

Bayesian Persuasion with Externalities: Exploiting Agent Types

Jonathan Shaki¹, Jiarui Gan², Sarit Kraus¹

¹Bar-Ilan University

²University of Oxford

jonath.shaki@gmail.com, jiarui.gan@cs.ox.ac.uk, sarit@cs.biu.ac.il

Abstract

We study a Bayesian persuasion problem with externalities. In this model, a principal sends signals to inform multiple agents about the state of the world. Simultaneously, due to the existence of externalities in the agents' utilities, the principal also acts as a correlation device to correlate the agents' actions. We consider the setting where the agents are categorized into a small number of types. Agents of the same type share identical utility functions and are treated equitably in the utility functions of both other agents and the principal, while taking actions independently. We study the problem of computing optimal signaling strategies for the principal, under three different types of signaling channels: public, private, and semi-private. Our results include revelation-principle-style characterizations of optimal signaling strategies, linear programming formulations, and analysis of in/tractability of the optimization problems. It is demonstrated that when the maximum number of deviating agents is bounded by a constant, our LP-based formulations compute optimal signaling strategies in polynomial time. Otherwise, the problems are NP-hard.

1 Introduction

Bayesian persuasion, introduced by (Kamenica and Gentzkow 2011), is a framework where a persuader (the principal) seeks to design a signaling scheme that conveys information to a decision-maker (the agent) to influence their choices in the principal's desired direction (Kamenica 2019; Fujii and Sakaue 2022; Celli, Coniglio, and Gatti 2020). The computational analysis of multi-agent Bayesian persuasion has been studied extensively in recent years, considering various domains, such as congestion games (Vasserman, Feldman, and Hassidim 2015; Castiglioni et al. 2021; Bhaskar et al. 2016; Azaria et al. 2014; Das, Kamenica, and Mirka 2017; Griesbach et al. 2022), auctions (Li and Das 2019; Bacchiocchi et al. 2022), security (Rabinovich et al. 2015; Zhou et al. 2023), and advertising (Kumar et al. 2023). In this paper, we consider persuading multiple agents in the presence of *externalities* in the agents' utilities—that is, the utility of an agent may depend on other agents' actions as well, instead of just their own—and we consider a fairly general model that is not tied to specific applications (e.g., congestion games).

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Persuading multiple agents in the presence of externalities is computationally challenging for several reasons. First, even when there are no externalities, public persuasion (i.e., when signals are sent through a public channel visible to every agent) is known to be intractable. It is shown that no polynomial time approximation scheme (PTAS) exists for computing optimal public persuasion policies unless $P = NP$ (Bhaskar et al. 2016; Rubinstein 2017; Dughmi 2019). When introducing externalities, the problem does not become computationally easier, as it is a strict generalization of the case without externalities. For public persuasion, the above hardness result holds automatically, whereas for private persuasion, while the case without externalities can be trivially solved as independent single-agent persuasion problems against individual agents, this is not possible when externalities are present. In the presence of the externalities, the actions of the agents need to be coordinated, too. Indeed, coordination is necessary even when the principal does not possess any private information about the hidden state, whereby the goal of the principal is essentially to induce a correlated equilibrium among the agents that is optimal for the principal. Hence, the problem we address is essentially a combination of optimal persuasion and coordination.

Given the computational barriers, we consider a more amenable but fairly natural special case of the model where the agents can be classified into a small number of *types*. Similarly to other works in game theory (Shrot, Aumann, and Kraus 2010; Lovejoy 2006), agents of the same type share identical utility function and are treated equitably in the utility functions of the other players (however, they still act independently and need not take the same action). Hence, the agents' joint action can be succinctly represented as a vector that describes, for each agent type T and action a , how many agents of type T are performing action a . Succinct representations are necessary for many applications as they avoid explicitly defining the utilities of possible joint actions, which grow exponentially with the number of agents. From a computational perspective, however, this means that designing an efficient algorithm, which runs in time polynomial in the size of the problem representation, becomes harder. Indeed, it has been shown that correlated equilibria are hard to optimize in many succinctly represented games (Papadimitriou and Roughgarden 2005; Barman and Ligett 2015).

Contributions and Technical Novelties

We make the following contributions in this paper.

- We show that the classical revelation principle does not hold if agents can deviate jointly. We provide alternative revelation-principle-style characterizations for our model. Speaking informally, the new characterizations show that the principal can restrict herself to sending the agents the actions they are recommended to play, along with a blocking profile, which encodes necessary information that explains, for every possible deviation, why the deviating agents will not benefit from this deviation.
- Based on the characterizations, we present polynomial-time algorithms to compute optimal policies of the principal when joint deviations of the agents are restricted among a constant number of agents. In the case without this restriction, we show that computing optimal policies is intractable.
- We consider both public and private persuasion as the two most commonly studied forms of persuasion in the literature. Additionally, we also introduce a semi-private interaction framework where, while the principal publicly recommend joint actions for the agents, additional private information is used to further reshape the agents' beliefs. This form of persuasion occurs naturally when agents can see or know about each other's actions, while the principal can communicate separately with each agent through private channels.

To obtain the above results the main challenges lie in finding ways to reduce the representations of several key element that would otherwise be exponential if not handled carefully. In the public and semi-private cases, we introduce *representative action vectors* to reduce the exponentially large space of the agents' joint actions to a polynomial space, without loss of optimality. This technique fails in the private case, as we show through a counterexample. Hence, we further introduce the concept of *lottery policy* by exploiting a symmetry in the agents' roles. Such lottery policies allow representative action vectors to be applicable for private persuasion, too. Besides the representations of joint actions, concise representations of information encoded in the principal's signals are also necessary. In the semi-private and private cases, different agents may form different posterior beliefs about the state of the world (even if they are of the same type and are recommended to perform the same action). Since our characterizations indicate that the principal's signals need to encode sufficient information that explains, to every agent, why it will not be beneficial to deviate, it appears that jointly the space of signals for us to consider grows exponentially with the number of agents. We nevertheless show a way to represent the explanations that reduce the size of the signal space to a polynomial, by defining the explanations for each deviation instead of each agent.

Motivating Examples

Beyond the theoretical interests, the problem of persuading multiple agents in the presence of externalities is motivated by many real-world scenarios. Consider for example a navigation app (the principal) that would like to minimize its

users' overall travel time. Each user is associated with an origin and a destination (which indicates their type) and needs to choose a route (their possible actions) between these two points. The traffic load in each route depends on the state of the world. There is a prior distribution over possible states, but on any given day, the app possesses information of the actual traffic load, while individual users only have a prior knowledge about the state.

As another example, consider authorities, such as the Food and Drug Administration (FDA) in the US, that are responsible for evaluating the efficacy of new drugs. The interaction between the FDA and a company that wants their new drug to be approved can be captured by the persuasion model (Wang 2013). In this case, the company (the principal) persuades the FDA committee members (the agents) to accept the drug. The drug could be effective or ineffective, and with side effects or without (i.e., four possible states). The company needs to decide on the clinical trial to carry out, and its results will inform the committee members about the effectiveness of the drug. A committee member can vote to approve the drug, disapprove it, or ask for additional trials, which are their actions in the model. Their utilities can depend on the actual efficacy of the drug (the state) and the impact of a right/wrong decision on their own reputation and the FDA's reputations. The weights the members put on their own and the FDA's overall reputations may differ depending on their seniority (see more details in Section 2). The decision is made by majority rule, so there are externalities in the members' utilities.

Related Work

Bayesian persuasion against a single agent is tractable (Kamenica 2019; Dughmi 2017), but the more general problem with multiple agents becomes harder. Hence, most work on multi-agent Bayesian persuasion has been done under special game structures. The constraints of each domain were used in order to find tractable solutions. For example, in the context of singleton congestion games, with both private and public signals, (Zhou, Nguyen, and Xu 2022) proved that efficient computation of optimal policy for the principal is achievable when the number of resources is held constant. (Xu 2020) explored public persuasion without externalities, featuring binary actions but with a larger number of possible states. The study established that the optimal public signaling policy can be efficiently computed in polynomial time for arbitrary principal utility functions, given a constant number of states. (Castiglioni, Celli, and Gatti 2023) showed that it is possible to compute bi-criteria approximations of optimal public signaling schemes in arbitrary persuasion problems but, again, without externalities in quasi-poly-time.

In terms of persuasion with externalities, it has been shown that even in Bayesian zero-sum games no polynomial time approximation scheme (PTAS) is available for finding optimal public persuasion policy unless $P = NP$ (Rubinstein 2017). To deal with the challenges, we introduce the concept of types. The concept of type has also appeared in many previous works. In mechanism design settings, a type means a possible type of utility function an agent may have. The

mechanism designer does not know the agent’s exact type but maintains a probabilistic belief about what it might be. For example, (Castiglioni, Marchesi, and Gatti 2022) considered a persuasion scenario where the principal elicits the agent’s type first. The agent reports their type to the principal’s mechanism, and the principal commits to a signaling policy according to the agent’s report. No externalities are assumed in their model. Indeed, in our model, we assume that the types of the agents are known to the principal in advance, similarly to the standard persuasion model. The concept of type is more of a structure to group agents into small categories to allow us to derive efficient algorithms.

Our work considers three different types of signaling channels, including semi-private persuasion, which lies in between public and private persuasion. Similar perspectives have been adopted in the work of (Babichenko et al. 2021), who introduced a model that smoothly transitions between public and private persuasion. More specifically, they introduced a multi-channel communication structure, where each agent observes a subset of the principal’s communication channels. The authors provided a complete characterization of the conditions under which one communication method outperforms another. They provided characterizations applicable to scenarios both with and without externalities among the agents. However, their computational studies focus on the case with: a constant number of worlds, no externalities and an additive utility function of the principal. For future work, extending our model to a multi-channel structure with externalities seems very promising. Our semi-private model is the first step in this direction.

2 The Problem

A persuasion problem is given by a tuple $\langle \Omega, \mu, N, A, \mathcal{T}, (u_T)_{T \in \mathcal{T} \cup \{0\}}, d \rangle$. Specifically: Ω is the set of the possible worlds. $\mu \in \Delta(\Omega)$ is the prior distribution of the worlds. $N = \{1, \dots, n\}$ contains the indices of the n agents involved. A is the set of actions available to the agents (w.l.o.g. it is the same for all agents). $\mathcal{T} \subseteq 2^N$ is a collection of disjoint subsets—referred to as *types*—of the agents, which defines a partition of N . Finally, d is the maximum number of agents that can deviate jointly from the principal’s recommendation of actions. When $d = 1$, only unilateral deviations are possible, and when $d > 1$, groups deviations are allowed.

We denote by $\mathcal{A} = A^n$ the set of all possible joint actions of the agents. Given a joint action $\mathbf{a} \in \mathcal{A}$, we call $\rho_{\mathbf{a}} \in \mathbb{N}_{>0}^{|\mathcal{T}| \times |A|}$ the *action profile* corresponding to \mathbf{a} , which anonymizes the agents’ identities in \mathbf{a} and encodes only the number of agents of certain types who perform certain actions: $\rho_{\mathbf{a}}(T, a) = |\{i \in T : a_i = a\}|$ for all $T \in \mathcal{T}$ and $a \in A$. The utility $u_T(a, \rho | \omega)$ of a type- T agent depends on the action a performed by the agent, the action profile ρ encoding the joint action of the agents, as well as the current world ω . Alternatively, given a joint action \mathbf{a} , we can write the utility of an agent $i \in T$ as:

$$u_i(\mathbf{a} | \omega) = u_T(a_i, \rho_{\mathbf{a}} | \omega).$$

Similarly, we let $u_0(\rho | \omega)$ be the principal’s utility for an

action profile ρ in world ω (the principal does not perform actions herself).

The Principal’s Policy

As in standard persuasion models, the principal has an information advantage over the agents: she is the only player who can observe the world. The principal reveals this private information to influence the agents’ decision-making (choices of actions). Additionally, in our model the principal also has the power to correlate the agents’ actions and serves as a correlation device. To perform both signaling and correlation, we consider policies of the form $\sigma : \Omega \rightarrow \Delta(\mathcal{A} \times \mathcal{G})$, which map the world observed by the principal to a meta-signal $s = (\mathbf{a}, \mathbf{g}) \in \mathcal{A} \times \mathcal{G}$. The joint action \mathbf{a} functions as a correlation signal and encodes the action a_i each agent is recommended to perform, while $\mathbf{g} = (g_i)_{i \in N}$ provides additional information about the world to influence the agents’ beliefs, with g_i sent to each agent i .

We consider both public and private persuasion. In public persuasion, the communication channel is public and the entirety of the meta-signal s is seen by every agent. In private persuasion, each agent i only observes the part of signal (a_i, g_i) that is sent to them. In general, we use s_i to denote the part of s observed by agent i . Hence, for every $i \in N$, we have

- $s_i = (\mathbf{a}, \mathbf{g})$ in public persuasion, and
- $s_i = (a_i, g_i)$ in private persuasion.

The policies used in these two cases are referred to as public and private policies, respectively.

Besides these two most commonly studied persuasion modes, we also consider an in-between mode, which we call *semi-private* persuasion. In semi-private persuasion, the joint action \mathbf{a} in the principal’s signal is sent publicly while \mathbf{g} remains private. Essentially, private persuasion becomes semi-private when the agents can see the actions to be performed by the other agents. Hence,

- $s_i = (\mathbf{a}, g_i)$ in semi-private persuasion.

For ease of description, we write $s = (s_i)_{i \in N} = (\mathbf{a}, \mathbf{g})$.

Agents’ Belief and Stability of the Policy

Upon receiving a meta-signal s_i , the agent performs a belief update according to Bayes’ rule and derives the following posterior probability for each world $\omega \in \Omega$, as well as each possible joint action $\mathbf{a} \in \mathcal{A}$:

$$\mathbb{P}(\mathbf{a}, \omega | s_i) = \frac{\mu(\omega) \cdot \sigma(s_i, \mathbf{a} | \omega)}{\sum_{\omega' \in \Omega} \mu(\omega') \cdot \sigma(s_i | \omega')},$$

where \mathbb{P} denotes the probability measure induced by the policy, and we use the notation:

$$\sigma(x | \omega) = \sum_{s' \in \mathcal{S}: s' \vdash x} \sigma(s' | \omega),$$

where $s' \vdash x$ means s' satisfies x , e.g., $s' = (\mathbf{a}', \mathbf{g}')$ $\vdash (s_i, \mathbf{g})$ means $s'_i = s_i$ and $\mathbf{g}' = \mathbf{g}$, and $s' \vdash s_i$ means $s'_i = s_i$. In the private case, the agents also maintain a probabilistic belief about the other agents’ actions.

The stability of a policy is concerned with whether some subset of the agents can simultaneously benefit from some joint deviation. We only consider *non-transferable* utility in this paper, and we do not consider further communication between the agents (which may further change their posterior beliefs). Under these conditions, the stability of a policy is defined as follows.

Definition 1 (Stable policy). *Given a policy $\sigma : \Omega \rightarrow S$, a signal $s = (\mathbf{a}, \mathbf{g}) \in S$ is unstable if there exists a set $N' \subseteq N$ containing at most d agents and an action vector $\mathbf{a}' = (a'_i)_{i \in N'}$, such that for all $i \in N'$:*

$$\sum_{\tilde{\mathbf{a}} \in \mathcal{A}, \omega \in \Omega} \mathbb{P}(\tilde{\mathbf{a}}, \omega | s_i) \cdot \left(u_i(\tilde{\mathbf{a}} \oplus \mathbf{a}' | \omega) - u_i(\tilde{\mathbf{a}} | \omega) \right) > 0, \quad (1)$$

where $\tilde{\mathbf{a}} \oplus \mathbf{a}'$ denotes the joint action resulting from replacing \tilde{a}_i with a'_i for all $i \in N'$. The signal is stable otherwise. The policy σ is stable if all signals in $\text{supp}(\sigma)$ are stable.¹

In the public and semi-private cases, since agents know the joint action deterministically, the instability constraint in the definition simplifies to

$$\sum_{\omega \in \Omega} \mathbb{P}(\omega | s_i) \cdot \left(u_i(\mathbf{a} \oplus \mathbf{a}' | \omega) - u_i(\mathbf{a} | \omega) \right) > 0 \quad (2)$$

We will sometimes abuse the notation slightly and write

$$u_i(\mathbf{a} | p) = \mathbb{E}_{\omega \sim p} u_i(\mathbf{a} | \omega)$$

for a posterior $p \in \Delta(\Omega)$. So, Eq. (2) further simplifies to $u_i(\mathbf{a} \oplus \mathbf{a}' | p_i) > u_i(\mathbf{a} | p_i)$ by letting $p_i = \mathbb{P}(\omega | s_i)$.

Hence, a stable policy incentivizes the agents to perform the recommended actions. It yields utility

$$u_0(\sigma) = \mathbb{E}_{\omega \sim \mu} \mathbb{E}_{(\mathbf{a}, \mathbf{g}) \sim \sigma(\cdot | \omega)} u_0(\mathbf{a} | \omega)$$

for the principal. Our goal is to find, among all stable policies, one that maximizes the principal's utility (or decide correctly that no stable policy exists).

The FDA Example

Before we present our results, we describe a concrete example to illustrate the model defined above. Consider a simplified version of the FDA example we mentioned in the introduction. A drug can be either *safe* or *unsafe*, and the FDA decides whether to *accept* or *reject* a new drug. The utility of the company (the principal) is 0 if their drug is rejected, 2 if their drug is accepted and safe, 1 if it is accepted but unsafe. The FDA uses the majority rule to make their decision, where each of their senior members has two votes and each junior member has one.

The states of the drugs, safe or unsafe, correspond to the two possible worlds ω_+ and ω_- , with priors $\mu(\omega_+) = 0.25$ and $\mu(\omega_-) = 0.75$. Suppose that there are 2 senior members ($T_1 = \{1, 2\}$) and 3 junior members ($T_2 = \{3, 4, 5\}$). The possible actions of the agents include voting for rejection (a^{rej}) and for acceptance (a^{acc}). For an action profile ρ , the outcome of the vote is denoted by $\text{acc}(\rho)$.

¹ $\text{supp}(\cdot)$ denotes the support set, i.e., for a policy $\sigma : \Omega \rightarrow S$, $\text{supp}(\sigma) = \{s \in S : \sigma(s) > 0\}$.

The senior members place high importance on the reputation of the FDA. Each of them gets a high penalty (-3) if s/he votes in favor of an unsafe drug and the drug is accepted according to the majority vote. A smaller penalty (-2) is incurred if the senior member votes against an unsafe drug but the drug is accepted. Similarly, utilities for the other possible cases are listed in the table below. The junior members, on the other hand, are more concerned with their own reputation than the FDA's reputation. They get a positive utility when warning against an unsafe drug, even though the drug was eventually accepted by the majority.

	ω_-	ω_+		ω_-	ω_+
$a^{\text{rej}}, \text{acc}(\rho)$	-2	-1	$a^{\text{rej}}, \text{acc}(\rho)$	2	-2
$a^{\text{acc}}, \text{acc}(\rho)$	-3	1	$a^{\text{acc}}, \text{acc}(\rho)$	-3	1
$\cdot, \neg\text{acc}(\rho)$	0	0	$\cdot, \neg\text{acc}(\rho)$	0	0

$u_{T_1}(a, \rho | \omega) \qquad \qquad \qquad u_{T_2}(a, \rho | \omega)$

Consider the action vector $\mathbf{a} = (a^{\text{rej}}, a^{\text{acc}}, a^{\text{acc}}, a^{\text{rej}}, a^{\text{acc}})$. The corresponding action profile $\rho_{\mathbf{a}}$ is the function such that $\rho_{\mathbf{a}}(T_1, a^{\text{rej}}) = \rho_{\mathbf{a}}(T_1, a^{\text{acc}}) = \rho_{\mathbf{a}}(T_2, a^{\text{rej}}) = 1$ and $\rho_{\mathbf{a}}(T_2, a^{\text{acc}}) = 2$. Since the majority votes for acceptance in this profile, the utility of the principal is $u_0(\rho_{\mathbf{a}} | \omega_+) = 2$ if the drug is safe; otherwise, $u_0(\rho_{\mathbf{a}} | \omega_-) = 1$.

3 Public Persuasion

To explain our algorithm for public persuasion, we first analyze the structure of optimal policies and prove a result similar to the *revelation principle* (Kamenica and Gentzkow 2011). In the case *without* externalities, the revelation principle states that it is without loss of optimality to consider incentive compatible (IC) and direct policies. Namely, the policy directly recommends actions to the agents and incentivizes them to perform these recommended actions. The additional signal \mathbf{g} in the meta-signal is unnecessary in this case. This elegant result simplifies the design of policies: in particular, it sets an upper bound to the size of the meta-signal space. However, it does not hold any more when there are externalities and multiple agents can deviate together (i.e., $d > 1$). We provide an example to illustrate this.

An Example

Suppose that there are two worlds ω_1 and ω_2 , with equal prior $\mu(\omega_1) = \mu(\omega_2) = 0.5$. There are only two agents and they are of the same type, i.e., $\mathcal{T} = \{T\}$ and $T = \{1, 2\}$. There are three actions available to them: $A = \{a_1, a_2, a_3\}$. We have $d = 2$, so the agents can deviate jointly.

The agents' utilities are -1 for all the action profiles except two special profiles ρ_1 and ρ_2 , where $\rho_1(a_1) = 2$ and $\rho_2(a_2) = \rho_2(a_3) = 1$ (we omit types since there is only one type). The agents' utilities for these two profiles are given in the tables below. Irrespective of the world, the principal's utility is 1 for ρ_1 and 0 for all other profiles.

Hence, in this example, the maximum possible expected utility of the principal is 1, which can be obtained only if the agents are incentivized to follow ρ_1 . It can be seen that this is achieved by the policy which always reveals the world to

	ω_1	ω_2
a_1	1	1
$u_T(a, \rho_1 \omega)$		

	ω_1	ω_2
a_2	0	10
a_3	10	0
$u_T(a, \rho_2 \omega)$		

the agents while recommending $\mathbf{a}^1 = (a_1, a_1)$ (which is the only action vector corresponding to ρ_1). This policy is also stable. Indeed, given the utility definition, it would only be beneficial for the agents to deviate to ρ_2 . However, in ρ_2 , one of them will have to receive a lower utility of 0 and hence block this joint deviation.

However, recommending \mathbf{a}^1 without using any additional information is not stable. Indeed, this makes \mathbf{a}^1 the only signal used in the policy, irrespective of the persuasion mode (public, semi-private, or private). The signal is uninformative as a result, and the agents' beliefs remain the same as the prior after receiving the recommendation. In this case they will always prefer to deviate to ρ_2 and obtain an expected utility of 5. Therefore, for all the persuasion modes considered, signaling only the recommended actions to the agents may only produce suboptimal policies.

It turns out that while the revelation principle breaks in its original form, a variant of it can be established to bound the size of the meta-signal space. We present a new revelation-principle-style characterization next.

Revelation Principle for Public Persuasion

We identify meta-signals that can be merged without changing the induced outcome to reduce the meta-signal space as much as possible. To this end, we introduce the *signatures* of meta-signals. We begin by defining *representatives* and *blocking profiles*.

We fix an arbitrary *representative set* $\bar{\mathcal{A}} \subseteq \mathcal{A}$, in which every profile $\rho \in P$ finds exactly one joint action $\bar{\mathbf{a}} \in \bar{\mathcal{A}}$ such that $\rho_{\bar{\mathbf{a}}} = \rho$. Namely, $\bar{\mathbf{a}}$ is representative of the joint actions whose action profile is ρ . Since the representative set defines a bijection between the action profiles and the representative actions, we will use a representative action $\bar{\mathbf{a}}$ and its corresponding action profile $\rho_{\bar{\mathbf{a}}}$ interchangeably.

Definition 2 (Representative). *A joint action $\bar{\mathbf{a}}$ is representative of another joint action \mathbf{a} if $\bar{\mathbf{a}} \in \bar{\mathcal{A}}$ and $\rho_{\bar{\mathbf{a}}} = \rho_{\mathbf{a}}$.*

In an action profile, we refer to the tuple $(T, a) \in \mathcal{T} \times A$ as the *subtype* of the agents who are of type T and perform action a . Let D_* denote the set of all possible deviations:

$$D_* = \left\{ \delta : \mathcal{T} \times A^2 \rightarrow \mathbb{Z}_{\geq 0} \mid \begin{array}{l} 1 \leq \sum_{T, a, a'} \delta(T, a, a') \leq d, \\ \delta(T, a, a) = 0 \forall a, T \end{array} \right\}.$$

Each deviation δ specifies the number $\delta(T, a, a')$ of agents of subtype (T, a) who deviate to action a' . The total number of deviating agents is bounded by d by assumption, while there is no need to deviate from an action to itself. Next, let D_ρ contain deviations that are feasible from ρ :

$$D_\rho = \left\{ \delta \in D_* \mid \sum_{a' \in A} \delta(T, a, a') \leq \rho(T, a) \forall a, T \right\}, \quad (3)$$

i.e., the number of type- T agents who deviate from action a must be consistent with the number of type- T agents who originally perform this action in ρ .

The *blocking profile* of a public signal is then defined as follows. It describes how each possible deviation is blocked by agents of certain subtypes.

Definition 3 (Blocking profile). *The blocking profile of a meta-signal $s = (\mathbf{a}, \mathbf{g})$ is the set $\beta =$*

$$\left\{ (\delta, T, a, a') \in D_\rho \times \mathcal{T} \times A^2 \mid \begin{array}{l} \delta(T, a, a') > 0, \\ u_T(a, \rho_{\mathbf{a}} | p_s) \geq u_T(a', \rho_{\mathbf{a}} \oplus \delta | p_s) \end{array} \right\},$$

where $p_s = \mathbb{P}(\cdot | s)$ denotes the posterior induced by s , and $\rho_{\mathbf{a}} \oplus \delta$ denotes the action profile resulting from applying deviation δ to $\rho_{\mathbf{a}}$.

Namely, $(\delta, T, a, a') \in \beta$ means that the agents of subtype (T, a) are unwilling to deviate to playing a , hence blocking δ . It is straightforward that if a meta-signal is stable, then its blocking profile covers every $\delta \in D_\rho$, i.e., it contains at least one element involving δ . We denote by B_ρ^{pub} the set of subsets of $D_\rho \times \mathcal{T} \times A^2$ that covers every $\delta \in D_\rho$ in the public case. In other words, $B_{\rho_{\mathbf{a}}}^{\text{pub}}$ consists of all possible blocking profiles of stable meta-signals $s = (\mathbf{a}, \mathbf{g})$. For simplicity, we will also write $B_{\rho_{\mathbf{a}}}^{\text{pub}}$ as $B_{\bar{\mathbf{a}}}^{\text{pub}}$.

Combining a representative and a blocking profile gives the signature of a meta-signal.

Definition 4 (Signature). *The public signature of a meta-signal $s = (\mathbf{a}, \mathbf{g})$, denoted $\phi^{\text{pub}}(s)$, is the tuple $(\bar{\mathbf{a}}, \beta)$ where $\bar{\mathbf{a}}$ is the representative of \mathbf{a} and β is the blocking profile of s .*

Intuitively, meta-signals with the same signature can be merged without affecting the outcome induced by the policy. The support of the resulting policy is then bounded by the number of distinct signatures there are. This leads to the following revelation-principle-style characterization for public persuasion.

Theorem 3.1. *For any stable public policy $\sigma : \Omega \rightarrow \Delta(\mathcal{A} \times G)$, there exists a stable public policy $\bar{\sigma} : \Omega \rightarrow \Delta(C)$ that yields as much utility for the principal as σ does, where $C = \{(\bar{\mathbf{a}}, \beta) : \bar{\mathbf{a}} \in \bar{\mathcal{A}}, \beta \in B_{\bar{\mathbf{a}}}^{\text{pub}}\}$. Moreover, the signature of each meta-signal $c \in C$ is exactly c .²*

The above result indicates that it suffices to encode in each meta-signal the recommended actions, as well as a blocking profile to explain to the agents why it would not be beneficial for them to deviate in any possible way. The signature of every meta-signal used in the policy is exactly the meta-signal itself. Based on the theorem, we design an algorithm to compute optimal public policies next.

Computing an Optimal Public Policy

We use the following LP to compute an optimal public policy, with variables $\sigma(s_{\bar{\mathbf{a}}, \beta} | \omega)$ for $\bar{\mathbf{a}} \in \bar{\mathcal{A}}, \beta \in B_{\bar{\mathbf{a}}}^{\text{pub}}, \omega \in \Omega$.

$$\max_{\omega \in \Omega} \sum_{\omega \in \Omega} \mu(\omega) \sum_{\bar{\mathbf{a}} \in \bar{\mathcal{A}}} \sum_{\beta \in B_{\bar{\mathbf{a}}}^{\text{pub}}} \sigma(s_{\bar{\mathbf{a}}, \beta} | \omega) \cdot u_0(\bar{\mathbf{a}} | \omega), \quad (4)$$

²Omitted proofs are in the appendix: arxiv.org/abs/2412.12859

subject to the following constraint for every $\bar{\mathbf{a}} \in \bar{\mathcal{A}}$, $\beta \in B_{\bar{\mathbf{a}}}^{\text{pub}}$, and every tuple $(\delta, T, a, a') \in \beta$:

$$\sum_{\omega \in \Omega} \mu(\omega) \cdot \sigma(s_{\bar{\mathbf{a}}, \beta} | \omega) \cdot u_T(a', \rho_{\bar{\mathbf{a}}} \oplus \delta | \omega) \leq \sum_{\omega \in \Omega} \mu(\omega) \cdot \sigma(s_{\bar{\mathbf{a}}, \beta} | \omega) \cdot u_T(a, \rho_{\bar{\mathbf{a}}} | \omega). \quad (5)$$

Additionally, for all $\omega \in \Omega$:

$$\sum_{\bar{\mathbf{a}} \in \bar{\mathcal{A}}} \sum_{\beta \in B_{\bar{\mathbf{a}}}^{\text{pub}}} \sigma(s_{\bar{\mathbf{a}}, \beta} | \omega) = 1 \quad (6)$$

$$\sigma(s_{\bar{\mathbf{a}}, \beta} | \omega) \geq 0 \quad \text{for all } \bar{\mathbf{a}} \in \bar{\mathcal{A}}, \beta \in B_{\bar{\mathbf{a}}}^{\text{pub}} \quad (7)$$

so that $\sigma(\cdot | \omega)$ is a valid distribution.

In other words, the variables encode a policy characterized by Theorem 3.1. Under the assumption that this policy is stable, the objective function Eq. (4) captures exactly the expected utility of the principal. Eq. (5) further ensures that the signature of each meta-signal is exactly itself, which follows by (the negation of) Eq. (2). Consequently, the policy is indeed stable since by definition B_{ρ}^{pub} only contains blocking profiles under which ρ is stable.

Time Complexity Given a constant d , and a constant number of types and actions in \mathcal{T} and A , the size of the above LP grows only polynomially in the problem size. For general d , however, the formulation can grow exponentially. Indeed, this is inevitable: when d is part of the input, computing an optimal public policy is NP-hard. The reduction also applies to the semi-private and private cases.

Theorem 3.2. *For constant d , $|\mathcal{T}|$, and $|A|$, an optimal public policy can be computed in polynomial time.*

Theorem 3.3. *When d is an input to the problem, computing an optimal policy is NP-hard for public, semi-private, and private persuasion.*

4 Semi-private Persuasion

We now consider the semi-private case. Similarly to the public case, a blocking profile should contain information about how each possible deviation is blocked. However, in the semi-private case we need to define the profiles separately for different agents because of divergent beliefs. This may result in an exponential growth of blocking profiles when the number of agents increases (i.e., each profile corresponds to a possible combination of n sets of deviations).

Representing Semi-Private Blocking Profiles

To address the above issue, we explore the minimum possible representation of a blocking profile. Key to our approach is Lemma 4.1, a generalization of Hall's theorem (Hall 1934). With this result, we can prove Lemma 4.2 to find a function γ that encodes sufficient information for us to derive a concise representation of blocking profiles.

Lemma 4.1. *Suppose that $B_1, \dots, B_m \subseteq \{1, \dots, \ell\}$, and $r_1, \dots, r_m \in \mathbb{R}$ are non-negative. There exist m disjoint sets $\bar{B}_1, \dots, \bar{B}_m$, $\bar{B}_i \subseteq B_i$ and $|\bar{B}_i| = r_i$ for all $i = 1, \dots, m$, if and only if it holds for every set $M \subseteq \{1, \dots, m\}$ that $|\bigcup_{i \in M} B_i| \geq \sum_{i \in M} r_i$.*

Lemma 4.2. *If a meta-signal $s = (\mathbf{a}, \mathbf{g})$ is stable, then there exists a function $\gamma : D_* \rightarrow 2^A \times 2^N$ such that for every δ , it holds for $(A', N') = \gamma(\delta)$ that:*

- 1) *agents in N' are all of the same subtype, say (T, a) ;*
- 2) *$\delta(T, a, a') > 0$, for all $a' \in A'$;*
- 3) *$\rho_{\mathbf{a}}(T, a) - |N'| < \sum_{a' \in A'} \delta(T, a, a')$; and*
- 4) *for every $a' \in A', i \in N'$,*

$$u_T(a, \rho_{\mathbf{a}} | p_i) \geq u_T(a', \rho_{\mathbf{a}} \oplus \delta | p_i), \quad (8)$$

where $p_i = \mathbb{P}(\cdot | s_i)$ denotes the posterior induced by s_i .

Intuitively, for a deviation to be successful, a matching that fulfills this deviation needs to be established between agents of certain subtypes and the actions they are willing to deviate to. A stable meta-signal prevents such matchings. Hence, for every possible deviation $\delta \in D_{\rho_{\mathbf{a}}}$, the set N' stated in the lemma contains a set of agents of subtype (T, a) who are unwilling to deviate to any of the actions in A' according to Eq. (8). According to the third condition in the lemma, N' is large enough, so that the remaining agents of this subtype are insufficient for fulfilling δ . Based on γ , we define the semi-private blocking profiles as follows.

Definition 5 (Semi-private blocking profile). *Let $s = (\mathbf{a}, \mathbf{g})$ be a stable meta-signal and γ be the function in Lemma 4.2.³ The semi-private blocking profile of each s_i is*

$$\beta_i = \{(\delta, A') \in D_{\rho_{\mathbf{a}}} \times 2^A \mid i \in N' \text{ for } (A', N') = \gamma(\delta)\}.$$

The semi-private blocking profile of s is $\beta = (\beta_i)_{i \in N}$.

Namely, β_i describes, for every blocked deviation δ , a set A' of actions disliked by agent i (as well as other agents in some N') that results in δ being blocked. This is the minimum information that “witnesses” the stability of s while ensuring that merging meta-signals with the same semi-private signature (Definition 6) does not change the outcome of the policy. While the sets N' are not listed in any β_i , it is implicitly defined through the joint profile β . This is key to avoiding the exponential growth of the profile space.

Definition 6 (Semi-private signature). *The semi-private signature of a meta-signal $s = (\mathbf{a}, \mathbf{g})$, denoted $\phi^{\text{sem}}(s)$, is the tuple $(\bar{\mathbf{a}}, \beta)$ where $\bar{\mathbf{a}}$ is a representative of \mathbf{a} , and $\beta = (\beta_i)_{i \in N}$ is the semi-private blocking profile of s . For each $s_i = (\mathbf{a}, g_i)$, the semi-private signature of s_i is defined as $\phi^{\text{sem}}(s_i) = \phi_i^{\text{sem}}(s) = (\bar{\mathbf{a}}, \beta_i)$.*

We can now prove a new revelation principle for the semi-private case. For every $\mathbf{a} \in \mathcal{A}$, we let $B_{\mathbf{a}}^{\text{sem}}$ be the set of blocking profiles defined by all possible meta-signals according to Definition 5, when the joint action in the meta-signal is fixed to \mathbf{a} (which also appear in the last two conditions about γ in Lemma 4.2).

Theorem 4.3. *For any stable semi-private policy $\sigma : \Omega \rightarrow \Delta(\mathcal{A} \times G)$, there exists a stable semi-private policy $\bar{\sigma} : \Omega \rightarrow \Delta(C)$ that yields as much utility for the principal as σ does, where $C = \{(\bar{\mathbf{a}}, \beta) : \bar{\mathbf{a}} \in \bar{\mathcal{A}}, \beta \in B_{\bar{\mathbf{a}}}^{\text{sem}}\}$. Moreover, the semi-private signature of each meta-signal $c \in C$ is c .*

³If there are multiple such functions, we choose one according to an arbitrary tie-breaking rule. The same applies to Definition 7.

It can be shown that when d , $|\mathcal{T}|$, and $|A|$ are constants, $B_{\mathbf{a}}^{\text{sem}}$ grows only polynomially with the size of the problem. Similarly to the public case, we obtain a polynomial-size LP formulation. This leads to a polynomial-time algorithm for semi-private persuasion.

Theorem 4.4. *For constant d , $|\mathcal{T}|$, and $|A|$, an optimal semi-private policy can be computed in polynomial time.*

5 Private Persuasion

In private persuasion, even the action recommendations are sent privately. This means that the agents maintain probabilistic beliefs about both the world and the actions of the other agents. We follow the same routine as the previous sections, deriving a revelation principle for the private case and then an LP formulation based on the result. A crucial difference in the private case is that it is no longer w.l.o.g. to restrict the policy to representative action vectors (Proposition 5.1).⁴ To solve the problem requires a novel concept called the *lottery policy*.

Proposition 5.1. *An optimal private policy supported only on representative action vectors need not exist.*

Before we define lottery policies, we first extend the notions of blocking profile and signature to the private case. We can replicate Lemma 4.2 and prove the following lemma, with slight changes in the domain of γ and the stability constraint in the last condition. In the private case, the agents' expected utilities are functions of their beliefs over all possible joint actions.

Lemma 5.2. *If a meta-signal $s = (\mathbf{a}, \mathbf{g})$ is stable, then there exists a function $\gamma : D_* \rightarrow 2^A \times 2^{N'}$ such that for every δ , it holds for $(A', N') = \gamma(\delta)$ that:*

- 1) *agents in N' are all of the same subtype, say (T, a) ;*
- 2) *$\delta(T, a, a') > 0$, for all $a' \in A'$;*
- 3) *$\rho_{\mathbf{a}}(T, a) - |N'| < \sum_{a' \in A'} \delta(T, a, a')$; and*
- 4) *for every $a' \in A'$, $i \in N'$,*

$$\sum_{\omega \in \Omega} \sum_{\tilde{\mathbf{a}} \in \mathcal{A}: \tilde{a}_i = a} \mathbb{P}(\tilde{\mathbf{a}}, \omega | s_i) \cdot u_T(a, \rho_{\tilde{\mathbf{a}}} | \omega) \geq \sum_{\omega \in \Omega} \sum_{\tilde{\mathbf{a}} \in \mathcal{A}: \tilde{a}_i = a} \mathbb{P}(\tilde{\mathbf{a}}, \omega | s_i) \cdot u_T(a', \rho_{\tilde{\mathbf{a}}} \oplus \delta | \omega).$$

This leads to the definitions of the private blocking profile and the private signature.

Definition 7 (Private blocking profile). *Let $s = (\mathbf{a}, \mathbf{g})$ be a stable meta-signal. The private blocking profile of each s_i is*

$$\beta_i = \{(\delta, A') \in D_* \times 2^A \mid i \in N' \text{ for } (A', N') = \gamma(\delta)\},$$

where γ is the function satisfying the conditions in Lemma 5.2.

Definition 8 (Private signature). *The private signature of a meta-signal $s = (\mathbf{a}, \mathbf{g})$, denoted $\phi^{\text{prv}}(s)$, is (\mathbf{a}, β) where $\beta = (\beta_i)_{i \in N}$ is the private blocking profile of s . For each $s_i = (a_i, g_i)$, the private signature of s_i is defined as $\phi_i^{\text{prv}}(s_i) = \phi_i^{\text{prv}}(s) = (\mathbf{a}_i, \beta_i)$.*

⁴This also implies that private persuasion can achieve strictly higher utilities than semi-private and public persuasion, since in the latter two it is w.l.o.g. to consider representative action vectors.

Lottery Policy and Revelation Principle

We now introduce lottery policies. Lottery policies utilize a symmetry in the agents' roles to avoid explicit representation of the original policies. Given a private policy σ , a lottery policy $\lambda(\sigma)$ "lotterizes" σ by uniformly randomizing the signals of σ among agents of the same type.

Definition 9 (Lottery policy). *Let \mathcal{M} be the set of all permutations $m : N \rightarrow N$ such that $m(i) \in T$ if and only if $i \in T$. When σ signals each meta-signal (\mathbf{a}, \mathbf{g}) , the lottery policy $\lambda(\sigma)$ draws $m \sim \text{Uniform}(\mathcal{M})$ and signals $(\mathbf{a}', \mathbf{g}')$, where $a'_i = a_{m(i)}$, $g'_i = g_{m(i)}$.*

For example, consider an instance with two types of agents: $T_1 = \{1, 2\}$ and $T_2 = \{3\}$; and a private policy σ such that:

$$\begin{aligned} \sigma((a_1, a_2, a_3), (g_1, g_1, g_1) | \omega_1) &= 0.4 \\ \text{and } \sigma((a_1, a_2, a_1), (g_1, g_2, g_2) | \omega_1) &= 0.6. \end{aligned}$$

The corresponding lottery policy gives:

$$\begin{aligned} \lambda(\sigma)((a_1, a_2, a_3), (g_1, g_1, g_1) | \omega_1) &= 0.2, \\ \lambda(\sigma)((a_2, a_1, a_3), (g_1, g_1, g_1) | \omega_1) &= 0.2, \\ \lambda(\sigma)((a_1, a_2, a_1), (g_1, g_2, g_2) | \omega_1) &= 0.3, \\ \text{and } \lambda(\sigma)((a_2, a_1, a_1), (g_2, g_1, g_2) | \omega_1) &= 0.3. \end{aligned}$$

Using lottery policies, we can derive a concise characterization of optimal private policies, as stated in Theorem 5.3. We let $B_{\mathbf{a}}^{\text{prv}}$ be the set of blocking profiles induced by all possible meta-signals according to Definition 8.

Theorem 5.3. *There exists a private policy $\sigma : \Omega \rightarrow \Delta(C)$ where $C = \{(\tilde{\mathbf{a}}, \beta) : \tilde{\mathbf{a}} \in \tilde{\mathcal{A}}, \beta \in B_{\tilde{\mathbf{a}}}^{\text{prv}}\}$, such that $\lambda(\sigma)$ is an optimal private policy. Moreover, the signature of each meta-signal $c \in C$ is exactly the c .*

Similarly to the semi-private case, when d , $|\mathcal{T}|$, and $|A|$ are constants, $B_{\mathbf{a}}^{\text{prv}}$ grows only polynomially with the size of the problem, so we can derive a polynomial-size LP formulation, which gives a polynomial-time algorithm for private persuasion.

Theorem 5.4. *For constant d , $|\mathcal{T}|$, and $|A|$, an optimal private policy can be computed in polynomial time.*

6 Conclusion

We studied the problem of Bayesian persuasion with externalities. We showed that when multiple agents can coordinate their deviation the classical revelation principle does not hold, and presented alternative characterizations of optimal policies for public, private, and semi-private persuasion. The concept of agent types is introduced as a succinct representation of the problem. We presented polynomial time algorithms when only a constant number of agents can jointly deviate, and we proved that the problem is hard otherwise.

Acknowledgments

This research has been partially supported by the Israel Science Foundation under grant 2544/24 and the Israel Ministry of Innovation, Science & Technology grant 1001818511.

References

- Azaria, A.; Rabinovich, Z.; Goldman, C. V.; and Kraus, S. 2014. Strategic information disclosure to people with multiple alternatives. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4): 1–21.
- Babichenko, Y.; Talgam-Cohen, I.; Xu, H.; and Zabarnyi, K. 2021. Multi-channel bayesian persuasion. *arXiv preprint arXiv:2111.09789*.
- Bacchiocchi, F.; Castiglioni, M.; Marchesi, A.; Romano, G.; and Gatti, N. 2022. Public Signaling in Bayesian Ad Auctions. In *IJCAI 2022*.
- Barman, S.; and Ligett, K. 2015. Finding any nontrivial coarse correlated equilibrium is hard. *ACM SIGecom Exchanges*, 14(1): 76–79.
- Bhaskar, U.; Cheng, Y.; Ko, Y. K.; and Swamy, C. 2016. Hardness results for signaling in bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 479–496.
- Castiglioni, M.; Celli, A.; and Gatti, N. 2023. Public bayesian persuasion: being almost optimal and almost persuasive. *Algorithmica*, 1–37.
- Castiglioni, M.; Celli, A.; Marchesi, A.; and Gatti, N. 2021. Signaling in bayesian network congestion games: the subtle power of symmetry. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 5252–5259.
- Castiglioni, M.; Marchesi, A.; and Gatti, N. 2022. Bayesian Persuasion Meets Mechanism Design: Going Beyond Intractability with Type Reporting. In *Proceedings of the 21st International Conference on Autonomous Agents and Multi-agent Systems*, 226–234.
- Celli, A.; Coniglio, S.; and Gatti, N. 2020. Private Bayesian persuasion with sequential games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 1886–1893.
- Das, S.; Kamenica, E.; and Mirka, R. 2017. Reducing congestion through information design. In *2017 55th annual allerton conference on communication, control, and computing (allerton)*, 1279–1284. IEEE.
- Dughmi, S. 2017. Algorithmic information structure design: a survey. *ACM SIGecom Exchanges*, 15(2): 2–24.
- Dughmi, S. 2019. On the hardness of designing public signals. *Games and Economic Behavior*, 118: 609–625.
- Fujii, K.; and Sakaue, S. 2022. Algorithmic Bayesian persuasion with combinatorial actions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5016–5024.
- Griesbach, S. M.; Hoefer, M.; Klimm, M.; and Koglin, T. 2022. Public signals in network congestion games. *arXiv preprint arXiv:2205.09823*.
- Hall, P. 1934. On Representation of Subsets. *J. of London Math. Soc.*, 10: 26–30.
- Kamenica, E. 2019. Bayesian persuasion and information design. *Annual Review of Economics*, 11: 249–272.
- Kamenica, E.; and Gentzkow, M. 2011. Bayesian persuasion. *American Economic Review*, 101(6): 2590–2615.
- Kumar, Y.; Jha, R.; Gupta, A.; Aggarwal, M.; Garg, A.; Malyan, T.; Bhardwaj, A.; Shah, R. R.; Krishnamurthy, B.; and Chen, C. 2023. Persuasion strategies in advertisements. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 57–66.
- Li, Z.; and Das, S. 2019. Revenue enhancement via asymmetric signaling in interdependent-value auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2093–2100.
- Lovejoy, W. S. 2006. Optimal mechanisms with finite agent types. *Management Science*, 52(5): 788–803.
- Papadimitriou, C. H.; and Roughgarden, T. 2005. Computing equilibria in multi-player games. In *SODA*, volume 5, 82–91. Citeseer.
- Rabinovich, Z.; Jiang, A. X.; Jain, M.; and Xu, H. 2015. Information disclosure as a means to security. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 645–653.
- Rubinstein, A. 2017. Honest Signaling in Zero-Sum Games Is Hard, and Lying Is Even Harder. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*.
- Shrot, T.; Aumann, Y.; and Kraus, S. 2010. On agent types in coalition formation problems. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-agent Systems: volume 1-Volume 1*, 757–764.
- Vasserman, S.; Feldman, M.; and Hassidim, A. 2015. Implementing the wisdom of waze. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Wang, Y. 2013. Bayesian persuasion with multiple receivers. Available at SSRN 2625399.
- Xu, H. 2020. On the tractability of public persuasion with no externalities. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2708–2727. SIAM.
- Zhou, C.; Nguyen, T. H.; and Xu, H. 2022. Algorithmic information design in multi-player games: Possibilities and limits in singleton congestion. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 869–869.
- Zhou, C.; Spivey, A.; Xu, H.; and Nguyen, T. H. 2023. Information Design for Multiple Interdependent Defenders: Work Less, Pay Off More. *Games*, 14(1): 12.