



On the Benefits of Being Constrained When Receiving Signals

Shih-Tang Su^{1(✉)}, David Kempe², and Vijay G. Subramanian¹

¹ Electrical Engineering and Computer Science, University of Michigan,
Ann Arbor, USA
shihtang@umich.edu

² Department of Computer Science, University of Southern California,
Los Angeles, USA

Abstract. We study a Bayesian persuasion setting in which the receiver is trying to match the (binary) state of the world. The sender’s utility is partially aligned with the receiver’s, in that conditioned on the receiver’s action, the sender derives higher utility when the state of the world matches the action.

Our focus is on whether in such a setting, being constrained helps a receiver. Intuitively, if the receiver can only take the sender’s preferred action with smaller probability, the sender might have to reveal more information, so that the receiver can take the action more specifically when the sender prefers it. We show that with a binary state of the world, this intuition indeed carries through: under very mild non-degeneracy conditions, a more constrained receiver will always obtain (weakly) higher utility than a less constrained one. Unfortunately, without additional assumptions, the result does not hold when there are more than two states in the world, which we show with an explicit example.

Keywords: Bayesian persuasion · Information design · Signaling games

1 Introduction

In this paper, we study situations akin to the following stylized dialog, which will likely be familiar to anyone who has ever served on hiring committees:

ALICE: I see that you wrote strong recommendation letters for your Ph.D. graduates Carol and Dan. Can you compare them for us?

BOB: They are both great! Carol made groundbreaking contributions to ...; Dan made groundbreaking contributions to

ALICE: Which of the two would you say is stronger?

S.-T. Su—Supported in part by NSF grant ECCS 2038416 and MCubed 3.0.

V. G. Subramanian—Supported in part by NSF grants ECCS 2038416, CCF 2008130, and CNS 1955777.

BOB: They are hard to compare. You really need to interview both of them!

ALICE: We can only invite one of them for an interview.

BOB: I guess Carol is a bit stronger.

What happened in this example? Alice and Bob were involved in a signaling setting, in which Bob had an informational advantage. Bob’s goal was to get as many of his students interviews as possible, while Alice wanted to only invite the strong students. While Bob knew which of his students were strong (or how strong), Alice had to rely on information from Bob. As is standard in signaling settings, Bob could use this fact to improve his own utility. In this sense, the example initially was virtually identical to the standard “judge/prosecutor” example in the seminal paper of Kamenica and Gentzkow [21].

However, a change happened along the way. When Alice revealed that she was constrained in her actions (to one interview at most), this changed the utility that Bob could obtain from his previous strategy. For example, if he had insisted on not ranking the students, Alice might have flipped a coin. Implicitly, while Bob wanted *both* of his students to obtain interviews, when forced to choose, he knew he would obtain higher utility from the stronger of his students being interviewed. In this sense, his utility function was “partially aligned” with Alice’s; this partial alignment, coupled with Alice’s constraint, resulted in Alice obtaining more information, and thus higher utility.

The main goal of the present paper is to investigate to what extent the behavior illustrated informally in the dialog above arises in a standard model of Bayesian persuasion. Specifically, if the utilities of the sender and receiver are “partially aligned,” will it always benefit a receiver to be more constrained in how she can choose her actions?

1.1 The Model: An Overview

Our model—described fully in Sect. 2—is based on the standard Bayesian persuasion model of Kamenica and Gentzkow [21]. For our main result, we assume that the state space is binary: $\Theta = \{\theta_1, \theta_2\}$. These states could correspond to a student being bad/good in our introductory example, a defendant being innocent/guilty in the example of Kamenica and Gentzkow [21], or a stock about to go up or down. The sender and receiver share a common prior Γ for the distribution of the state θ . In addition, the sender will observe the actual state θ , but only after committing to a signaling scheme (also called information structure).

A signaling scheme is a (typically randomized) mapping $\phi : \Theta \rightarrow \Sigma$. The receiver observes the (typically random) signal $\sigma = \phi(\theta)$; based on this observation, she takes an action $a \in A$. Here, we assume that—like the state space—the action space is binary, i.e., $A = \{a_1, a_2\}$. Based on the true state of the world and the action taken by the receiver, both the sender and receiver derive utilities $U_S(\theta, a)$, $U_R(\theta, a)$. The receiver will choose her action (upon observing σ) to maximize her own expected utility; the sender, knowing that the receiver is rational, will commit to a signaling scheme to maximize his expected utility under rational receiver behavior.

Motivated by many practical applications, we assume that the receiver prefers to match the state of the world, in the sense that $U_R(\theta_1, a_1) \geq U_R(\theta_1, a_2)$ and $U_R(\theta_2, a_2) \geq U_R(\theta_2, a_1)$. For instance, in our introductory example, Alice prefers to interview strong candidates and to not interview weak ones; in the judge-prosecutor example, the judge prefers convicting exactly the guilty defendants; and an investor prefers to buy stocks that will go up and sell stocks that will go down. Our assumption about the “partial alignment” of the sender and receiver utilities is formalized as an *action-matching* preference of the sender, stated as follows: $U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2)$. That is, if a candidate is being interviewed, Bob prefers it to be a strong candidate over a weak one (but may still prefer a weak candidate being interviewed over a strong/weak candidate *not* being interviewed); similarly, if a prosecutor sees a defendant convicted, he would prefer the defendant to be guilty (but may still prefer an innocent defendant being convicted over going free); similarly, an investment platform may want to entice a client to buy stock, but conditioned on the client buying stock, the platform may prefer for the stock to go up.

In addition to the assumption of partial alignment, our main addition to the standard Bayesian persuasion model is to consider constraints on the receiver’s actions. Specifically, we assume that there are (lower and upper) bounds \underline{b}, \bar{b} on the probability with which the receiver is allowed to take action¹ a_1 . Such a constraint corresponds to a department only being willing to interview at most 10% of their applicants, a judge having a quota for how many defendants (at most) to convict, or a conference having an upper bound on its number/fraction of accepted papers. Such a constraint creates dependencies between the receiver’s actions under different received signals, and may force her to randomize between actions, contrary to the standard Bayesian persuasion setting in which the receiver may deterministically choose any utility-maximizing action conditioned on the observed signal σ . To see this, consider a prior under which a candidate is strong with probability $\frac{1}{3}$, and the receiver obtains utility 1 from interviewing a strong candidate and -1 from interviewing a weak candidate (and 0 from not interviewing). If the sender reveals no information, the receiver would prefer to interview no candidates, but a lower-bound constraint may force her to do so, in which case she would randomize the decision to interview the smallest total number of candidates. We write $\pi : \Sigma \rightarrow A$ for the receiver’s (typically randomized) mapping from signals to actions. Note that the constraint applies across all sources of randomness (the state of the world, the sender’s randomization, and the receiver’s randomization), so it is required that $\underline{b} \leq \mathbb{P}_{\Gamma, \phi, \pi}[\pi(\phi(\theta)) = a_1] \leq \bar{b}$.

To avoid trivialities, we assume that $\mathbb{P}_{\Gamma}[\theta = \theta_1] \in [\underline{b}, \bar{b}]$, that is, if the sender revealed the state of the world perfectly, the receiver would be allowed to match it. We say that a receiver with constraints $(\underline{b}', \bar{b}')$ is *more constrained* than one with constraints (\underline{b}, \bar{b}) iff $\underline{b}' \geq \underline{b}$ and $\bar{b}' \leq \bar{b}$.

¹ This implies constraints of $1 - \bar{b}, 1 - \underline{b}$ on the probability of taking action a_2 . A more general model and its specialization to binary actions is discussed in Sect. 2.3.

1.2 Our Results

Our main result is that when the state of the world is binary, a receiver is always (weakly) better off when more constrained. We state this result here informally, and revisit it more formally (and prove it) in Sect. 3.

Theorem 1 (Main Theorem (stated informally)). *Consider a Bayesian persuasion setting in which the state and action spaces are binary, the receiver is trying to match the state of the world, and the sender is action-matching. Then, a more constrained receiver always obtains (weakly) higher utility than a less constrained one.*

Unfortunately, this insight does not extend to more fine-grained states of the world: even for a ternary state of the world, there are examples with partially aligned sender and receiver in which a more constrained receiver is strictly worse off. We discuss such an example in depth in Sect. 4. It is possible to obtain some positive results recovering versions of Theorem 1 by imposing additional constraints on the sender’s and receiver’s utility functions. However, many of these constraints are strong, and have only limited applicability to real-world settings. We discuss some of these approaches in Sect. 5—whether there are less restrictive conditions recovering Theorem 1 for more states of the world is an interesting direction for future work.

1.3 Related Work

In general, information design as an area is concerned with situations in which a better-informed sender or information designer can influence the behavior of other agents via the provision of information. The literature generally studies problems in which the underlying game between the agents is given and fixed, but where the sender can influence the outcome by an appropriate choice of information to be disclosed. The core difference between Bayesian persuasion [3, 5, 6, 20, 21] and other standard paradigms that study information transmission (such as cheap talk [11], verifiable messages [17, 27] or signaling games [33]) is the commitment power of the sender. In Bayesian persuasion models, the sender moves first and commits to a (typically randomized) mapping from states of the world to signals. Subsequently, the sender observes the state of the world and applies the mapping. Based on the mapping and the observed signal, the rational recipients (called receivers) choose actions.

The study of Bayesian persuasion was initiated in the seminal work of Kamenica and Gentzkow [21] and Rayo and Segal [31]. In their work, the sender can *commit to sending any distribution of messages* before (accurately) observing the state of the world; the receiver, on the other hand, only has knowledge of the prior. The full commitment setting allows for an equivalence to an alternate model where the sender publicly chooses the amount of information (regarding the state of the world) he will privately observe and then (strategically) decides how much of this information to share with the receiver via verifiable messages. Follow-up work of Bergemann and Morris [3, 5] established a useful

and important equivalence between the set of outcomes achievable via information design and Bayes correlated equilibria. Since these seminal works, there has been a large body of work on Bayesian persuasion with theoretical developments as well as a multitude of applications. To keep our discussion focused, for the broader literature, we refer the reader to survey articles [6, 20].

The literature closest to our work studies information design with a constrained sender: the constraints arise through a diverse set of assumptions. The work in [29, 30] shows that pooling equilibria result if the receiver either prefers lower complexity (for a certification process) or performs a validation of the sender's signal; this holds whether the signals of the sender are exogenously constrained or not. A growing body of work considers constraints on the sender that arise either due to communication costs for signaling [10, 16, 19, 28], capacity limitations for signaling [13, 25], the sender's signal serving multiple purposes (such as convincing a third party to take a payoff-relevant action) [7], or costs to the receiver for acquiring additional information [26]. The contributions are then to characterize either the applicability of the concavification approach [21], the optimal signaling structure, or the conditions for the optimality of certain signaling structures. In [22], constraints on the sender arise from the receiver having access to some publicly available information. Within this context, Kolotilin [22] studies comparative statics on the sender's utility based on the quality of the sender's information or the public information. There is also a burgeoning literature on constraints on the sender arising from a privately informed receiver (e.g., [8, 9, 12, 18, 24]). The main contributions in this line of research are to characterize the optimal signaling structure with a key aspect being the fact that the sender constructs a different signal for each receiver type.

Based on the discussion above, clearly there is significant literature studying a constrained sender's optimal signaling scheme and utility. However, work that studies constraints on the *receiver*, or their impact on the receiver's utility, is extremely limited. To the best of our knowledge, [2] is the only work to analyze a constrained receiver problem. The authors impose *ex ante* and *ex post* constraints on the receiver's posterior beliefs, characterize the dimensionality of the optimal signaling structure and develop low-complexity approximate welfare maximizing algorithms. In our work, we have two important differences: first, we impose constraints on the receiver's *actions* as opposed to posterior beliefs; and second, we explore when these constraints result in increased utility for the receiver.

2 Problem Formulation

Our model is based on the standard Bayesian persuasion model [21]. Two players, a sender and a receiver, interact in a signaling game. The sender can observe the state of the world, while the receiver can take an action. The sender can convey information about the state of the world to the sender. Both players receive utility as a function of both the state of the world and the action chosen by the receiver. Since their utility functions typically do not align, the sender will be strategic in the information he reveals to the receiver.

2.1 State of the World, Actions, and Utilities

The (random) state of the world θ is drawn from a state space Θ . For our main result, we assume that the state space is binary ($\Theta = \{\theta_1, \theta_2\}$); however, we define the model in more generality. The sender and receiver share a common-knowledge prior distribution $\Gamma \in \Delta(\Theta)$ for θ . When the state space is binary, this prior is fully characterized by $p = \mathbb{P}_\Gamma[\theta = \theta_1]$.

Only the receiver can take an action $a \in A$. Again, for our main result, we assume that the action space is binary: $A = \{a_1, a_2\}$. Both the sender's and the receiver's utilities are functions of the true state θ and the action taken; they are captured by the functions $U_S : \Theta \times A \rightarrow \mathbb{R}$ and $U_R : \Theta \times A \rightarrow \mathbb{R}$. As discussed in Sect. 1.1, we assume that the receiver tries to match the state of the world with her action.

Definition 1 (State-Matching Receiver). *We say that the receiver's utility function is state-matching if it satisfies the following: for all i, j, k with $i \leq j \leq k$ or $i \geq j \geq k$, we have that*

$$U_R(\theta_i, a_j) \geq U_R(\theta_i, a_k). \quad (1)$$

When the state of the world is binary, the condition simplifies to:

$$U_R(\theta_1, a_1) \geq U_R(\theta_1, a_2) \text{ and } U_R(\theta_2, a_2) \geq U_R(\theta_2, a_1). \quad (2)$$

In words, a state-matching receiver always prefers an action closer to the true state; however, the definition does not enforce any comparisons between choosing an action that is too high vs. too low compared to the true state.

The key notion for our analysis is a partial alignment of the sender's utility with the receiver's. This is captured by the fact that the sender, given any fixed action, would prefer states closer to the action, expressed in Definition 2:

Definition 2 (Action-Matching Sender). *We say that the sender's utility function is action-matching if it satisfies the following: for all i, j, k with $i \leq j \leq k$ or $i \geq j \geq k$, we have that*

$$U_S(\theta_j, a_i) \geq U_S(\theta_k, a_i). \quad (3)$$

When the state of the world is binary, the condition simplifies to:

$$U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1) \text{ and } U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2). \quad (4)$$

In words, an action-matching sender always prefers a state of the world closer to the action chosen by the receiver; again, we do not enforce any comparisons between states that are higher vs. lower than the chosen action.

Notice the difference between Inequalities (3) and (4) vs. (1) and (2): (1) and (2) compare the receiver's utilities when the state of the world is fixed and the action is changed, while (3) and (4) compare the sender's utilities when the action is fixed and the state of the world is changed. That is, given that the receiver takes a particular action, the sender derives higher utility when that action more closely matches the state of the world than when it does not. Again, a justification for this assumption is discussed in Sect. 1.1.

2.2 Signaling Schemes

Before the receiver takes her action, the sender can send a signal σ to reveal (partial) information about the state of the world. More precisely, prior to observing the state of the world θ , the sender commits to a signaling scheme ϕ , which is a mapping $\phi : \Theta \rightarrow \Delta(\Sigma)$. For our purposes ϕ is conveniently characterized by the probability with which each signal is sent conditional on the state. We write $\phi_{i,j} = \mathbb{P}[\phi(\theta) = \sigma_j \mid \theta = \theta_i] \in [0, 1]$ for the probability that signal σ_j is sent conditional on the state of the world being θ_i . We write $\bar{\phi}_j = \sum_i \mathbb{P}_\Gamma[\theta = \theta_i] \cdot \phi_{i,j}$ for the probability of sending the signal σ_j .

The receiver is Bayes-rational, and her objective is to maximize her expected utility after observing the signal. The expected utility derived from action a when observing σ_j can be written as

$$\bar{U}_R(\sigma_j, a) = \sum_{\theta_i \in \Theta} \mathbb{P}[\theta = \theta_i \mid \phi(\theta) = \sigma_j] \cdot U_R(\theta_i, a) = \sum_{\theta_i \in \Theta} \frac{\mathbb{P}[\theta = \theta_i] \cdot \phi_{i,j}}{\bar{\phi}_j} \cdot U_R(\theta_i, a).$$

Thus, barring other constraints (which we will introduce below), the receiver chooses an action a in $\operatorname{argmax}_a \bar{U}_R(\sigma_j, a)$. Following most of the literature in the field of information design, we assume that the receiver breaks ties in favor of an action most preferred by the sender. The following very useful alternative view has been observed in the prior literature (see, e.g., [4]): instead of sending abstract signals, the sender can without loss of generality send the receiver a recommended action a_j . The sender must ensure that ϕ is such that the receiver will always voluntarily follow the recommendation. In other words, the recommended action a_j must always be in $\operatorname{argmax}_a \bar{U}_R(\sigma_j, a)$. This constraint ensures *ex-post incentive compatibility (EPIC)* of the signaling scheme, and is often referred to as an *obedience constraint*.

We write $\pi : \Sigma \rightarrow \Delta(A)$ for the receiver's (possibly randomized) best-response function. In the setting described so far, there is actually no need for the receiver to randomize, and she can always choose any arbitrary deterministic $\pi(\sigma_j) \in \operatorname{argmax}_a \bar{U}_R(\sigma_j, a)$. However, as we will see in Sect. 2.3, the situation changes when the receiver is constrained. For a receiver strategy π , we write $\pi_{i,j} = \mathbb{P}[\pi(\sigma_j) = a_i]$ for the probability that the receiver, upon observing signal σ_j , chooses action a_i .

The sender's objective is to design a signaling strategy which maximizes his expected utility in the subgame perfect equilibrium. That is, he chooses ϕ so as to maximize his expected utility (under all sources of randomness)

$$\mathbb{E}_{\theta \sim \Gamma, \sigma \sim \phi(\theta), a \sim \pi(\phi)(\sigma)} [U_S(\theta, a)],$$

assuming a best response $\pi^{(\phi)}$ from the receiver.

2.3 Constrained Receiver

Our main conceptual departure from prior work is that we consider constraints on the receiver, restricting the probability with which actions can be chosen. In a

general setting, such constraints are lower and upper bounds on the probability of taking each action, i.e., \underline{b}_a and \bar{b}_a for each a . Formally, we require that for each action a_i , the combination of the sender's signaling scheme ϕ and the receiver's response π satisfy

$$\underline{b}_{a_i} \leq \sum_j \bar{\phi}_j \cdot \pi_{i,j} \leq \bar{b}_{a_i}. \quad (5)$$

The constraints are common knowledge among the sender and receiver. When the state space is binary, the constraints can be simplified: they are fully characterized by the lower and upper bounds $\underline{b} = \max(\underline{b}_{a_1}, 1 - \bar{b}_{a_2})$, $\bar{b} = \min(\bar{b}_{a_2}, 1 - \underline{b}_{a_1})$ for the probability with which the receiver can choose action a_1 .

The focus of our work is on whether being (more) constrained helps the receiver, by forcing an action-matching sender to disclose “more” information. Without any further assumptions, this is trivially false. For example, suppose that the state of the world is known to be θ_1 with probability 1, and both the sender and the receiver obtain utility 1 when the receiver chooses action a_1 , and 0 otherwise. If the constraint specified that a_1 must be taken with probability 0, and a_2 with probability 1, then of course, the receiver (and the sender) would be worse off. In order to allow us to clearly articulate the question of whether a constrained receiver obtains more information, we require that perfect state matching would always be feasible for the receiver, if the true state were revealed:

Definition 3 (Implementable and Feasible Constraints). *Consider constraints $\langle \underline{b}_{a_i}, \bar{b}_{a_i} \rangle$ for all $a_i \in A$. We say that the constraints are **implementable** iff $\sum_i \underline{b}_{a_i} \leq 1 \leq \sum_i \bar{b}_{a_i}$.*

*The constraints are **feasible** iff $\underline{b}_{a_i} \leq \mathbb{P}_T[\theta = \theta_i] \leq \bar{b}_{a_i}$ for all i .*

*For the special case of a binary state space, a constraint $\langle \underline{b}, \bar{b} \rangle$ is **feasible** iff $\underline{b} \leq p \leq \bar{b}$.*

Notice that when constraints are not implementable, there is no strategy for the receiver to satisfy all constraints. When constraints are feasible, then with full information, perfect state matching can be implemented by the receiver.

We say that the constraints $\langle \underline{b}_{a_i}, \bar{b}_{a_i} \rangle$ are *more binding* (or the receiver is more constrained by them) than $\langle \underline{b}'_{a_i}, \bar{b}'_{a_i} \rangle$ if and only if $\underline{b}'_{a_i} \leq \underline{b}_{a_i}$ and $\bar{b}_{a_i} \leq \bar{b}'_{a_i}$ for all i . When the state space is binary, the condition simplifies: the constraint $\langle \underline{b}, \bar{b} \rangle$ is more binding than $\langle \underline{b}', \bar{b}' \rangle$ if and only if $\underline{b}' \leq \underline{b}$ and $\bar{b} \leq \bar{b}'$.

We note that the presence of a constraint may force the receiver to randomize between actions, even possibly actions that are not optimal. For a simple example, suppose that the state of the world is binary and determined by a fair coin flip, and the receiver obtains utility 2 from matching state θ_2 , 1 from matching state θ_1 , and 0 for not matching the state. If the sender reveals no information, then a receiver constrained by—say— $\underline{b} = \bar{b} = \frac{1}{2}$, would have to flip a fair coin to decide which action to choose, even though the optimal strategy would be to always choose a_2 .

While the receiver's best response π may in general (have to) be randomized, we show that there is always an optimal signaling strategy for the sender such

that the receiver will play a deterministic strategy π . Notice that the following proposition does not even require feasibility in the sense that the prior distribution satisfies the constraints: it merely requires that the constraints allow for the existence of *any* signaling scheme and corresponding receiver strategy.

Proposition 1. *Assume that $|\Sigma| \geq |A|$, and let $\langle b_{a_i}, \bar{b}_{a_i} \rangle$ (for all i) be implementable constraints on the receiver. Then, for any signaling scheme $\hat{\phi}$, there exists another signaling scheme ϕ under which the sender has at least the same utility as under $\hat{\phi}$, and such that the receiver's best response $\pi^{(\phi)}$ is deterministic. In particular, there is a sender-optimal strategy under which the receiver responds deterministically.*

Proof. We will give an explicit construction of such a strategy. Let $\hat{\phi}$ be any signaling scheme. Let $\pi^{(\hat{\phi})}$ be the receiver's (randomized) best response. Recall that $\pi_{i,j}$ is the probability with which the receiver plays a_i when receiving the signal σ_j . We will first construct an intermediate signaling scheme ϕ' , and from it the final signaling scheme ϕ .

As a first step, the signaling scheme maps to an expanded space $\Sigma' = \Sigma \times A$. When observing the state θ_k , the sender sends the signal (σ_j, a_i) with probability $\phi_{k,j} \cdot \pi_{i,j}$. In other words, the sender performs exactly the randomization that the receiver would perform, and makes the corresponding recommendation to the receiver. Conditioned on the signal σ_j , the signal's second component a_i reveals no information about the state of the world. Therefore, because the distribution of a_i is exactly the distribution that $\pi^{(\hat{\phi})}(\sigma_j)$ uses, it is a best response for the receiver (and satisfies the constraints) to deterministically² follow the sender's "recommendation" a_i when receiving the signal (σ_j, a_i) .

Then, following the standard approach for reducing the size of the signal space, we "compress" all signals under which the receiver chooses the same action into one signal. That is, under the final signaling scheme ϕ , whenever the sender was going to send (σ_j, a_i) for any j under ϕ' , the sender simply sends a_i . Because it is a best response for the receiver to deterministically choose a_i for all received (σ_j, a_i) , it is still a best response to follow the recommendation a_i .

Thus, we have constructed a signaling scheme ϕ such that the receiver plays a deterministic best response, and the number of signals employed by the sender is at most $|A|$.

Finally, to prove the existence of a sender-optimal signaling scheme with deterministic receiver response, let $\hat{\phi}$ be any sender-optimal signaling scheme. The existence of a signaling scheme, and thus a sender-optimal one, follows because the constraints are implementable by assumption. Then, applying the previous argument to $\hat{\phi}$ gives the desired optimal signaling scheme with deterministic receiver responses. \square

² Note that it is optimal for the receiver to follow the recommendation due to the overall constraints. In isolation, the receiver may be better off deviating for some signals—however, doing so would violate a constraint, or come at the expense of having to choose an even more suboptimal action under another signal.

In general, most of the literature on Bayesian persuasion assumes that the signal space is at least as large as the action space (which is enough to obtain sender-optimal strategies, and find them via an LP [23] when EPIC holds). Hence, we make the same assumption that $|\Sigma| \geq |A|$ in Proposition 1.

Henceforth, we will restrict attention to signaling schemes with deterministic best response functions π without loss of optimality. However, the sender still has to ensure that following the deterministic recommendation is incentive compatible for the receiver. Since the receiver is constrained, her space of deviations is only to best-response functions satisfying the constraints. This is captured by the following definition:

Definition 4. *Let $\phi : \Theta \rightarrow \Sigma$ be a direct signaling scheme for the sender, i.e., making action recommendations and assuming $\Sigma = A$. Let Π be the set of all randomized mappings $\pi : \Sigma \rightarrow A$ (characterized by $\pi_{i,j}$) satisfying the following inequalities for all actions a_j :*

$$\underline{b}_{a_j} \leq \sum_i \bar{\phi}_i \cdot \pi_{i,j} \leq \bar{b}_{a_j}.$$

Then, ϕ is ex ante incentive compatible iff for all feasible response functions $\pi \in \Pi$,

$$\sum_i \bar{\phi}_i \cdot \bar{U}_R(\sigma_i, a_i) \geq \sum_i \sum_j \bar{\phi}_i \cdot \pi_{i,j} \cdot \bar{U}_R(\sigma_i, a_j).$$

Note that the presence of constraints forces us to deviate from the standard EPIC requirement in the literature. Definition 4 bears similarity to definitions in [2, 9, 12], where ex ante constraints are considered.

3 Our Main Result

In this section, we present the main result of this paper.

Theorem 2. *Consider a Bayesian persuasion setting in which the state and action spaces are binary. The receiver is state-matching, and the sender is action-matching. Let $\langle \underline{b}, \bar{b} \rangle$ and $\langle \underline{b}', \bar{b}' \rangle$ be two feasible constraints such that $\langle \underline{b}, \bar{b} \rangle$ is more binding than $\langle \underline{b}', \bar{b}' \rangle$, and let Φ, Φ' be the set of all sender-optimal signaling schemes under these constraints, respectively.*

Let $\phi \in \operatorname{argmax}_{\phi \in \Phi} \mathbb{E}_{\theta \sim \Gamma, \sigma \sim \phi(\theta), a \sim \pi^{(\phi)}(\sigma)} [U_R(\theta, a)]$ maximize the receiver's utility over Φ , and $\phi' \in \operatorname{argmax}_{\phi' \in \Phi'} \mathbb{E}_{\theta \sim \Gamma, \sigma \sim \phi'(\theta), a \sim \pi^{(\phi')}(\sigma)} [U_R(\theta, a)]$ maximize the receiver's utility over Φ' . Then the receiver is no worse off under ϕ than under ϕ' , i.e., $\mathbb{E}_{\theta \sim \Gamma, \sigma \sim \phi(\theta), a \sim \pi^{(\phi)}(\sigma)} [U_R(\theta, a)] \geq \mathbb{E}_{\theta \sim \Gamma, \sigma \sim \phi'(\theta), a \sim \pi^{(\phi')}(\sigma)} [U_R(\theta, a)]$.

Proof. At a high level, the intuition for the proof is as follows. Based on the discussion in Sect. 2.3, the constraints on the receiver actually translate into constraints on the sender in the optimization problem. Because the sender's signaling schemes are more constrained, he has to reveal more information. However, this intuition is not complete—after all, the constraints may entice the

sender to reveal *less* information. Furthermore, as we see in Sect. 4, when the state space is not binary, a more constrained receiver may be worse off.

Let ϕ, ϕ' be as defined in the statement of the theorem, and let $\phi_{i,j}, \phi'_{i,j}$ be their corresponding conditional probabilities of sending the signal σ_j in state θ_i . By Proposition 1, w.l.o.g., under the sender-optimal strategies ϕ, ϕ' , the sender recommends an action to the receiver, and the receiver deterministically follows the recommendation. That is, the signal σ_i can be associated with the action a_i for $i = 1, 2$. Our proof is based on distinguishing four cases, depending on the sender's utility:

1. $U_S(\theta_1, a_1) \geq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_2, a_1)$

In this case, for every state, the sender prefers the same action as the receiver. Since the sender's and the receiver's preferences are fully aligned, the sender's optimal strategy is to fully reveal the state of the world. Since the constraints are feasible, the receiver can perfectly match the state of the world under both constraints, and hence obtains the same utility under both constraints.

2. $U_S(\theta_1, a_1) \geq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \leq U_S(\theta_2, a_1)$

In this case, the sender always prefers action a_1 . Since the sender is action-matching, $U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2)$. Combining these inequalities, we obtain that the sender's utility function satisfies the following total order:

$$U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1) \geq U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2).$$

This implies that

$$U_S(\theta_1, a_1) - U_S(\theta_1, a_2) \geq U_S(\theta_2, a_1) - U_S(\theta_2, a_2). \quad (6)$$

We now show that $\phi_{1,2} = 0$. An identical proof also shows that $\phi'_{1,2} = 0$. We distinguish two cases:

- (a) If $\phi_{1,2} > 0$ and $\phi_{2,1} > 0$, then the sender could move some probability mass $\epsilon > 0$ from recommending a_2 under θ_1 to recommending a_1 , and in return move the same amount from recommending a_1 under θ_2 to recommending a_2 . Because the receiver is state-matching, she will still follow the sender's recommendation, and the total probability with which each action is played stays unchanged, so the strategy is still feasible. By Eq. (6), the sender's utility (weakly) increases. By choosing ϵ as large as possible, we arrive at the claim or at the following case.
- (b) If $\phi_{1,2} > 0$ and $\phi_{2,1} = 0$, then $\bar{\phi}_{a_1} = p \cdot \phi_{1,1} < p \leq \bar{b}$. Therefore, it is feasible for the sender to always send the signal σ_1 when the state is θ_1 (i.e., decrease $\phi_{1,2}$ to 0 and increase $\phi_{1,1}$ by the same amount). Again, because the receiver is state-matching, she will still follow the sender's recommendation, and the sender is weakly better off because $U_S(\theta_1, a_1) \geq U_S(\theta_1, a_2)$.

Because $U_S(\theta_2, a_1) \geq U_S(\theta_2, a_2)$ and $\phi_{1,1} = 1$ (as proved above), the sender will also send σ_1 with as much probability as possible when the state is θ_2 , subject to not violating the receiver's incentive to play a_1 and not exceeding

the upper bound \bar{b} (or \bar{b}'). In other words, the sender maximizes $\phi_{1,2}$ subject to $\mathbb{E}_{\theta \sim \Gamma, \sigma \sim \phi(\theta)} [U_R(\theta, a_1) | \sigma_1] \geq \mathbb{E}_{\theta \sim \Gamma, \sigma \sim \phi(\theta)} [U_R(\theta, a_2) | \sigma_1]$ and $\bar{b} \geq \bar{\phi}_{a_1}$ (or $\bar{b}' \geq \bar{\phi}_{a_1}$). Using $\phi_{1,1} = 1$, the incentive constraint is equivalently expressed as $\phi_{2,1} \leq \frac{p \cdot (U_R(\theta_1, a_1) - U_R(\theta_1, a_2))}{(1-p) \cdot (U_R(\theta_2, a_2) - U_R(\theta_2, a_1))}$. Since this inequality is independent of the bound and \bar{b}' is more restricted than \bar{b} , the receiver is weakly better off under the constraint \bar{b} than under \bar{b}' .

3. $U_S(\theta_1, a_1) \leq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_2, a_1)$

This case is symmetric to the previous one. Here, the roles of a_1 and a_2 (and θ_1 and θ_2) are reversed, and the important constraint becomes the lower bound \underline{b} (and \underline{b}') rather than the upper bound \bar{b} .

4. $U_S(\theta_1, a_1) \leq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \leq U_S(\theta_2, a_1)$

In this case, the fact that the sender is action-matching together with the assumed inequalities implies that

$$U_S(\theta_2, a_2) \stackrel{\text{AM}}{\geq} U_S(\theta_1, a_2) \geq U_S(\theta_1, a_1) \stackrel{\text{AM}}{\geq} U_S(\theta_2, a_1) \geq U_S(\theta_2, a_2).$$

Thus, the sender's utility is the same, regardless of the state and action. As a result, the sender is indifferent between all signaling schemes. In particular, fully revealing the state is an optimal strategy for the sender for any constraint; clearly, this would be best for the receiver.

Thus, for all four cases, the receiver will be no worse off under the more binding constraint. \square

3.1 Necessity of Partial Alignment

Our main Theorem 2 assumes that the sender is partially aligned with the receiver (in addition to the state space being binary). One may ask whether the partial alignment is necessary, or whether a more constrained receiver is *always* better off with binary state and action spaces. Here, we show that the assumption is necessary, by giving a 2×2 example under which the receiver is worse off when more constrained.

The sender's and receiver's utility functions are given in Table 1. Here, $0 < \epsilon \ll 1$. The prior distribution over states is $p = \frac{1}{4}$.

Table 1. Sender's and Receiver's Utility in the example without partial alignment

	θ_1	θ_2
a_1	2	3
a_2	1	0

(a) Sender's Utility

	θ_1	θ_2
a_1	1	ϵ
a_2	0	1

(b) Receiver's Utility

The sender prefers action a_1 in both states, and the receiver is state-matching. Notice that the sender is not action-matching: when the receiver plays a_1 , the sender prefers the state θ_2 over θ_1 . We write σ_i for the sender's signal suggesting action a_i , $i \in \{1, 2\}$.

We will compare the receiver's expected utilities in the following two settings:

1. There are (effectively) no constraints, i.e., $\bar{b}_{a_1} = 1, \underline{b}_{a_1} = 0, \bar{b}_{a_2} = 1, \underline{b}_{a_2} = 0$.
2. The constraint profile binds the sender-preferred action to at most its prior probability, i.e., $\bar{b}'_{a_1} = \frac{1}{4}, \underline{b}'_{a_1} = 0, \bar{b}'_{a_2} = 1, \underline{b}'_{a_2} = 0$.

The first setting is the classical Bayesian persuasion problem: the sender's optimal signaling strategy can be obtained by the concavification approach presented in [21], and is the following: Send σ_1 with $\phi_{1,1} = 1$ and $\phi_{2,1} = \frac{1}{3}$; send σ_2 with $\phi_{1,2} = 0$ and $\phi_{2,2} = \frac{2}{3}$. Given this commitment, the receiver's expected utility is $\frac{1+\epsilon}{2}$ when receiving σ_1 (because θ_1 and θ_2 are equally likely to occur), and her expected utility is 1 when receiving σ_2 . Thus, the receiver's overall expected utility is $\frac{3+\epsilon}{4}$.

In the second setting, the sender cannot send the signal σ_1 as frequently as in the unconstrained case. When the sender is forced to reduce $\mathbb{P}[\sigma_1]$, he prefers to reduce the probability $\phi_{1,1}$ instead of $\phi_{1,2}$. This is because $U_S(\theta_2, a_1) - U_S(\theta_2, a_2) > U_S(\theta_1, a_1) - U_S(\theta_1, a_2)$. However, reducing $\phi_{1,1}$ solely may cause the signal σ_1 to not be persuasive any more, when the posterior belief violates the incentive constraint. Hence, the sender's optimal signaling strategy requires him to maximize the total probability of σ_1 , under the constraint that the receiver is still willing to take action a_1 under σ_1 . Thus, the sender's optimal signaling scheme is the following: Send σ_1 with $\phi_{1,1} = \frac{1}{2}$ and $\phi_{2,1} = \frac{1}{6}$; send σ_2 with $\phi_{1,2} = \frac{1}{2}$ and $\phi_{2,2} = \frac{5}{6}$.

Against this signaling scheme, the receiver's best response to σ_1 is taking action a_1 , with an expected utility of $\frac{1+\epsilon}{2}$. Her best response to σ_2 is taking action a_2 , with an expected utility of $\frac{5}{6}$. Hence, the receiver's expected utility is $\frac{6+\epsilon}{8}$ under the constraints $\langle \underline{b}', \bar{b}' \rangle$.

In summary, the receiver's expected utility of $\frac{3+\epsilon}{4}$ in the first setting is higher than her utility of $\frac{6+\epsilon}{8}$ in the second setting. Thus, we have exhibited an example where a more constrained receiver is worse off than a less constrained one.

4 Failure of the Main Result with Larger State Spaces

Unfortunately, contrary to the case of binary state and action spaces, when the state and action spaces are larger, a state-matching receiver and action-matching sender (and feasible constraints) are not enough to ensure that the receiver is always better off when more constrained. Consider the utilities given in Table 2. There are three states in the world, and correspondingly three actions. The prior over the states is uniform.

Table 2. An example where a constrained receiver is worse off

	θ_1	θ_2	θ_3
a_1	10	10	0
a_2	0	2	2
a_3	0	0	1

(a) Sender's Utility

	θ_1	θ_2	θ_3
a_1	4	2	0
a_2	0	3	1
a_3	0	1	3

(b) Receiver's Utility

	θ_1	θ_2	θ_3
σ_1	1	$\frac{1}{2}$	0
σ_2	0	$\frac{1}{2}$	$\frac{1}{2}$
σ_3	0	0	$\frac{1}{2}$

(c) Sender-optimal signaling scheme when $\bar{b}_{a_1} = \frac{1}{2}$

Notice that the receiver is state-matching, and the sender is action-matching.

Unconstrained Receiver. First, consider an unconstrained receiver. The sender's optimal signaling scheme ϕ is to recommend action a_1 whenever the state of the world is θ_1 or θ_2 , and recommend action a_3 otherwise.

To verify that the receiver follows the recommendation, one simply compares the utility from the alternative actions: when the sender recommends a_1 , following the recommendation gives the receiver expected utility $\frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 2 = 3$, while a_2 would give utility $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 3 = \frac{3}{2}$, and a_3 would give $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 1 = \frac{1}{2}$. For the recommendation of a_3 , the receiver gets to match the state deterministically, so following the recommendation is optimal. Because the signaling scheme is even ex post incentive compatible for the receiver, it is most definitely ex ante incentive compatible.

To see that this signaling scheme is optimal for the sender, first observe that for states θ_1 and θ_2 , the sender obtains the maximum possible utility of 10 over all actions. For state θ_3 , the sender would prefer the receiver to play action a_2 . However, the only way to get the sender to play a_2 is to mix at least one unit of probability of θ_2 per unit of probability of θ_3 . While this increases the sender's utility for the unit of probability from θ_3 from 1 to 2, it decreases his utility for the unit of probability from θ_2 from 10 (since the receiver played a_1) to 2. Thus, the given signaling scheme is sender-optimal.

Under this signaling scheme, the receiver's expected utility can be calculated as $\frac{2}{3} \cdot (\frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 2) + \frac{1}{3} \cdot 3 = 3$.

Adding a Non-trivial Constraint. Now, consider a receiver constrained by an upper bound $\bar{b}_{a_1} = \frac{1}{2}$. Table 2c shows the sender-optimal signaling scheme. Here, the entries show the conditional probability $\phi_{i,j}$ of recommending action a_j (i.e., sending signal σ_j) when the state is θ_i .

First, notice that action a_1 is recommended with probability $\frac{1}{2}$, so the constraint is satisfied. Second, the receiver will follow the sender's recommendation, as can be checked by comparing her utility from each of the three actions conditioned on any signal. (In the case of receiving σ_2 , she is indifferent between a_2 and a_3 —recall that we assume tie breaking in favor of the sender.) Again, the given signaling scheme is even ex post incentive compatible, so in particular, it is also ex ante incentive compatible.

To see that the signaling scheme is optimal for the sender, first notice that he induces action a_1 (under states θ_1 or θ_2) with the maximum probability of $\frac{1}{2}$.

Also, notice that using all of the probability from θ_1 to induce a_1 is optimal for the sender, because under θ_1 , if any action other than a_1 is played, the sender's utility is 0. Because $\frac{1}{6}$ unit of probability from θ_2 yields a recommendation of a_1 , at most $\frac{1}{6}$ can yield a recommendation of a_2 , which gives the next-highest utility for the sender. And because the receiver will choose a_2 only when the conditional probability of θ_2 is at least as large as that of θ_3 , action a_2 is induced with the maximum possible probability of $\frac{1}{3}$. Inducing any other actions for any of the states would yield the sender utility 0. Hence, the given signaling scheme is optimal for the sender.

Under this signaling scheme, the receiver's expected utility is $\frac{1}{2} \cdot (\frac{2}{3} \cdot 4 + \frac{1}{3} \cdot 2) + \frac{1}{3}(\frac{1}{2} \cdot 3 + \frac{1}{2} \cdot 1) + \frac{1}{6} \cdot 3 = \frac{17}{6}$.

Thus, the constrained receiver's utility of $\frac{17}{6}$ is lower than the unconstrained receiver's of 3.

5 Discussion

We showed that a state-matching receiver, facing an action-matching sender under a binary state space, obtains weakly higher utility when more constrained. We believe that such behavior is in fact observed in the real world: for example, recommenders tend to be more careful in whom they nominate for particularly selective awards or positions.

5.1 Larger State/Action Spaces

As we discussed in Sect. 4, our results do not carry over to larger state spaces. Indeed, even for state spaces with three states, in which the receiver tries to minimize the distance between the action and the state of the world, there are counter-examples under which a constrained receiver is worse off.

While the result does not hold in full generality with three (or more) states, by imposing additional conditions, a positive result can be recovered:

Proposition 2. *Assume that the state space has size $|\Theta| = 3$, and that the receiver is state-matching and the sender is action-matching. In addition, assume that the following two conditions are satisfied.*

1. *The sender has a monotone³ preference over actions across all states, i.e., $U_S(\theta_i, a_1) \geq U_S(\theta_i, a_2) \geq U_S(\theta_i, a_3)$ for all i .*
2. *For every state i , the receiver is worse off choosing an action $j < i$ that is too low compared to choosing an action $k > i$ that is too high⁴: that is, $U_R(\theta_i, a_j) \leq U_R(\theta_i, a_k)$ for all $j < i < k$.*

Then, a more constrained receiver is never worse off than a less constrained one.

³ The result holds symmetrically if the order is reversed.

⁴ Notice that in the case $|\Theta| = 3$, this constraint only applies to $i = 2, j = 1, k = 3$.

We phrase it more generally to set the stage for a further generalization below.

The additional assumptions on the sender side capture a stronger version of the utility relationship of the interesting cases in the proof of Theorem 2. They are motivated in many of our cases: for instance, a letter writer may want to obtain the highest possible honor (or salary) for a student, or a prosecutor may want to maximize the sentence of a defendant.

The additional assumption on the receiver side would capture a cautious department or judge, who would prefer to err on the side of not inviting weak candidates (or giving awards to undeserving candidates), or giving the defendant a sentence that is too low rather than ever giving too high of a sentence.

While Proposition 2 shows that with enough assumptions, a positive result can be recovered, we believe that the assumptions are still rather restrictive, meaning that the proposition is likely of limited interest. The proof involves a long and tedious case distinction, and we therefore do not include it in the paper.

For fully general state spaces (i.e., $n = |\Theta| \geq 3$), we can currently obtain a positive result only by imposing even more assumptions on the utility functions. In addition to the (generalization of) the assumptions from Proposition 2, we can make the following assumptions: (1) Whenever $j < i$, the sender’s utility difference between actions $j < j'$ is larger under state θ_i than under state $\theta_{i'}$ for $i' > i$. In other words, when the state of the world is smaller, the sender is more sensitive to changes in the receiver’s action. (2) For any fixed state θ_i , the receiver’s utility as a function of j (the action) is increasing and *convex* for $j \leq i$, and decreasing and convex for $j \geq i$. By adding these two assumptions, we can again obtain a result that a constrained receiver is always weakly better off than an unconstrained one. While it is possible to construct reasonably natural applications which satisfy these conditions, the conditions are far from covering a broad class of Bayesian persuasion settings. For this reason, we are not including a proof of this result, instead considering the discussion as a point of departure towards identifying less stringent assumptions that may enable positive results.

Whether there is a broad and natural class of Bayesian persuasion instances with more than two states of the world in which the insight “A more constrained receiver is better off” from Theorem 2 carries over is an interesting direction for future research.

5.2 Finding Optimal Signaling Schemes

While the main focus of our work is on the receiver’s utility when more constrained, our model also raises an interesting computational question, as briefly discussed in Sect. 2.3. In particular, we do not know whether there is a polynomial-time algorithm which—given the sender’s and receiver’s utility functions as well as the constraints on the receiver—finds a sender-optimal signaling scheme. Since probability constraints on receivers (quotas) are quite natural in many signaling settings, this constitutes an interesting direction for future work.

The main difficulty in applying standard techniques is that the constraints may force the receiver to play an ex post suboptimal action. The standard LP for the sender’s optimization problem [14] maximizes the sender’s expected utility subject to the constraint that the receiver is incentivized to play the sender’s

recommended action. To appreciate the difference, consider a setting in which the state of the world is uniform over $\{\theta_1, \theta_2\}$, and the sender and receiver both obtain utility 1 if the receiver plays action a_1 , and 0 otherwise. Without any constraints, the sender need not send any signal, and the receiver would simply play action a_1 . But if the receiver is constrained to playing action a_1 with probability exactly $\frac{1}{2}$, then she must randomize, including the (always suboptimal) action a_2 with probability $\frac{1}{2}$. By Proposition 1, the randomization can be pushed to the sender instead, but when the sender recommend action a_2 , it will be ex post suboptimal for the receiver to follow the recommendation. Indeed, an LP requiring deterministic ex post obedience from the sender would become infeasible for this setting. Whether the sender's optimization problem can still be cast as a different LP, or solved using other techniques, is an interesting direction.

We remark here that the preceding example does not have a state-matching receiver. If the receiver is state-matching and the constraints are feasible, then full revelation of the state is ex post incentive compatible for the receiver. This implies that the linear program for optimizing the sender's utility over ex post incentive compatible signaling schemes has a feasible solution. However, since the LP is more restricted, it is not at all clear that its optimum solution maximizes the sender's utility when the recommendation does not have to be ex post incentive compatible.

5.3 Receiver's Strategic Behaviors on Constraint Enforcement

We assumed throughout the paper that the receiver's constraints are common knowledge, and that enforcing the constraints is indeed required of the receiver (or in her best interest). Aside from the interview example provided in Sect. 1, such constraints are encountered in real-world scenarios such as a patient's dietary restrictions, the salary cap for a sports team, or the capacity limit of an event or facility.

Given that we showed constraints to be beneficial for the receiver, one may suspect that a receiver could strategically misrepresent how harsh her constraints are, or—along the same lines—claim to be constrained, but not enforce the claimed constraints. This would allow the receiver to obtain more information from a sender. In other words, when constraints are not common knowledge, they become private information of the receiver, which could be strategically manipulated; for instance, in the interview example, Alice could indicate a constraint just to force Bob's hand.

Naturally, allowing strategic manipulation in the model will significantly complicate the problem, either making it a dynamic information design problem [15] with multiple senders [1] or a mechanism design problem with incorporated information design modules [32]. Analyzing a model with private receiver constraints thus constitutes an interesting directions for future work.

References

1. Ambrus, A., Takahashi, S.: Multi-sender cheap talk with restricted state spaces. *Theor. Econ.* **3**(1), 1–27 (2008)
2. Babichenko, Y., Talgam-Cohen, I., Zabarnyi, K.: Bayesian persuasion under ex ante and ex post constraints. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, no. 6, pp. 5127–5134 (2021)
3. Bergemann, D., Morris, S.: Robust predictions in games with incomplete information. *Econometrica* **81**(4), 1251–1308 (2013)
4. Bergemann, D., Morris, S.: The comparison of information structures in games: bayes correlated equilibrium and individual sufficiency. Technical report 1909R, Cowles Foundation for Research in Economics, Yale University (2014)
5. Bergemann, D., Morris, S.: Bayes correlated equilibrium and the comparison of information structures in games. *Theor. Econ.* **11**(2), 487–522 (2016)
6. Bergemann, D., Morris, S.: Information design: a unified perspective. *J. Econ. Lit.* **57**(1), 44–95 (2019)
7. Boleslavsky, R., Kim, K.: Bayesian persuasion and moral hazard (2018). Available at SSRN 2913669
8. Candogan, O.: Reduced form information design: persuading a privately informed receiver (2020). Available at SSRN 3533682
9. Candogan, O., Strack, P.: Optimal disclosure of information to a privately informed receiver (2021). Available at SSRN 3773326
10. Carroni, E., Ferrari, L., Pignataro, G.: Does costly persuasion signal quality? (2020). Available at SSRN
11. Crawford, V.P., Sobel, J.: Strategic information transmission. *Econometrica* **50**(6), 1431–1451 (1982)
12. Doval, L., Skreta, V.: Constrained information design: Toolkit (2018). arXiv preprint [arXiv:1811.03588](https://arxiv.org/abs/1811.03588)
13. Dughmi, S., Kempe, D., Qiang, R.: Persuasion with limited communication. In: Proceedings 17th ACM Conference on Economics and Computation, pp. 663–680 (2016)
14. Dughmi, S., Xu, H.: Algorithmic Bayesian persuasion. *SIAM J. Comput.* **50**(3), 68–97 (2019)
15. Farhadi, F., Teneketzis, D.: Dynamic information design: a simple problem on optimal sequential information disclosure. *Dyn. Games Appl.*, 1–42 (2021)
16. Gentzkow, M., Kamenica, E.: Costly persuasion. *Am. Econ. Rev.* **104**(5), 457–62 (2014)
17. Grossman, S.J.: The informational role of warranties and private disclosure about product quality. *J. Law Econ.* **24**(3), 461–483 (1981)
18. Guo, Y., Shmaya, E.: The interval structure of optimal disclosure. *Econometrica* **87**(2), 653–675 (2019)
19. Hedlund, J.: Persuasion with communication costs. *Games Econom. Behav.* **92**, 28–40 (2015)
20. Kamenica, E.: Bayesian persuasion and information design. *Ann. Rev. Economics.* **11**, 249–272 (2019)
21. Kamenica, E., Gentzkow, M.: Bayesian persuasion. *Am. Econ. Rev.* **101**(6), 2590–2615 (2011)
22. Kolotilin, A.: Experimental design to persuade. *Games Econom. Behav.* **90**, 215–226 (2015)

23. Kolotilin, A.: Optimal information disclosure: a linear programming approach. *Theor. Econ.* **13**(2), 607–635 (2018)
24. Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., Li, M.: Persuasion of a privately informed receiver. *Econometrica* **85**(6), 1949–1964 (2017)
25. Le Treust, M., Tomala, T.: Persuasion with limited communication capacity. *J. Econ. Theor.* **184**, 104940 (2019)
26. Matyskova, L.: Bayesian persuasion with costly information acquisition. In: cERGE-EI Working Paper Series (2018)
27. Milgrom, P.R.: Good news and bad news: representation theorems and applications. *Bell J. Econ.* **12**(2), 380–391 (1981)
28. Nguyen, A., Tan, T.Y.: Bayesian persuasion with costly messages. *J. Econ. Theor.* **193**, 105212 (2021)
29. Perez-Richet, E.: Interim Bayesian persuasion: First steps. *Am. Econ. Rev.* **104**(5), 469–74 (2014)
30. Perez-Richet, E., Prady, D.: Complicating to persuade? (2012, unpublished Manuscript)
31. Rayo, L., Segal, I.: Optimal information disclosure. *J. Polit. Econ.* **118**(5), 949–987 (2010)
32. Roesler, A.K.: Mechanism design with endogenous information. Technical report, Working Paper, Bonn Graduate School of Economics, Bonn, Germany (2014)
33. Spence, M.: Job market signaling. *Q. J. Econ.* **87**(3), 355–374 (1973)