

Mediated Cheap Talk Design

Itai Arieli, Ivan Geffner, Moshe Tennenholtz*

Technion - Israel Institute of Technology

iarieli@technion.ac.il, ieg8@cornell.edu, moshet@technion.ac.il

Abstract

We study an information design problem with two informed senders and a receiver in which, in contrast to traditional Bayesian persuasion settings, senders do not have commitment power. In our setting, a trusted mediator/platform gathers data from the senders and recommends the receiver which action to play. We characterize the set of implementable action distributions that can be obtained in equilibrium, and provide an $O(n \log n)$ algorithm (where n is the number of states) that computes the optimal equilibrium for the senders. Additionally, we show that the optimal equilibrium for the receiver can be obtained by a simple revelation mechanism.

Introduction

The extensive literature on information design and Bayesian persuasion studies optimal information revelation policies for the informed player. The two leading models of information revelation are cheap talk (Crawford and Sobel 1982) and Bayesian persuasion (Kamenica and Gentzkow 2011). The main distinction between these models is the underlying assumption that in the Bayesian persuasion models the sender has commitment power in the way she discloses the information.

Commitment power in the Bayesian persuasion model is crucial (see, e.g., (Conitzer and Sandholm 2006)) and, while it may hold in some real-world settings,¹ it is often considered strong. Another fundamental assumption in Bayesian persuasion models is that the informed player is also the one that designs the information revelation policy. In practice, however, information revelation can be determined by other external or legal constraints. For example, information revealed to a potential customer about a product is determined by the commerce platform based on information submitted by different suppliers.

On the other hand, in the cheap talk model there is no commitment power. However, in many cases, this lack of

commitment leads to a lack of expressive power from the sender and, as a result, may induce highly inefficient outcomes. For example, consider the classic market for lemons example in (Akerlof 1978), in which a marketer tries to sell a product to a customer that plays the role of the receiver. The product can be of either good quality or bad quality with equal probability. The customer would like to buy the product only if he believes that it is of good quality with a probability of at least $\frac{3}{4}$, and the sender always prefers that the product is bought. In this case, under a cheap talk equilibrium, the sender has no way to signal credible information to the receiver and the receiver never buys the product. This is a typical situation that arises in marketplaces that match sellers with buyers, or advertisers with consumers.

In our model there is a finite state space of size n , two informed players (senders), and an additional uninformed player (the receiver) that determines the outcome of the game by playing a binary action from the set $A := \{0, 1\}$ (this could represent buying a product or not, passing a law or not, etc.). The utility of each player is determined by the game state and by the action played by the receiver, and the incentives of the senders may not necessarily be aligned (e.g., senders can be a car seller and a technician that tested the car, or two parties who studied the monetary value of a law, or two suppliers of a product, etc.). The state of the game is drawn from a prior distribution that is common knowledge among the players, but only the senders know its exact value. Thus, the senders' purpose is to reveal information to the receiver in such a way that the receiver plays the action that benefits them the most. Since the senders have no commitment power we are interested in cheap talk equilibria, in which it is never in the interest of the senders to be dishonest, and it is always in the interest of the receiver to play the action suggested by the protocol. As we show in this work, the existence of a second informed sender dramatically enriches the set of cheap talk equilibria that can be obtained.

We consider a *mediated cheap talk* setting of communication between the senders and the receiver. In this setting, the senders communicate with a trusted mediator, and as a function of the two messages that the mediator receives, he sends an action recommendation (possibly at random) to the receiver. Our first result provides a characterization of the truthful equilibria that are implementable in the mediator

*The work by Ivan Geffner and Moshe Tennenholtz was supported by funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 740435).

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The leading motivating example of (Kamenica and Gentzkow 2011) is of a prosecutor persuading a judge.

setting, in which it is always beneficial for the senders to report truthful information and for the receiver to play whatever is suggested by the mediator. We then analyze the case where the two senders have aligned preferences and provide an algorithm with $O(n \log n)$ steps to calculate the best equilibrium outcome and payoff for the senders. We later extend this algorithm to find the optimal outcome for one of the senders whenever the senders have different incentives. Finally, we study the best equilibrium for the receiver and show that the optimal revelation policy lies within a finite set of mechanisms.

A major motivation for our work is data-driven decision making. Recommendation systems and classifiers are at the heart of many systems and determine the offering for or grouping of users based on data collected and provided by data sources. While in the early days these systems were based on data aggregated from their users and from previous interactions with them, the explosion of data facilitated professional data aggregation, and systems are designed to work with external data sources. Needless to say, data sources may be strategic and may attempt to influence the system’s decisions. In abstract terms, the system acts as a mediator aiming in implementing a policy, that is a mapping from a state (e.g., type of user) to an action (recommendation, group assignment) based on messages received by data sources that can access the state. In general there may be several data sources, each of which has access to the required data but they have different preferences regarding the policy to be implemented by the system. Therefore, the main theoretic question is to know which policies can be implemented in the strategic game between the data providers. Given that knowledge, we can tackle the question of what would be a mechanism that maps data sources’ messages to actions that are optimal when the system aims to implement a particular policy. Our work provides rigorous answers to the above question. This is complementary to work exploiting commitment power in data-intensive tasks, such as segmentation (e.g. (Emek et al. 2012)) and incentive-compatible exploration and exploitation (Kremer, Mansour, and Perry 2013; Bahar, Smorodinsky, and Tennenholtz 2019); in our setting, information providers do not have commitment power.

Finally, as shown in (Abraham et al. 2006), the assumption of communicating with a trusted mediator in most cases can be replaced with the assumption that both the senders and the receiver can communicate via private authenticated channels. This is true as long as (a) there exists a punishment strategy for the senders/receiver, or (b) we allow an arbitrarily small probability of error. This means that if the receiver or the senders can punish other participants when they are caught deviating (e.g., by quitting the game if all outcomes give positive utilities to the players), the same equilibria that can be achieved in the mediator setting can also be achieved in a cheap talk equilibrium without a mediator. If there is no such punishment strategy, the sets of equilibria in both settings might not be equal, but for any equilibrium in the mediator setting, there exist equilibria in the unmediated setting that are arbitrarily close.

Related Literature

The literature on information design is too vast to address all the related work. We will therefore mention some key related papers. The work by (Krishna and Morgan 2001) considers a setting that is similar to the one considered by (Crawford and Sobel 1982), where a real interval represents the set of states and actions. In this setting the receiver’s and the senders’ utilities are *biased* by some factor that affects their incentives and utility. In (Krishna and Morgan 2001) there are two informed senders that reveal information sequentially to the receiver. They consider the best receiver equilibrium and show that, when both senders are *biased* in the same direction, it is never beneficial to consult both of them. By contrast, when senders are biased in opposite directions, it is always beneficial to consult both of them. Our setting is different than theirs as we consider a finite state space and a binary action set for the receiver. In addition, our focus is the best equilibrium for either both or one of the senders.

In another work (Salamanca 2021) characterizes the optimal mediation for the sender in a sender-receiver game. Relatedly, (Lipnowski and Ravid 2020), and (Kamenica and Gentzkow 2011) provide a geometric characterization of the best cheap talk equilibrium for the sender under the assumption that the sender’s utility is state-independent. In (Gan et al. 2022) and (Fujii and Sakaue 2022), the authors study the complexity of finding equilibria in sequential decision-making settings and in settings where the receiver’s actions are specified by combinatorial constraints, respectively.

(Kamenica and Gentzkow 2017) consider a setting with two senders in a Bayesian persuasion model. The two senders, as in the standard Bayesian persuasion model, have commitment power and they compete over information revelation. The authors characterize the equilibrium outcomes in this setting.

Integrating mediators into a strategic setting is common in many game-theoretical works (e.g., (Aumann 1987), (Morgan and Morrison 1999)). In more recent work, (Kosenko 2018) and (Arieli, Babichenko, and Sandomirskiy 2022) study mediators in a Bayesian persuasion model. In these works the mediators are strategic players that may affect information revelation to the receiver. By contrast, in this work we remain agnostic to the incentives of the mediator and, as in (Aumann 1987), the mediator only serves as a correlation device.

The Model

Throughout the rest of the paper we will focus on the mediator setting since the mechanisms involved are much more simple than those required for the cheap talk setting, since the latter require non-trivial distributed computing primitives. Fortunately, as shown in (Abraham et al. 2006), all results obtained in the mediator setting also apply in the cheap talk setting, except for an arbitrarily small probability of error.

We start by suppressing the incentive-compatibility constraints of the receiver and instead study the mechanism with only two players. Consider a finite state space Ω with a com-

mon prior $\mu \in \Delta(\Omega)$. There are two players that will later play the role of the senders but can also be considered as two political parties. There is a bill that can be either approved or rejected and the utilities of the two players are $u_1, u_2 : A \times \Omega \rightarrow \mathbb{R}$, where $A = \{0, 1\}$. We assume that both players observe the realized state. We call A the action set.

Define a communication protocol (M_1, M_2, τ) that is implemented by a mediator as follows. For $i = 1, 2$ the set M_i is the finite message space of player i . The function $\tau : M_1 \times M_2 \rightarrow \Delta(A)$, which is implemented by the mediator, maps a pair of messages to a probability over the action $a \in A$ (although, for simplicity, for the rest of the paper we will associate $\tau(m_1, m_2)$ with the probability that the mediator suggests 0). A communication protocol defines a game between the two players, where a (behavioral) strategy of player i is a mapping $\sigma_i : \Omega \rightarrow \Delta(M_i)$. A communication protocol (M_1, M_2, τ) together with a profile of strategies $\sigma = (\sigma_1, \sigma_2)$ induces a probability measure $\mathbf{P}_\sigma \in (A \times \Omega)$.

We call a Nash equilibrium of the game induced by a communication protocol a *cheap talk equilibrium*. A *policy* is a mapping $p : \Omega \rightarrow \Delta(A)$. For simplicity we identify $p(\omega)$ with the probability of recommending action $a = 0$ in state ω . Say that a policy p is *cheap talk implementable* (or simply implementable) if there exists a communication protocol (M_1, M_2, τ) and a corresponding cheap talk equilibrium σ such that $\mathbf{P}_\sigma(a = 0|\omega) = p(\omega)$ for every $\omega \in \Omega$.

Our first main goal is to characterize the set of implementable policies p . To do this, we define a binary relation \prec over Ω as follows. For two states ω, ω' we define the relation $\omega' \prec \omega$ iff $u_i(0, \omega) > u_i(1, \omega)$ and $u_j(0, \omega') < u_j(1, \omega')$ for $i, j \in \{1, 2\}$ and $i \neq j$. That is $\omega' \prec \omega$ holds if one of the senders strictly prefers action 0 at ω and the other strictly prefers action 1 at ω' . Note that \prec is not necessary symmetric. To see this assume that both players prefers action 0 to 1 at ω and action 1 to 0 at ω' . Thus we have that $\omega' \prec \omega$ but not vice versa.

Our characterization goes as follows

Theorem 1. $p : \Omega \rightarrow \Delta(A)$ is an implementable policy iff $p(\omega') \leq p(\omega)$ for every pair $\omega, \omega' \in \Omega$ such that $\omega' \prec \omega$.

Surprisingly, the set of implementable policies is independent of the prior μ . We note that in the case where the two players are never indifferent between the outcomes in A , our characterization gets even simpler form. Specifically, the set Ω can be partitioned into four kinds of states: $\Omega_{1,1} \subseteq \Omega$, where both players prefer action 1; $\Omega_{0,0}$, where both prefer action 0; $\Omega_{1,0}$, where player 1 prefers action 1 and player 2 prefers action 0; and $\Omega_{0,1}$, where player 1 prefers action 0 and player 2 prefers action 1.

Corollary 1. In cases where indifference never holds, a policy $p : \Omega \rightarrow [0, 1]$ is implementable iff the following conditions hold:

- (1) The minimum of p among all states in $\Omega_{0,0}$ is (weakly) larger than the maximum of p among all other states.
- (2) The maximum of p among all states in $\Omega_{1,1}$ is (weakly) smaller than the minimum of p among all other states.
- (3) p is constant on $\Omega_{1,0}$.
- (4) p is constant on $\Omega_{0,1}$.

Player 2

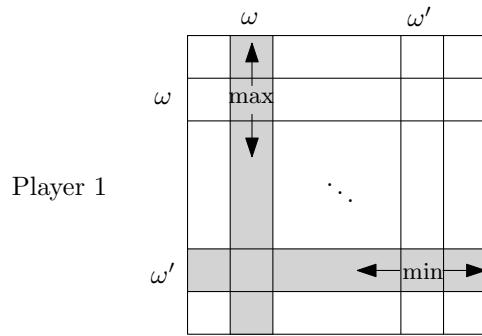


Figure 1: Matrix with the values of τ

Proof of Theorem 1. By the revelation principle, we can transform any cheap talk equilibrium into a *truthful* cheap talk equilibrium, in which both players send the current state to the mediator. Thus, for simplicity, we can assume that $M_1 = M_2 = \Omega$, and that $\sigma_1 \equiv \sigma_2 \equiv \text{Id}_\Omega$. Suppose that player 1 prefers action 0 in state ω and player 2 prefers action 1 in state ω' . Then, player 1 cannot increase the probability that the receiver plays action 0 by sending another state to the mediator, and player 2 cannot decrease such probability. If we plot the values of $\tau(\cdot, \cdot)$ in an $n \times n$ matrix as in Figure 1, we get the following insight: $\tau(\omega, \omega)$ must be the maximum value in the ω column and, simultaneously, $\tau(\omega', \omega')$ must be the minimum value in the ω' row. Thus, it must hold that $\tau(\omega, \omega) \geq \tau(\omega', \omega) \geq \tau(\omega', \omega')$, and therefore that $p(\omega) \geq p(\omega')$.

Conversely, we can check that if $p(\omega) \leq p(\omega')$ whenever $\omega \prec \omega'$, there is a way to construct τ such that $\tau(\omega, \omega) = p(\omega)$ for all $\omega \in \Omega$ and $(\text{Id}_\Omega, \text{Id}_\Omega)$ is a cheap talk equilibrium in (Ω, Ω, τ) . The high-level idea is that the problem reduces to assigning a value to all entries in the matrix of Figure 1 such that (a) $\tau(\omega, \omega)$ is always the greatest (resp., the smallest) entry in the ω column whenever player 1 prefers 0 (resp., prefers 1), and (b) $\tau(\omega, \omega)$ is always the greatest (resp., the smallest) entry in the ω row whenever player 2 prefers 0 (resp., prefers 1). We can get an idea about how to fill this matrix by looking at Figure 1: if player 1 prefers 0 in ω , player 2 prefers 1 in ω' , and $p(\omega) \geq p(\omega')$, then we should simply assign a value to $\tau(\omega, \omega')$ such that $\tau(\omega, \omega) \geq \tau(\omega', \omega) \geq \tau(\omega', \omega')$, e.g., $\tau(\omega', \omega) := \frac{\tau(\omega, \omega) + \tau(\omega', \omega')}{2}$. The same value works if the preferences between player 1 and 2 are reversed. If both have the same preferences, we have that $\tau(\omega', \omega)$ should be either smaller or greater than both $\tau(\omega, \omega)$ and $\tau(\omega', \omega')$, which means that we can set $\tau(\omega', \omega) := 0$ or $\tau(\omega', \omega) := 1$ respectively. More precisely, let τ be such that

1. $\tau(\omega, \omega) = p(\omega)$ for all $\omega \in \Omega$.

2. $\tau(\omega', \omega) = \frac{p(\omega) + p(\omega')}{2}$ if player 1 prefers 0 in ω and player 2 prefers 1 in ω' or, vice-versa, if player 1 prefers 1 in ω' and player 2 prefers 0 in ω .
3. $\tau(\omega', \omega) = 0$ if player 1 prefers 0 in ω and player 2 prefers 0 in ω' .
4. $\tau(\omega', \omega) = 1$ if player 1 prefers 1 in ω and player 2 prefers 1 in ω' .

If a player is indifferent between actions 1 and 0, we assume that she prefers 0. Note that with this assumption, all possible cases are covered, and thus τ is defined for all pairs (ω', ω) . We next show that (Id_Ω, Id_Ω) is a cheap talk equilibrium in (Ω, Ω, τ) .

Suppose that the game state is ω ; we will show that neither player 1 nor player 2 can increase their utility by misreporting the state to the mediator. Suppose that player 1 prefers action 0. If she were to lie and tell the mediator that the state is ω' , there are two possibilities: if player 2 prefers action 1 in ω' , then $\omega' \prec \omega$ and thus, by construction (case 2), $\tau(\omega', \omega) = \frac{p(\omega) + p(\omega')}{2} \leq \tau(\omega, \omega)$, where the last inequality derives from the fact that $\omega' \prec \omega \implies p(\omega') \leq p(\omega)$. This means that, in this case, player 1 cannot increase its utility by reporting ω' instead of ω . If, instead, player 2 prefers 0 in ω' , we are in case 3 and thus $\tau(\omega', \omega) = 0 \leq \tau(\omega, \omega)$ as before. If player 1 prefers 1 in ω , there are again two possibilities: if player 2 prefers 0 in ω' , then again $\tau(\omega', \omega) = \frac{p(\omega) + p(\omega')}{2}$, but in this case $\omega \prec \omega'$ and thus $p(\omega) \leq p(\omega')$, which means that $\tau(\omega', \omega) \geq \tau(\omega, \omega)$ and, therefore, player 1 cannot increase its utility by reporting ω' instead of ω . The remaining possibility is the one in which player 2 prefers action 1 in ω' . However, in this case we have that, by construction, $\tau(\omega', \omega) = 1 \geq \tau(\omega, \omega)$. An analogous argument shows that player 2 cannot increase her utility by misreporting (note that, by construction, τ is symmetric with respect to the preferences of player 1 and player 2). \square

Application to Information Design

In this section we study the implication of Theorem 1 for information design problems. In particular, in addition to the two informed players, who henceforth will be called senders, we have an uninformed receiver with a utility function $v : A \times \Omega \rightarrow \mathbb{R}$. In this setting, the mediator does not play the action immediately, but instead suggests it to the receiver. The receiver plays the action if it is incentive-compatible to do so. More precisely, given protocol (M_1, M_2, τ) with strategy profile (σ_1, σ_2) , the mediator plays the action suggested by the mediator if and only if it gives the receiver a better expected utility than playing any other action. This is formalized in the definition below.

Definition 1. Given communication protocol (M_1, M_2, τ) , a pair of behavioral strategies $\sigma = (\sigma_1, \sigma_2)$ induces a cheap talk equilibrium if, for all $i \in \{1, 2\}$, the strategy σ_i maximizes the utility of sender i given τ and σ_{-i} , and, in addition, $\mathbf{P}_\sigma \in \Delta(A \times \Omega)$ induces an incentive-compatible recommendation for the receiver.

Note that $\mathbf{P}_\sigma \in \Delta(A \times \Omega)$ induces an incentive-compatible recommendation for the receiver if and only if

the following inequality holds:

$$\sum_{\omega \in \Omega} \mathbf{P}_\sigma(\omega | a) v(a, \omega) \geq \sum_{\omega \in \Omega} \mathbf{P}_\sigma(\omega | a) v(1 - a, \omega).$$

That is, τ and σ generate an action recommendation for the receiver. As is standard in the literature, such a recommendation is incentive compatible if, conditional on the recommendation on an action $a \in A$, the receiver is better off accepting the recommendation than playing action $1 - a$. With this, we can define an *implementable policy*.

Definition 2. A policy $p : A \rightarrow [0, 1]$ is implementable if there exists a communication protocol (M_1, M_2, τ) and a cheap talk equilibrium σ such that $\mathbf{P}_\sigma(a = 0 | \omega) = p(\omega)$ for every $\omega \in \Omega$.

Intuitively, p is implementable if there exists a communication protocol and a cheap talk equilibrium that induce p .

Common Interest among Senders

One of the goals of this work is to characterize implementable policies. We first consider the case where the two senders have a common interest. That is, $u_1 = u_2 = u$.

Let β be the utility that the receiver can guarantee when she has no information about the current state (i.e., $\beta = \max(E_\mu(v(1, \omega)), E_\mu(v(0, \omega)))$, let $\Omega_0 = \Omega_{0,0}$ and $\Omega_1 = \Omega_{1,1}$, and, for each policy p , let $q_p \in \Delta(A \times \Omega)$ be the corresponding distribution that is generated by p and the prior μ . The following lemma gives a first characterization of the implementable policies whenever both senders have the same utilities.

Lemma 1. A policy $p : \Omega \rightarrow [0, 1]$ is implementable iff $E_{q_p}[v(a, \omega)] \geq \beta$ and $\min_{\omega \in \Omega_0} p(\omega) \geq \max_{\omega \in \Omega_1} p(\omega)$

Proof. By Theorem 1, the condition $\min_{\omega \in \Omega_0} p(\omega) \geq \max_{\omega \in \Omega_1} p(\omega)$ is necessary and sufficient for p to be implementable by the senders. Clearly, since information cannot harm the receiver, we must have $E_{q_p}[v(a, \omega)] \geq \beta$ for any implementable policy. We now show that the condition $E_{q_p}[v(a, \omega)] \geq \beta$ induces an incentive-compatible recommendation for the receiver. We note that

$$\begin{aligned} E_{q_p}[v(a, \omega)] &= q_p(a = 0) E_{q_p}[v(0, \omega) | a = 0] \\ &\quad + q_p(a = 1) E_{q_p}[v(1, \omega) | a = 1]. \end{aligned}$$

If, by way of contradiction, the policy is not incentive compatible, then either $E_{q_p}[v(1, \omega) | a = 0] > E_{q_p}[v(0, \omega) | a = 0]$ or $E_{q_p}[v(0, \omega) | a = 1] > E_{q_p}[v(1, \omega) | a = 1]$ (or both). In any case, this means that always playing 0 or always playing 1 yields an expected payoff which is higher than β . This contradicts the definition of β . \square

This lemma shows that, for a policy to be implementable, (a) it has to satisfy that 0 is always played with more probability whenever both senders prefer 0 than when both senders prefer 1, and (b) the expected utility of the receiver under this policy should always be better than when she receives no information at all.

It is interesting to see how this setting compares with the one in which there is only one sender but with commitment power, as in (Kamenica and Gentzkow 2011). It turns

out that, in our setting, the set of implementable policies is more restrictive, and thus there are cases in which the optimal Bayesian persuasion mechanism with one sender with commitment power is not implementable, as the following example shows.

Example 1. Consider a state space $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$ with six states. The utility for the senders and the receiver is given in the following table:

	ω_1	ω_2	ω_3	ω_4	ω_5	ω_6
v	0, 2	1, 0	0, 1.9	0, 2.1	1, 0	2, 0
u	0, 1	1, 0	1, 0	1, 0	0, 1	0, 1

For each state ω_i , the left-hand number represents the utility from action 0 and the right-hand number represents the utility from action 1. The prior is the uniform distribution $\mu = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})$. One can show that the optimal Bayesian persuasion policy generates the following conditional probability of action² 0: $p = (0, 1, 1, 0, 0, 0.45)$. Note that $\Omega_0 = \{\omega_2, \omega_3, \omega_4\}$ and $\Omega_1 = \{\omega_1, \omega_5, \omega_6\}$. Thus, $\min_{\omega \in \Omega_0} p(\omega) = 0$ while $\max_{\omega \in \Omega_1} p(\omega) = 0.45$. Lemma 1 implies that p is not implementable.

In the following section, we provide an efficient algorithm that finds the best equilibrium for the senders. Note that this is equivalent to finding the best implementable policy since we can always efficiently construct the protocol and the equilibrium that induces the resulting policy as in the proof of Theorem 1.

Best Sender Equilibrium

Common Interest among Senders In this section we show how to compute the best implementable policy for the senders when both senders have the same utilities. Note that this is equivalent to computing the best equilibrium for the senders since we can efficiently compute the equilibrium that induces the resulting policy as in the proof of Theorem 1. This result will serve as a stepping-stone to the more general algorithm in Appendix that outputs the best policy for the first sender, even in the case where senders don't have the same utilities. The algorithm provided runs in $O(n \log n)$ operations, where n is the number of states in Ω . Before we start, note that we can refine Lemma 1 and get a better characterization of optimal implementable policies.

Lemma 2. Let p^s be the optimal policy for the senders (i.e., $p^s(\omega) = 1$ for every $\omega \in \Omega_0$ and $p^s(\omega) = 0$ for every $\omega \in \Omega_1$), and let β be the receiver's expected utility with no information. If $E_{q_{p^s}}[v(a, \omega)] \geq \beta$, then p^s is the implementable policy that is optimal for the senders. Otherwise, all optimal implementable policies p for the senders satisfy $E_{q_p}[v(a, \omega)] = \beta$.

²The calculation is based on the standard algorithm for computing the optimal persuasion policy for the sender in the case where the action set for the receiver is binary. See, e.g., (Arieli and Danino 2019) and (Renault, Solan, and Vieille 2017).

Proof. If $E_{q_{p^s}}[v(a, \omega)] \geq \beta$, then it is also implementable by Lemma 1 and, by construction, it is also optimal for the senders. Otherwise, by Lemma 1, any optimal implementable policy p must satisfy $E_{q_p}[v(a, \omega)] \geq \beta$. Suppose that $E_{q_p}[v(a, \omega)] > \beta$. By assumption, this means that p is not equal to p^s and, therefore, there exists either $\omega \in \Omega_0$ such that $p(\omega) < 1$ or $\omega \in \Omega_1$ such that $p(\omega) > 0$. Consider the former case. We can define a new policy p' by setting $p'(\omega) = p(\omega) + \delta$ for some small $\delta > 0$, and setting $p'(\omega') = p(\omega')$ for all other $\omega' \neq \omega$. By choosing a value of δ that is small enough, the inequality $E_{q_{p'}}[v(a, \omega)] \geq \beta$ is still satisfied. In addition, since p is incentive-compatible for the senders, so is p' . Therefore, p' is implementable and yields a higher payoff than p to the senders, which contradicts the assumption that p is optimal. A similar construction can be applied when $p(\omega) > 0$ for some $\omega \in \Omega_1$. Hence, $E_{q_p}[v(a, \omega)] = \beta$, as desired. \square

This lemma shows that we can restrict our search to implementable policies that give exactly β utility to the receiver (in addition to the policy in which the receiver always plays according to the senders' preferences). We divide the rest of this section into two parts. First, we focus on a simpler problem that will be needed for the final algorithm and, second, we use the solution to this problem as a primitive for the final algorithm. The following notation will be useful.

Let Ω_C be the set of states where the senders and the receiver agree on the identity of the optimal action. That is, $\omega \in \Omega_C$ iff $[u(1, \omega) - u(0, \omega)][v(1, \omega) - v(0, \omega)] \geq 0$. Let $\Omega_D = \Omega \setminus \Omega_C$ be the set of disagreement states. We distinguish between $\Omega_{D,0} = \Omega_D \cap \Omega_0$ and $\Omega_{D,1} = \Omega_D \cap \Omega_1$. Let $\Omega_{C,0}$ and $\Omega_{C,1}$ be similarly defined.

Because of Theorem 1, we know that any solution must satisfy that, for any $\omega \in \Omega_0$, $p(\omega)$ is greater than all $p(\omega')$ for $\omega' \in \Omega_1$. Given $\alpha \in [0, 1]$, consider the problem of finding the best sender equilibrium that is constrained to $p(\omega) \geq \alpha$ for all $\omega \in \Omega_0$ and $p(\omega) \leq \alpha$ for $\omega \in \Omega_1$. Let p_α denote the solution of this problem for a particular α . By the above property, there exists an alpha such that p_α is the actual best equilibrium for the senders (with no constraints).

Next, we show how to compute p_α . Clearly, in the states ω in which the senders and the receiver have the same action preference a , the mechanism should satisfy $p(\omega) = 1 - a$ (i.e., $p(\omega) = 1$ whenever they all prefer 0 and $p(\omega) = 0$ whenever they all prefer 1). The main difficulty is finding the correct configuration for the disagreement states. Suppose that $\Omega_D = \{\omega_1, \dots, \omega_k\}$ is sorted according to its states' resistance $r(\omega) := \frac{v(1, \omega) - v(0, \omega)}{u(0, \omega) - u(1, \omega)}$ from largest to smallest (note that these values are always positive). Consider the following algorithm that computes p_α (whenever it exists):

- Step 1:** To each state $\omega \in \Omega$, assign $p(\omega) = 1$ if $\omega \in \Omega_0$ and assign $p(\omega) = 0$ otherwise. If the receiver's utility this way is larger than β , return this configuration and terminate.
- Step 2:** Iterate through $\omega_1, \dots, \omega_k$. If $p(\omega_i) = 1$, decrease this value until either $p(\omega_i) = \alpha$ or the receiver's utility equals β . If $p(\omega_i) = 0$, increase this value until either $p(\omega_i) = \alpha$ or the receiver's utility equals β . After

each step, if the utility of the receiver is equal to β , return this configuration and terminate.

3. **Step 3:** If no solution was found in Step 2 (which can only happen if, after iterating through all k elements, the utility of the sender never reached β), there is no incentive-compatible configuration for α .

We claim that this protocol returns p_α if it exists. In fact, given the configuration at the end of Step 1, note that decreasing $p(\omega)$ by Δ for $\omega \in \Omega_{D,0}$ or increasing $p(\omega)$ by Δ for $\omega \in \Omega_{D,1}$ increases the receiver's utility by $\Delta|v(1, \omega) - v(0, \omega)|$ while it decreases the senders' utilities by $\Delta|u(1, \omega), u(0, \omega)|$. By Lemma 1, p is implementable if and only if the receiver gets a utility of at least β . Therefore, in order for the receiver to get utility β while minimizing the losses for the senders, it is optimal to decrease/increase as much as possible the values of p with the highest resistance, as in Step 2. This shows that the algorithm above returns the correct solution.

Now that we are able to compute p_α , it remains to find for what value of α the policy p_α maximizes the senders' utilities. We claim that we can restrict the search to values of α such that p_α attains values only in $\{0, \alpha, 1\}$.

Lemma 3. *There exists $\alpha \in [0, 1]$ such that p_α is optimal for the senders and $p_\alpha(\omega) \in \{0, \alpha, 1\}$ for all $\omega \in \Omega_D$.*

Proof. Suppose that there exists a solution p_α such that j is the first index such that $p_\alpha(\omega_j) \notin \{0, \alpha, 1\}$. By construction, the solution provided by the algorithm satisfies that $p_\alpha(\omega_i) = \alpha$ for all $i < j$, and for $i > j$ $p_\alpha(\omega_i)$ is 1 or 0 depending on whether $\omega_i \in \Omega_0$ or Ω_1 , respectively. Therefore, we can compute the value of $p_\alpha(\omega_j)$ by solving the following equation:

$$\begin{aligned} \beta &= \sum_{i=1}^{j-1} \mu(\omega_i) (\alpha v(0, \omega_i) + (1 - \alpha) v(1, \omega_i)) \\ &+ p_\alpha(\omega_j) \mu(\omega_j) v(0, \omega_j) \\ &+ \sum_{\omega \in \Omega_{j,0}} \mu(\omega) v(0, \omega) + \sum_{\omega \in \Omega_{j,1}} \mu(\omega) v(1, \omega), \end{aligned}$$

where $\Omega_{j,0} := (\Omega \setminus \{\omega_1, \dots, \omega_j\}) \cap \Omega_0$ and $\Omega_{j,1} := (\Omega \setminus \{\omega_1, \dots, \omega_j\}) \cap \Omega_1$. Note that this means that $p_\alpha(\omega_j)$ is locally linear in α , and thus that the senders' expected utility given by p_α is also locally linear in α whenever $p_\alpha(\omega_j)$ attains values outside of $\{0, \alpha, 1\}$. Therefore, the senders' expected utility given by p_α is piecewise linear as a function of α , and hence its maximum lies in one of the segment's endpoints (i.e., in one of the values of α such that $p_\alpha(\omega_j) \in \{0, 1, \alpha\}$ for all $\omega \in \Omega_D$). \square

The algorithm we provide to find the best policy for the senders involves checking the senders' utilities at each of these endpoints. Note that a segment's endpoint is precisely a value of α such that the solution found by the above algorithm attains values only in $\{0, \alpha, 1\}$. By construction, the only possible preimages of α in p_α are $S_1 := \{\omega_1\}$, $S_2 := \{\omega_1, \omega_2\}$, \dots , $S_k := \{\omega_1, \dots, \omega_k\}$. Moreover, for each of these sets S_j , there is at most one value of $\alpha_j \in [0, 1]$ such that $p_{\alpha_j}^{-1}(\alpha_j) = S_j$, and α_j is given by

$$\begin{aligned} \beta &= \sum_{i=1}^j \mu(\omega_i) (\alpha_j v(0, \omega_i) + (1 - \alpha_j) v(1, \omega_i)) + \\ &+ \sum_{\omega \in \Omega_{j,0}} \mu(\omega) v(0, \omega) + \sum_{\omega \in \Omega_{j,1}} \mu(\omega) v(1, \omega). \end{aligned}$$

Isolating α_j from the equation we get

$$\alpha_j = \frac{Z_j}{\sum_{i=1}^j \mu(\omega_i) (v(0, \omega_i) - v(1, \omega_i))},$$

where

$$\begin{aligned} Z_j &= \beta - \sum_{i=1}^j \mu(\omega_i) v(1, \omega_i) \\ &- \sum_{\omega \in \Omega_{j,0}} \mu(\omega) v(0, \omega) \\ &- \sum_{\omega \in \Omega_{j,1}} \mu(\omega) v(1, \omega) \end{aligned}$$

Putting everything together, our algorithm is as follows:

1. **Step 1:** Compute the best possible mechanism for the senders (i.e., set $p(\omega) = 1$ for $\omega \in \Omega_0$ and $p(\omega) = 0$ for $\omega \in \Omega_1$). If this mechanism gives the receiver a utility greater than or equal to β , return this configuration and terminate.
2. **Step 2:** Compute p_0 and p_1 and set the best configuration p to be the one between p_0 and p_1 that reports the most utility to the senders.
3. **Step 3:** For $j = 1, 2, \dots, k$, compute α_j . If $\alpha_j \in [0, 1]$ and p_{α_j} is better for the senders than p , set p to p_{α_j} . Return p .

Note that if Ω_D is sorted by resistance beforehand, this algorithm takes $O(n)$ operations to compute p since all the partial sums used to compute α_j and the senders' utilities can either be precalculated in $O(n)$ operations, or they can simply be updated by adding one term to each sum at each iteration j . If Ω_D is not sorted, the algorithm takes $O(n \log n)$ operations since it needs to sort the states by resistance first.

We can now apply the algorithm to Example 1 above and get that the optimal policy for the sender is $p = (0, 1, \frac{10}{19}, 0, 0, 0)$. Interestingly, this policy yields a utility of $\frac{43}{57} \approx 0.754$ for the senders, as opposed to $\frac{91}{120} \approx 0.758$ in the single sender Bayesian persuasion (with commitment power). In contrast, a single sender's cheap talk equilibrium (with no commitment power) yields a utility of $\frac{1}{2}$ for the sender. An interesting follow-up question is what is the maximal (normalized) loss of the best cheap talk equilibrium over the Bayesian persuasion optimal revelation policy.

General Case In this section we sketch the construction of an $O(n \log n)$ algorithm that outputs the optimal equilibrium for the first sender in the general case, in which senders may not have common interests. The full construction can be found in the full version of the paper³.

Given $\alpha, \gamma \in [0, 1]$, consider the problem of finding the optimal policy $p_{\alpha, \gamma}$ for the first sender such that $p_{\alpha, \gamma}(\omega) = \alpha$ for all $\omega \in \Omega_{1,0}$ and $p_{\alpha, \gamma}(\omega) = \gamma$ for all $\omega \in \Omega_{0,1}$. By Theorem 1, the policy $p_{\alpha, \gamma}$ that gives the most utility to the first sender is also the actual implementable policy that is optimal for the first sender (with no constraints). It is straightforward to check that $p_{\alpha, \gamma}$ can be computed with a slight variation of the algorithm that outputs p_α in the common interest case. Thus, it remains to check which values of α and γ maximize the utility of the first sender.

A similar argument to the one used in the previous section shows that there is at least one optimal policy in which

³<https://arxiv.org/abs/2211.14670>

$p_{\alpha,\gamma}(\omega) \in \{0, \alpha, \gamma, 1\}$ for all $\omega \in \Omega$. Unfortunately, by contrast to the common interest case, this still leaves us with an infinite number of possibilities for α and γ . In fact, it can be shown that, if we assume that $\alpha \leq \gamma$, the possible solutions can be partitioned into several cases (at most n) in which a linear equation on α and γ must be satisfied. Since the expected utility of the first sender is also linear in α and γ , for each of these cases there always exists an optimal solution in the boundary, which is when $\alpha = \gamma$ or $(\alpha, \gamma) \in \{(0, 0), (0, 1), (1, 1)\}$. Considering also the cases in which $\alpha \geq \gamma$ gives us the additional solution where $\alpha = 1$ and $\gamma = 0$. This additional constraint allows us to reduce the number of possible optimal solutions to a finite number which is linear in n . Each of these solutions can be computed in constant time in the same fashion as in the common interest algorithm.

Best Receiver Equilibrium

Most literature on Bayesian persuasion focuses on the best sender equilibrium. This is because the informed sender can also decide on how information is revealed. In our case it makes sense to assume that the information designer that determines the equilibrium selection has the same incentive as the receiver. As we shall now show, determining the best receiver equilibrium is easy.

Here we take the general approach where the preferences of the two senders may not be aligned. For simplicity, we consider the case where no sender is indifferent between the two actions in any state. In this case, we can use Corollary 1 to determine the optimal policy. We call a policy p *pure* if $p : \Omega \rightarrow \Delta(A)$ is a Dirac measure on either 0 or 1. Let $\Omega_F \subseteq \Omega$ be the set of states where all three decision makers have the same preference. Let \mathcal{P} be the set of all pure policies such that (i) $p(\omega)$ recommends the commonly preferred action for every $\omega \in \Omega_F$ and (ii) every $p \in \mathcal{P}$ is constant across types of states for all four different types of states in $\Omega \setminus \Omega_F$ (recall Corollary 1). That is, $p(\omega) = p(\omega')$ for every $\omega, \omega' \in \Omega_{0,0} \setminus \Omega_F$, $p(\omega) = p(\omega')$ for every $\omega, \omega' \in \Omega_{1,0}$, etc.

We claim that by the incentive-compatibility constraints \mathcal{P} contains 6 policies. To see this, note that since any $p \in \mathcal{P}$ is pure and is fixed over states in Ω_F , it can be described by a vector in $\{0, 1\}^4$ according to its values in the four types of states: the first value represents its value in $\Omega_{0,0}$, the second represents its value in $\Omega_{1,0}$, the third represents its value in $\Omega_{0,1}$, and the fourth represents its value in $\Omega_{1,1}$. Corollary 1 asserts that the first value must be the global maximum across the four values and the last value must be the global minimum. We note that the policy $(1, 1, 1, 1)$ dominates the policy $(1, 1, 1, 0)$ for the receiver. This is true since, by definition, in $\Omega_{1,1}^0$ the receiver is in disagreement with both senders and prefers action 0 in all these states. Similarly, $(0, 0, 0, 0)$ dominates $(1, 0, 0, 0)$. We denote the remaining four policies, after omitting the two dominated policies by \mathcal{P}^* .

We have the following simple corollary to determine the optimal policy for the receiver.

Lemma 4. *There exists an optimal policy for the receiver that lies in \mathcal{P}^* .*

Proof. Consider first the set of all implementable policies for the sender. This set is a convex polytope in \mathbb{R}^Ω . The vertices of this polytope are pure policies. In addition, the utility of the receiver is linear over the polytope and so the maximum implementable policy for the receiver is attained as a pure policy. Moreover, we can assume that in Ω_F the recommended action is the consensus action. This is true since taking any policy and altering it in Ω_F according to the consensual action retains incentive-compatibility and increases the utility of the receiver.

To complete the proof we only need to show why a policy p that has both values 0 and 1 over states in $\Omega_{0,0} \setminus \Omega_F$ or two values over states in $\Omega_{1,1} \setminus \Omega_F$ can be improved by a policy in \mathcal{P} . To see this, note that if policy p has both values 0 and 1 over states in $\Omega_{0,0} \setminus \Omega_F$, then by Corollary 1 we must have that p is 0 across all states in $\Omega \setminus \Omega_F$ that are not in $\Omega_{0,0}$. Since, by definition, the receiver prefers action 1 in all states in $\Omega_{0,0} \setminus \Omega_F$, the policy $(0, 0, 0, 0)$ in \mathcal{P}^* dominates p . A similar consideration shows that $(1, 1, 1, 1)$ dominates all implementable policies that attain two values over states in $\Omega_{1,1} \setminus \Omega_F$. \square

Conclusion and Open Problems

In this work, we characterized all incentive-compatible policies in a setting with two senders and one receiver with no commitment power, in which all agents can communicate through a trusted mediator. This characterization is also valid in the cheap talk setting, where there is no mediator and all agents can communicate with each other through private authenticated channels. However, in the cheap talk setting, the implementable policies in general allow a small probability of error. We also provided an $O(n \log n)$ algorithm (where n is the number of states) that finds the optimal policy for each of the senders, and a very simple mechanism that is optimal for the receiver.

Our results show that when there are two senders the equilibrium outcomes are much richer and are closer to those of classical Bayesian persuasion but without the commitment power assumption. A natural question to ask is whether our results can be extended to a more general setting. In particular, it is still open whether one can find a similar characterization in the following settings:

- A setting in which the receiver can play more than two actions.
- A setting in which there are more than two senders, but each of them possesses only partial information about the state. Note that, if there are more than two senders and all of them are fully informed, then any policy p can be implemented by the mediator. This can be done by setting the set of messages equal to the set of states Ω . In this way, the mediator can fully deduce the state ω by taking the majority between the messages sent by the senders and then sampling an action from $p(\omega)$. Thus, for more than two senders the setting is interesting only if they are not fully aware of the state.
- A setting in which there are more than two senders, but up to k of them can collude and deviate from the proposed strategy in a coordinated way.

References

Abraham, I.; Dolev, D.; Gonen, R.; and Halpern, J. 2006. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, 53–62.

Akerlof, G. A. 1978. The market for “lemons”: Quality uncertainty and the market mechanism. In *Uncertainty in economics*, 235–251. Elsevier.

Arieli, I.; Babichenko, Y.; and Sandomirskiy, F. 2022. Bayesian Persuasion with Mediators. *arXiv preprint arXiv:2203.04285*.

Arieli, I.; and Danino, G. 2019. Delegated persuasion. *Available at SSRN 3421954*.

Aumann, R. J. 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica: Journal of the Econometric Society*, 1–18.

Bahar, G.; Smorodinsky, R.; and Tennenholtz, M. 2019. Social Learning and the Innkeeper’s Challenge. In Karlin, A.; Immorlica, N.; and Johari, R., eds., *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, 153–170. ACM.

Conitzer, V.; and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic commerce - EC ’06*. New York, USA: ACM Press.

Crawford, V. P.; and Sobel, J. 1982. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431–1451.

Emek, Y.; Feldman, M.; Gamzu, I.; Leme, R. P.; and Tennenholtz, M. 2012. Signaling schemes for revenue maximization. In Faltings, B.; Leyton-Brown, K.; and Ipeirotis, P., eds., *Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, Valencia, Spain, June 4-8, 2012*, 514–531. ACM.

Fujii, K.; and Sakaue, S. 2022. Algorithmic Bayesian persuasion with combinatorial actions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5): 5016–5024.

Gan, J.; Majumdar, R.; Radanovic, G.; and Singla, A. 2022. Bayesian persuasion in sequential decision-making. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5): 5025–5033.

Kamenica, E.; and Gentzkow, M. 2011. Bayesian persuasion. *American Economic Review*, 101(6): 2590–2615.

Kamenica, E.; and Gentzkow, M. 2017. Competition in persuasion. *Review of economic studies*, 84(1): 1.

Kosenko, A. 2018. Mediated persuasion. *Available at SSRN 3276453*.

Kremer, I.; Mansour, Y.; and Perry, M. 2013. Implementing the “Wisdom of the Crowd”. In Kearns, M. J.; McAfee, R. P.; and Tardos, É., eds., *Proceedings of the fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, 605–606. ACM.

Krishna, V.; and Morgan, J. 2001. A model of expertise. *The Quarterly Journal of Economics*, 116(2): 747–775.

Lipnowski, E.; and Ravid, D. 2020. Cheap talk with transparent motives. *Econometrica*, 88(4): 1631–1660.

Morgan, M. S.; and Morrison, M. 1999. *Models as mediators*. Cambridge University Press Cambridge.

Renault, J.; Solan, E.; and Vieille, N. 2017. Optimal dynamic information provision. *Games and Economic Behavior*, 104: 329–349.

Salamanca, A. 2021. The value of mediated communication. *Journal of Economic Theory*, 192: 105191.