{\frac{1}{n^2} \sum_{i \in [m]} a_i V(\{i\})}{\sum_T \bar{\nu}_T V(T)} \\
&\geq 1 - \frac{\frac{1}{n^2} \sum_{i \in [m]} a_i V(\{i\})}{\alpha \frac{1}{n} \sum_{i \in [m]} a_i V(\{i\})} = 1 - \frac{1}{n\alpha}
\end{aligned}$$

where the first inequality follows from (12), and the second one from the fact that $V_m^+(a) \geq \frac{1}{n} \sum_{i \in [m]} a_i V(\{i\})$. ◀

By Lemma 8 we can restrict attention to the discrete problem (11). Note that the discrete problem (11) is equivalent to

$$\begin{aligned}
\max_{S^1, \dots, S^k \subseteq [m]} \quad & \frac{1}{k} \sum_{j=1}^k V(S^j) \\
\text{subject to} \quad & |\{j \in [k] \mid i \in S^j\}| \leq k_i \quad \forall i \in [n]
\end{aligned} \tag{13}$$

where $k = n^4(n+1)$ and $k_i = a_i n^4(n+1)$ for all $i \in [n]$.

We complete the proof of the theorem by showing that (13) admits a $(1 - 1/e)$ approximation. We do so by showing that (13) corresponds to the problem of maximizing a monotone submodular function subject to a matroid constraint.

Consider base set $U = [m] \times [k]$. We get that the size of U is polynomially bounded. For a subset $R = \{(i_1, j_1), (i_2, j_2), \dots, (i_l, j_l)\}$ of U and $j \in [k]$, write R^j to denote the projected subset $\{i' \in [m] \mid (i', j) \in R\}$.

With this notation in hand, define function F for each subset $R = \{(i_1, j_1), (i_2, j_2), \dots, (i_l, j_l)\} \subset U$ as follows

$$F(R) := \frac{1}{k} \sum_{j=1}^k V(R^j). \quad (14)$$

We claim that F is submodular: consider subsets $X \subset Y \subset U$ and element $(i, j) \in U$. Note that $F(X + (i, j)) - F(X) = \frac{1}{k}V(X^j + i) - \frac{1}{k}V(X^j)$ and $F(Y + (i, j)) - F(Y) = \frac{1}{k}V(Y^j + i) - \frac{1}{k}V(Y^j)$. Since $X^j \subset Y^j$, the submodularity (monotonicity) of V implies the submodularity (monotonicity) of F .

Next we consider a partition matroid \mathcal{M} over U . Specifically, we say that a subset $R \subset U$ is independent (with respect to the matroid \mathcal{M}) iff $|\{(i', j') \in R \mid i' = i\}| \leq k_i$ for all $i \in [n]$. Note that this is a partition matroid where the disjoint partitions are $B_i := \{(i, 1), (i, 2), \dots, (i, k)\}$ and the cardinality bounds are k_i s. In other words, we obtain \mathcal{M} by defining R to be an independent subset iff $|R \cap B_i| \leq k_i$ for all i .

Note that if a subset $R \subset U$ is independent then R^1, R^2, \dots, R^k satisfy the constraints of the optimization problem (13), i.e., for an independent R we have $|\{j \in [k] \mid i \in S^j\}| \leq k_i$ for all i .

Overall, we get that optimization problem (13) is equivalent to the following problem:

$$\begin{aligned} & \max_{R \subset U} F(R) \\ & \text{subject to } R \in \mathcal{M} \end{aligned}$$

Since this is a submodular maximization problem subject to a matroid constraint, it admits a $(1 - \frac{1}{e})$ approximation; see [10]. This in turn implies that the original problem admits a $(1 - \frac{1}{n})(1 - \frac{1}{0.62n})(1 - \frac{1}{e}) = (1 - \frac{1}{e} - O(\frac{1}{n}))$ approximation. We can set parameters such that instead of a multiplicative factor of $(1 - \frac{1}{n})(1 - \frac{1}{0.62n})$ in the approximation we get a term that is arbitrarily close one. Hence, we get the desired result.

4.2 Hardness of Approximating the Concave Closure

This section shows that the $(1 - \frac{1}{e})$ approximation guarantee obtained in Section 4.1.1 is tight. In particular, applying the machinery developed by [24] leads us to the following theorem. We note that [24] establish the hardness of approximating maximum social welfare in combinatorial auctions and similar tools were developed in [17] for studying the inapproximability of the domatic number.

► **Theorem 9.** *Given a monotone, submodular function $V : 2^{[n]} \rightarrow \mathbb{R}_+$ and vector $a \in [0, 1]^n$, for any $\varepsilon > 0$, it is NP-hard to approximate the concave closure, $V^+(a)$, by a factor better than $(1 - \frac{1}{e} - \varepsilon)$.*

Proof Sketch. [24] study the combinatorial auction problem where n goods have to be partitioned among m agents whose utilities are submodular functions of the goods assigned

to them. In this problem, the objective is to maximize social welfare, i.e., the sum of the utilities of the receivers. It is shown in [24] that for this problem and any $\varepsilon > 0$ there does not exist a polynomial time algorithm that obtains an approximation ratio better than $(1 - \frac{1}{e} - \varepsilon)$, unless $P = NP$.

Specifically, [24] start with a *label-cover* problem where it is NP-hard to distinguish whether the optimal value, $OPT(L)$, is one or less than a particular constant, $c < 1$. From the given label cover problem they construct a combinatorial auction instance, I , wherein the maximum social welfare, $OPT(I)$ is greater than a threshold, τ if the label cover problem admits a solution of value one. Furthermore, if the optimal value of the label cover problem is less than c – i.e., $OPT(L) \leq c$ – then it must be the case that $OPT(I) \leq (1 - \frac{1}{e} - \varepsilon) \tau$. This, overall, establishes a $(1 - \frac{1}{e} - \varepsilon)$ hardness-of-approximation bound for the combinatorial auction problem.

Interestingly, in the constructed instance I all of the m receivers have the same monotone, submodular utility function, say, $f : 2^{[n]} \rightarrow \mathbb{R}_+$. We claim that approximating the concave closure of constructed function f at marginal vector $a := (\frac{1}{m}, \frac{1}{m}, \dots, \frac{1}{m})$ by a factor better than $(1 - \frac{1}{e} - \varepsilon)$ is NP-hard. In particular, if $OPT(L) = 1$ then $f^+(a) \geq \tau'$ and, moreover, if $OPT(L) \leq c$ then $f^+(a) \leq (1 - \frac{1}{e} - \varepsilon) \tau'$; here τ' is a fixed parameter.

The proof of this claim can be obtained by considering the subsets in the support of an optimal solution, μ^* , of problem (4) defined for function f . Note that the proof given in [24] proceeds by considering the subsets that constitute the partition of goods among receivers in I , instead we can focus on subsets in the support of μ^* to obtain the result for the concave closure. In particular, the arguments presented in [24] go through if, instead of cardinalities, we consider measure of sets and expected values of quantities with respect to μ^* .⁵ This, overall, establishes the desired inapproximability result for the concave closure. ◀

Acknowledgements. The authors thank Uriel Feige for helpful discussions and references.

References

- 1 Ricardo Alonso and Odilon Câmara. On the value of persuasion by experts. Technical report, University of Southern California, Marshall School of Business, 2014.
- 2 Ricardo Alonso and Odilon Câmara. Persuading voters. *Available at SSRN 2688969*, 2015.
- 3 Simon P Anderson and Régis Renault. The advertising mix for a search good. *Management Science*, 59(1):69–83, 2013.
- 4 Itai Arieli and Yakov Babichenko. Private bayesian persuasion. *Available at SSRN 2721307*, 2016.
- 5 Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information*. MIT press, 1995.
- 6 David Austen-Smith and Jeffrey S Banks. Information aggregation, rationality, and the condorcet jury theorem. *American political science review*, 90(01):34–45, 1996.
- 7 Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. *Journal of economic theory*, 137(1):580–609, 2007.

⁵ For example, in Lemma 5 in [24], we can redefine sets N_1^e (N_2^e) to be the collection of subsets – instead of collection of players – in the support of μ^* that cover (do not cover) an edge e in the label cover instance. Along these lines, instead of bounding the cardinalities of N_1^e and N_2^e (which are denoted by n_1^e and n_2^e in [24]), we can bound the measures $\sum_{S \in N_1^e} \mu_S^*$ and $\sum_{S \in N_2^e} \mu_S^*$. Among other things, these changes allow us to use Jensen’s inequality and bound Δ_e as specified in Lemma 5 by [24].

- 8 Fernando Branco, Monic Sun, and J Miguel Villas-Boas. Too much information? information provision and search costs. *Marketing Science*, 2015.
- 9 Peter Bro Miltersen and Or Sheffet. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 234–247. ACM, 2012.
- 10 Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 2011.
- 11 Jimmy Chan, Fei Li, and Yun Wang. Discriminatory persuasion: How to convince a group. *Available at SSRN*, 2015.
- 12 Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1982.
- 13 Shaddin Dughmi. Submodular functions: Extensions, distributions, and algorithms. a survey. *arXiv preprint arXiv:0912.0322*, 2009.
- 14 Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. Constrained signaling in auction design. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1341–1357. Society for Industrial and Applied Mathematics, 2014.
- 15 Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. *STOC 2016*, 2016.
- 16 Yuval Emek, Michal Feldman, Iftah Gamzu, Renato PaesLeme, and Moshe Tennenholtz. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation*, 2(2):5, 2014.
- 17 Uriel Feige, Magnús M Halldórsson, Guy Kortsarz, and Aravind Srinivasan. Approximating the domatic number. *SIAM Journal on computing*, 2002.
- 18 Matthew Gentzkow and Emir Kamenica. Bayesian persuasion. *American Economic Review*, 2011.
- 19 Matthew Gentzkow and Emir Kamenica. Competition in persuasion. Technical report, National Bureau of Economic Research, 2011.
- 20 Matthew Gentzkow and Emir Kamenica. Costly persuasion. *The American Economic Review*, 2014.
- 21 Itay Goldstein and Yaron Leitner. Stress tests and information disclosure. Technical report, FRB of Philadelphia Working Paper, 2015.
- 22 Mingyu Guo and Argyrios Deligkas. Revenue maximization via hiding item attributes. *arXiv preprint arXiv:1302.5332*, 2013.
- 23 David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, 2003.
- 24 Subhash Khot, Richard J Lipton, Evangelos Markakis, and Aranyak Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. In *Internet and Network Economics*. Springer, 2005.
- 25 Kazuo Murota. Convexity and steinitz’s exchange property. In *Integer Programming and Combinatorial Optimization*, pages 260–274. Springer, 1996.
- 26 Kazuo Murota. *Discrete convex analysis*. SIAM, 2003.
- 27 Keith E Schnakenberg. Expert advice to a voting body. *Journal of Economic Theory*, 2015.
- 28 Wataru Tamura. Optimal monetary policy and transparency under informational frictions. *Available at SSRN 2191900*, 2014.
- 29 Ina A Taneva. Information design. Technical report, University of Edinburgh, 2015.
- 30 Jan Vondrák. *Submodularity in combinatorial optimization*. PhD thesis, PhD thesis, Charles University, Prague, Czech Republic, 2007.
- 31 Yun Wang. Bayesian persuasion with multiple receivers. *Available at SSRN 2625399*, 2013.