

Information Design in Cheap Talk

Qianjun Lyu

University of Bonn

Wing Suen

University of Hong Kong

October 2, 2024

Abstract

An uninformed sender publicly commits to an informative experiment about an uncertain state, privately observes its outcome, and sends a cheap-talk message to a receiver. We provide an algorithm valid for arbitrary state-dependent preferences that will determine the sender's optimal experiment and his equilibrium payoff under binary state space. We give sufficient conditions for informative information transmission. These conditions depend more on marginal incentives—how payoffs vary with the state—than on the alignment of sender's and receiver's rankings over actions within a state. The algorithm can be easily modified to study the canonical cheap talk game with a perfectly informed sender.

Keywords: marginal incentives, common interest, concave envelope, quasiconcave envelope, double randomization

JEL Classification: D82, D83

This paper supersedes the project by the first author under the same title. We would like to thank Ricardo Alonso, Ettore Damiano, Francesc Dilme, Françoise Forges, Frédéric Koessler, Christoph Kuzmics, Stephan Lauer mann, Gilat Levy, Jin Li, Elliot Lipnowski, Francesco Nava, Matthew Mitchell, Xiaosheng Mu, Ronny Razin, Christopher Sandmann, Xianwen Shi, Dezső Szalay, Balázs Szentes, Jianrong Tian and Yimeng Zhang for helpful discussions at different stages of this research project. We would also like to thank seminar and conference participants at CETC, the Econometric Society meetings, the Parisian Seminar of Game Theory, 34th Stony Brook International Conference on Game Theory, Royal Economic Society Conference, University of Graz, and University of Bonn. Lyu acknowledges support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2126/1-390838866, and CRC TR 224 (Project B02). First version of this paper: 11 July 2022.

1. Introduction

Starting from Crawford and Sobel (1982), there is a large economics literature that studies how a biased sender can gain from strategic communication with an uninformed receiver. Much of this literature assumes that the sender is endowed with superior expertise. In many scenarios, however, the sender needs to learn about the payoff-relevant state before communicating with the receiver. For example, news media and think tanks that are biased for or against a political candidate or a government policy often collect information and conduct research in order to influence public opinion. Since the public may not have direct access to the data sources, nor the incentive to use time and effort to assess whether the conclusions drawn indeed follow from the original data, these conclusions effectively become cheap-talk messages. Similarly, financial institutions often have research departments whose works provide the basis for their portfolio recommendations to clients, but whether their investment advice is consistent with the findings of their research is often unverifiable. This paper studies optimal information acquisition when the sender cannot commit to communicating the outcome of his investigations in a verifiable way.

Specifically, we consider a strategic communication game where an uninformed sender publicly commits to an information structure (to learn about the state) and privately observes the information outcome before sending a cheap-talk message to a receiver, who then takes an action. One can interpret this game as a bridge between strategic communication (Crawford and Sobel, 1982) and Bayesian persuasion (Kamenica and Gentzkow, 2011), in the sense that the sender can commit to the information structure but not to truthful reporting.¹

We study this game with a binary state space and finite action space, while allowing the sender to have arbitrary state-dependent preferences. This means that his utility can depend on both the state and the receiver's action. State-dependent preferences create a tension between acquiring more information and alleviating conflict of interests. The first incentive is straightforward: acquiring more information enables the sender to make better use of the information. However, more information can also exacerbate the conflict, potentially intensifying the sender's incentive to misreport. Hence, when de-

¹ A canonical cheap talk game with a fully informed sender can be considered as a game where the sender has no commitment to both information structure and truthful reporting, because it is without loss of generality to assume that the sender will acquire perfect information in this case.

signing the optimal information structure, the sender needs to anticipate the credibility issues in the interim stage (when different information outcomes are realized). In general the sender may want to become only partially informed to alleviate the incentive compatibility problem.²

A generic feature of a model with discrete action space is that the receiver is indifferent between multiple actions at certain beliefs, even though the sender may not be indifferent over those actions. To enhance credibility, the sender sometimes benefits from the receiver’s randomization between his most and least preferred actions. Lipnowski and Ravid (2020) exploit this observation to the case where the sender has state-independent preferences (also referred to as transparent motives).³ They show that, under transparent motives, the receiver’s tie-breaking rules are determined by the sender’s indifference between reporting different messages. This further leads to their characterization that the sender’s highest equilibrium payoff is determined by the quasi-concave envelope of his indirect value function. However, in cases where the sender’s preference depends arbitrarily on the state, such tie-breaking rules are not necessary to guarantee incentive compatibility (as the sender’s preference can now vary across beliefs). Hence, a single geometric characterization, such as (quasi)-concavification, is insufficient to determine the highest equilibrium payoff for the sender.

Nevertheless, we show that with a discrete action space, the sender’s indirect value function remains sufficient to determine the optimal information structure and his highest equilibrium payoff. In Section 3, we present a finite algorithm to compute the optimal equilibrium outcome for the sender by searching the highest probability that the receiver can take the sender-preferred action without violating the sender’s incentive constraints. The optimal information structure generally induces two possible posterior beliefs, and the receiver may take pure or mixed strategies at each of these two beliefs. It turns out that for pure strategies, we can restrict the receiver to take the sender-preferred action. For mixed strategies, the mixing probability is determined by the sender’s indifference condition (in a specific way). Notably, “double randomization”—

² The result that an optimal information structure need not entail perfect information because of conflicts of interests is a recurrent theme in information economics; see, for example, Roesler and Szentes (2017) in the context of monopoly pricing and Bergemann et al. (2017) in the context of first-price auctions. Nevertheless, those papers assume full commitment, while in our model, the sender acquires only partial information to gain credibility.

³ Lipnowski and Ravid (2020) originally study a canonical cheap talk setting, where the sender privately observes the true state. Nevertheless, with state-independent preferences, the set of equilibrium outcomes in these two models coincide.

where the receiver takes mixed actions at both posterior beliefs—can be a part of the optimal information design. This contrasts with the case when the sender has transparent motives, where “double randomization” is never optimal.

To study the sender’s incentive compatibility, we call $m_s(a) = u_s(a, 1) - u_s(a, 0)$ the sender’s *marginal incentive* for action a . It is the difference in his utility of action a between state 1 and state 0. Graphically, the sender’s marginal incentives are the slopes of his piecewise indirect value function. Those slopes capture the marginal gain or marginal loss when the sender misreports his private beliefs, and therefore are crucial for incentive compatibility.

Intuitively one might expect that incentive issues are less severe if sender and receiver have “similar” preferences. However, the “similarity” we demand here is not captured by whether sender and receiver have the same ranking over actions given some particular state; but rather by the alignment of marginal incentives between them. We show in Section 4 that, when sender and receiver have opposite marginal incentives (i.e., if the receiver’s marginal incentives for action a is higher than a' , then the sender’s marginal incentives for action a' is higher than a), no information can be transmitted at all even if sender and receiver have identical ranking over actions in one state. Section 4 also provides the necessary and sufficient condition to guarantee informative information transmission when sender and receiver have aligned marginal incentives. Although it is possible that no incentive compatible information structure exists under aligned marginal incentives, this possibility only arises under rather stringent conditions. Randomization vastly expands the domain of preferences that would allow positive information transmission compared to a situation where the receiver only takes pure actions.

In Section 4.4, we discuss an application of our framework where sender and receiver have common interests in one of the states. That is, the sender’s optimal action in state 0 is also the receiver’s best response at that state. Then the optimal information structure generates a conclusive signal about state 0 under mild conditions. This application best disentangles the tension between acquiring more information and alleviating conflict of interests. Specifically, revealing state 0 (rather than pooling state 0 with state 1) generates more information for the sender and allows him to better use this information. On top of this, an experiment that identifies the common-interest state would align the two parties’ interests ex-post, which raises the sender’s ex-ante payoff.

This leads to an interesting result that, despite the sender and receiver having common interest in state 0, the optimal information structure does not necessarily reveal the true state with probability one when the true state is indeed of common interests.

As previously noted, our model represents a middle ground between Bayesian persuasion and canonical cheap talk. We hence explore its connections with Bayesian persuasion in Section 5 and with cheap talk in Section 6. Regarding Bayesian persuasion, the optimal experiment in our model can be more or less informative than that under Bayesian persuasion. Section 5 identifies a sensible set of preferences where the optimal experiment in our model is strictly more informative than the optimal experiment in Bayesian persuasion for certain prior beliefs. Then in Section 6, we link our model with the canonical cheap talk model where the sender is perfectly informed at the beginning. The equilibrium outcomes in the canonical cheap talk are a subset of the equilibrium outcomes in our model due to an additional constraint required for incentive compatibility there. Namely, the sender cannot gain by deviating to a more informative experiment than the one he commits to (in our model). This constraint allows us to simplify our algorithm further when searching for the sender-optimal equilibrium in the canonical cheap talk game. In general, the two algorithms can yield different solutions, suggesting that acquiring more information may indeed reduce the sender's equilibrium payoff.

Related literature. This paper describes a model of Bayesian persuasion with limited commitment, and is especially close to those papers in this literature that relax the commitment assumption at the communication stage. In Guo and Shmaya (2021) and Nguyen and Tan (2021), the sender cannot commit to reporting the true information outcomes but he incurs a cost of making incorrect claims. Alonso and Câmara (2024) allow the receiver to endogenously design an audit scheme, which in turn affects the sender's cost of misreporting. Lipnowski et al. (2022) discuss the situation where the sender can misreport the information outcomes at an exogenously given probability. In Krähmer (2021), the receiver can cross-check the sender's reports by privately randomizing over information structures. Regarding communication games with strategic information acquisition, Pei (2015) discusses a cheap talk game where the sender can acquire costly information that is unobserved by the receiver. Felgenhauer and Schulte (2014) consider a promotion game where the sender can privately and sequentially acquire signals generated from a binary experiment. Argenziano et al. (2016) allow the sender to choose the number of trials, which can be public information or the sender's

private information. In the latter two papers, though information cannot be falsified, its interpretation is subject to the sender’s disclosure policies. In contrast to these papers, we assume commitment on information structure and relax the commitment at the communication stage in the sense that the sender’s messages are pure cheap talk.

Our paper contributes to the literature on cheap talk with overt information acquisition. Ivanov (2010) investigates information design followed by cheap talk in a uniform-quadratic environment. He characterizes the optimal interval information structures. Deimen and Szalay (2019) consider a two-dimensional state space and the sender has access to a signal structure with elliptical distribution. In contemporaneous works, Kreutzkamp (2023) studies overt information acquisition with posterior-separable cost in a cheap talk model; Lou (2023) instead considers a costless environment with quadratic utilities. Both papers formulate their models as a linear persuasion problem with incentive constraints and utilize the recent developments on extreme points (Kleiner et al. (2021), Arieli et al. (2023)) to show the optimality of bi-pooling information structures. They characterize the optimal information structure under the uniform-quadratic setting. In contrast, we provide a complete characterization of the optimal information structure for arbitrary preferences within a finite model (binary states and finite actions). This cannot be otherwise accommodated with the existing results on linear persuasion.

Our paper is close to Lipnowski and Ravid (2020). They study the canonical cheap talk model where the sender has perfect private information *ex-ante*. They focus on situations where the sender’s preferences are state-independent, and find that the highest equilibrium payoff the sender can achieve is the quasiconcave envelope of his indirect value function. In contrast, our sender’s private information is endogenously determined by the information structure he commits to. In addition, we allow the sender to have arbitrary state-dependent preferences. It turns out that, when the sender has state-independent preferences, the equilibrium outcomes in these two models are equivalent. Therefore, the solution of our algorithm under state-independent preferences coincides with their characterization of the quasiconcave envelope. However, with state-dependent preferences, there is a fundamental difference between our model and the canonical cheap talk model as discussed in Section 6.⁴

Lin and Liu (2024) study the credibility of persuasion assuming that the sender’s

⁴ Other related papers are Lipnowski (2020) and Barros (2022). They provide conditions under which the optimal equilibrium outcome under cheap talk is equivalent to Bayesian persuasion.

deviation in messages is not detectable if the marginal distribution of messages remains the same. Their sender’s incentive constraints arrive at the ex-ante stage, in the sense that the gain from swapping messages in one state cannot outweigh the loss from that in another state.⁵ However, our sender’s incentive constraints arrive at the interim stage after the outcome of the experiment is privately revealed to the sender. The incentive constraints in these two papers are not nested. Salamanca (2021) studies a mediated communication game in which an informed sender sends a cheap talk message to a mediator, who can commit to a reporting rule based on the sender’s message. The receiver then takes an action based on the mediator’s report. Interestingly, under binary state space, our solution provides a lower bound to the sender’s highest achievable payoff in Salamanca (2021). The relationship is ambiguous for larger state space. We provide a more thorough discussion in Section 7.

Lastly, our paper contributes to the literature on algorithmic information design; see Dughmi and Xu (2016). In a recent work, Babichenko et al. (2023) discuss the algorithmic study of a canonical cheap talk game in a finite environment (finite states and finite actions). They show that with certain restrictions, e.g., when the cardinality of state space is constant, the computation will end in polynomial time. Instead, we study a different model with an additional layer of information acquisition. Unlike the classic algorithmic approach, our algorithm explicitly writes down the solution of the linear program, which turns out to have clean and simple geometric meanings.

2. The Model

A sender (S) and a receiver (R) initially share a common prior belief about some state θ . The state space $\Theta = \{0, 1\}$ is binary. We use $\mu \in \Delta\Theta$ to represent a probability distribution over the state, where $\mu(\theta)$ stands for the probability of state θ . The prior distribution about the state is denoted by μ_0 .

There is a finite set A of actions, with $|A| \geq 2$. We use a to represent a typical element of A , and use $\alpha \in \Delta A$ to represent a mixed action (i.e., a probability distribution over A). Each player $i \in \{S, R\}$ is an expected utility maximizer, whose utility $u_i(a, \theta)$ generally depends on both the action and the state. We assume that every action in A is the receiver’s unique best response to some belief. That is, for the receiver, no action is weakly dominated, and that no pure action is a duplicate of any alternative mixed

⁵ Lin and Liu (2024) focus on pure strategy equilibrium where the receiver cannot randomize.

action. The analysis in this paper can be suitably extended to handle situations when this assumption does not hold but at the cost of more clumsy notation.

The game consists of two stages. In the first stage, the sender commits to choosing a Blackwell experiment (a mapping from the state space to probability distributions over signals) and conducts the experiment at zero cost. As is standard in the Bayesian persuasion literature, this is equivalent to choosing a distribution of posterior beliefs induced by the experiment. In other words, the sender commits to a simple random posterior $P \in \Delta(\Delta\Theta)$ such that $\mathbb{E}_P[\mu] = \mu_0$, and P has a finite support.⁶ After the sender conducts the experiment, he privately observes the realization of the random posterior $\mu \in \text{supp}(P)$. We use $P(\mu)$ to denote the ex-ante probability that the experiment induces posterior μ for the sender (given the prior belief μ_0). The information structure chosen by the sender determines the distribution of his private information.

In the second stage, the sender interacts with the receiver in a game of strategic information transmission. Denote M as a rich finite message space. Given the random posterior P , the sender's reporting strategy, $\sigma_S : \text{supp}(P) \rightarrow \Delta M$, maps the realization of the random posterior to a distribution of messages. The receiver's decision rule, $\sigma_R : M \rightarrow \Delta A$, maps the sender's message to a distribution of actions. Each player i 's expected utility can be written as:

$$U_i(\sigma_S, \sigma_R, P) = \sum_{\mu \in \text{supp}(P), \theta \in \Theta, m \in M, a \in A} P(\mu)\mu(\theta)\sigma_S(m|\mu)\sigma_R(a|m)u_i(a, \theta).$$

In this framework the sender's posterior belief formation is trivial, and the receiver's posterior belief is obtained from P and σ_S using Bayes' rule. We focus on Perfect Bayesian Equilibrium, and call (σ_S, σ_R, P) an equilibrium strategy profile if σ_S and σ_R are mutual best responses given P and the belief system. The sender chooses the random posterior P to maximize his expected utility subject to an equilibrium. If there are multiple equilibria for a given P , we let the sender choose the one that gives him the highest expected utility.

Notice that each player's equilibrium payoff only depends on the joint distribution of the receiver's posterior belief and the action induced. Therefore, for every equilibrium such that the sender conceals information through a mixed reporting strategy, we can

⁶ See Denti et al. (2022). Because we are directly working with the random posterior induced by a Blackwell experiment, we implicitly assume, without loss of generality, that distinct signals induce different posterior beliefs.

find another truth-telling equilibrium where the sender directly coarsens the experiment in the first place and the equilibrium outcome remains the same. This is reminiscent of the revelation principle.

Lemma 1. *It is without loss of generality to focus on truth-telling equilibria and a random posterior with a binary support, i.e., $|\text{supp}(P)| = |\Theta| = 2$.*

The proof is provided in the Appendix.⁷ Because there are only two states, it is simpler to represent a probability distribution over the state by the probability of state 1. Henceforth, we use μ to stand for the probability of state 1. In addition, with binary states, a binary random posterior is completely pinned down by its support given a prior belief μ_0 .⁸ Therefore, we sometimes refer to a binary random posterior simply by its support.

With slight abuse of notation, let

$$u_i(a, \mu) := \mu u_i(a, 1) + (1 - \mu) u_i(a, 0)$$

be player i 's expected utility from action a when player i has posterior belief μ . Let

$$A_R(\mu) := \underset{a \in A}{\text{argmax}} u_R(a, \mu)$$

be the receiver's best-response correspondence, mapping from belief into a non-empty set of actions. We use $v(\mu) := \text{co}(u_S(A_R(\mu), \mu))$ to denote the sender's value correspondence given that both the sender and the receiver hold the same posterior belief μ and the receiver responds optimally to this belief. Finally, let

$$\bar{v}(\mu) := \max_{a \in A_R(\mu)} u_S(a, \mu)$$

be sender's value function when both sender and receiver hold the same belief μ and the receiver takes the sender-preferred action in his best response correspondence.

⁷ This result holds for any finite state space. In particular, the support of the optimal random posterior can have at most $|\Theta|$ elements.

⁸ For example, if $\text{supp}(P) = \{\mu', \mu''\}$, then the requirement that P is a mean-preserving spread of the prior belief μ_0 implies that μ' and μ'' are induced with probabilities $P(\mu')$ and $1 - P(\mu')$, where $P(\mu') = (\mu'' - \mu_0) / (\mu'' - \mu')$.

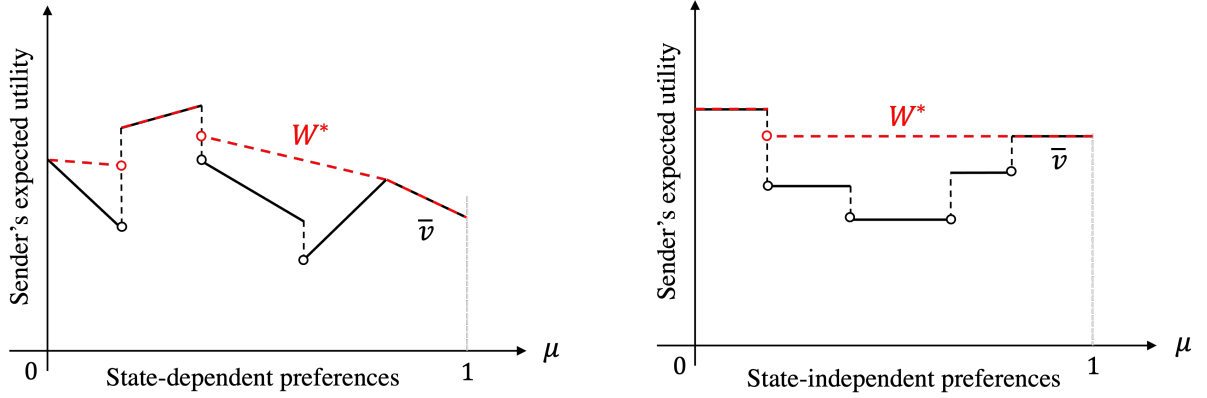


Figure 1: The sender's value function and his highest achievable payoff.

Given Lemma 1, the sender's information design problem can be written as:

$$\max_{P \in \Delta(\Delta\Theta), \sigma_R(a|\cdot) \in \Delta A_R(\cdot)} \sum_{\mu \in \text{supp } P} P(\mu) \sum_{a \in A_R(\mu)} \sigma_R(a|\mu) u_S(a, \mu),$$

subject to sender's incentive constraints: for every $\mu, \mu' \in \text{supp}(P)$,

$$\sum_{a \in A_R(\mu)} \sigma_R(a|\mu) u_S(a, \mu) \geq \sum_{a \in A_R(\mu')} \sigma_R(a|\mu') u_S(a, \mu), \quad (1)$$

and subject to the requirement that $|\text{supp}(P)| = 2$ and P is a mean-preserving spread of μ_0 . We denote $W^*(\mu_0)$ as the solution value to this program at the prior belief μ_0 .

Figure 1 give two examples of the sender's value function \bar{v} . The left panel refers to the case where the sender has state-dependent preferences (the piecewise slopes of \bar{v} are arbitrary). The right panel refers to the case where the sender has state-independent preferences (\bar{v} is piecewise constant). The red dashed curves W^* represent the highest equilibrium payoff the sender can achieve for each prior belief (we will elaborate the algorithm to determine W^* in the next section). The function W^* is piecewise affine.

If the sender with arbitrary preferences has full commitment power to truthfully report the outcome of the experiment, then the concave envelope of \bar{v} determines the highest equilibrium payoff the sender can achieve (Kamenica and Gentzkow, 2011). If the sender with state-independent preferences has no commitment power to truthful reporting, the quasiconcave envelope of \bar{v} determines the highest equilibrium payoff the sender can achieve (Lipnowski and Ravid, 2020). In our model, the sender has




Figure 2: The set of boundary beliefs.

arbitrary preferences and no commitment power. Therefore, $W^*(\cdot)$ is bounded above by the concave envelope of $\bar{v}(\cdot)$. The relationship between $W^*(\cdot)$ and the quasiconcave envelope of $\bar{v}(\cdot)$ is in general ambiguous (see the red curve in the left panel). We will elaborate more on this point later.

3. Optimal Information Design

Given the assumption that every element of A is a unique best response for the receiver at some belief, we have $|A_R(\mu)| \leq 2$ for all $\mu \in [0, 1]$. Moreover, we can order the actions in A in an increasing sequence, $\{a_{-J}, \dots, a_{-1}, a_0, a_1, \dots, a_K\}$, such that action a_n is receiver's best response on a closed interval of beliefs I_n , where the lowest belief in I_n is equal to the highest belief in I_{n-1} .⁹ Here, $a_0 = A_R(\mu_0)$ is the default action of the receiver when she has no information. Let μ_n be the point in the interval I_n which is closest to μ_0 . For completeness, also let $\mu_{K+1} = 1$ and $\mu_{-J-1} = 0$. We call $B := \{\mu_{-J-1}, \dots, \mu_{-1}, \mu_1, \dots, \mu_{K+1}\}$ the set of *boundary beliefs*. The notation adopted under this convention is illustrated by Figure 2. Elements of B are highlighted in red. We assume the prior μ_0 is in the interior of I_0 in the figure, but this is not important for our analysis.

Proposition 1. *For any prior belief, there exists an optimal binary random posterior whose support is a subset of the set of boundary beliefs.*

Proof. If $W^*(\mu_0) = \bar{v}(\mu_0)$, the random posterior with support $\{\mu_{-1}, \mu_1\}$ (which induces the default action a_0) is optimal. Suppose $W^*(\mu_0) > \bar{v}(\mu_0)$. Then there is an incentive compatible binary random posterior P with $\text{supp}(P) = \{\mu', \mu''\}$ that induces the receiver

⁹ Specifically, $I_n := \{\mu \in [0, 1] : a_n \in A_R(\mu)\}$.

to take different responses after different messages, i.e., $\sigma_R(a|\mu') \neq \sigma_R(a|\mu'')$. Suppose that at least one element of $\text{supp}(P)$ does not belong to B , say $\mu'' \in (\mu_k, \mu_{k+1})$. Then, the receiver takes pure action $\sigma_R(a|\mu'') = a_k$ after sender's message μ'' . Consider another random posterior P' with $\text{supp}(P') = \{\mu', \mu_{k+1}\}$, which is strictly more informative than P . The sender's incentive compatibility constraint implies that his utility difference from sending out the two messages (as a function of μ),

$$u_S(a_k, \mu) - \sum_a \sigma_R(a|\mu') u_S(a, \mu),$$

is nonpositive at $\mu = \mu'$ and nonnegative at $\mu = \mu''$. Since it is an affine function of μ , $\mu_{k+1} > \mu'' > \mu'$ implies that it is nonnegative at $\mu = \mu_{k+1}$, meaning that the sender's incentive compatibility constraint still holds at belief μ_{k+1} . Furthermore, the sender's payoff,

$$\max \left\{ u_S(a_k, \mu), \sum_a \sigma_R(a|\mu') u_S(a, \mu) \right\},$$

is convex in μ . Because P' is strictly more informative than P , his expected payoff is higher under P' (by Blackwell's theorem). A similar reasoning applies when μ' does not belong to B . \square

Proposition 1 suggests that a finite set of random posteriors is sufficient for determining the optimal information structure.¹⁰ It is driven by the observation that, for a given pair of actions, if a less informative information structure is incentive compatible, then the two parties' interests are aligned for each information outcome, which further implies that a more informative information structure is also incentive compatible and provides the sender with a higher expected utility conditional on that the more informative information structure induces the same pair of actions on path. Therefore, it is without loss of generality to consider the most informative information structure that can induce a given pair of actions. Every posterior belief induced by this information structure belongs to the set B . Henceforth, we can focus on binary random posterior P such that $\text{supp}(P) = \{\mu_{-j}, \mu_k\}$ for some j and k .

For a binary random posterior $\{\mu_{-j}, \mu_k\}$, use $\alpha_{-j} \in \Delta A_R(\mu_{-j})$ and $\alpha_k \in \Delta A_R(\mu_k)$ to

¹⁰ If babbling equilibrium is not optimal, then any equilibrium not supported under B can be strongly Pareto improved.

represent the mixed strategy taken after message μ_{-j} and μ_k , respectively. Let

$$\mathbb{E}_{\alpha_k}[u_S(a, \mu_k)] = \sum_a \alpha_k(a) u_S(a, \mu_k)$$

be the sender's expected utility if he has a posterior belief μ_k and the receiver takes the mixed strategy α_k , where $\alpha_k(a)$ stands for the probability of taking action a under the mixed strategy α_k . Define $\mathbb{E}_{\alpha_{-j}}[u_S(a, \mu_{-j})]$ similarly.

Starting with an initial belief $\mu \in (\mu_{-j}, \mu_k)$ (i.e., the expectation of the random posterior), the payoff from an experiment that generates posteriors μ_{-j} and μ_k and induces α_{-j} and α_k is:

$$W_{-j,k}(\mu; \alpha_{-j}, \alpha_k) := \frac{\mu_k - \mu}{\mu_k - \mu_{-j}} \mathbb{E}_{\alpha_{-j}}[u_S(a, \mu_{-j})] + \frac{\mu - \mu_{-j}}{\mu_k - \mu_{-j}} \mathbb{E}_{\alpha_k}[u_S(a, \mu_k)]. \quad (2)$$

This payoff is linear in μ with a constant derivative,

$$W'_{-j,k}(\cdot; \alpha_{-j}, \alpha_k) = \frac{\mathbb{E}_{\alpha_k}[u_S(a, \mu_k)] - \mathbb{E}_{\alpha_{-j}}[u_S(a, \mu_{-j})]}{\mu_k - \mu_{-j}}.$$

We provide a geometric illustration of this term in Figure 3.

If α puts probability one on an action $a \in A_R(\mu)$, then it represents a pure strategy. We sometimes replace α by a to emphasize the difference between a pure strategy and a mixed strategy.

To analyze incentive compatibility, we define the sender's *marginal incentive* corresponding to a pure action a as $m_S(a) := u_S(a, 1) - u_S(a, 0)$. More generally, the marginal incentive corresponding to a mixed strategy α is:

$$m_S(\alpha) = \mathbb{E}_\alpha[m(a)] = \sum_a \alpha(a) (u_S(a, 1) - u_S(a, 0)).$$

The sender's marginal incentives for pure actions are the slopes of his piecewise indirect value function. Hence, his marginal incentives for mixed actions are the weighted average of the slopes for the pure actions taken with positive probabilities. To generate credibility for a binary random posterior $\{\mu_{-j}, \mu_k\}$, both the slopes of $\mathbb{E}_{\alpha_{-j}}[u_S(a, \cdot)]$ and $\mathbb{E}_{\alpha_k}[u_S(a, \cdot)]$ and their values at μ_{-j} and μ_k matter. The next lemma provides a straightforward method for verifying incentive compatibility.

Lemma 2. *An information structure that generates posterior beliefs in $\{\mu_{-j}, \mu_k\}$, together with optimal actions $\alpha_{-j} \in \Delta A_R(\mu_{-j})$ and $\alpha_k \in \Delta A_R(\mu_k)$, satisfies sender's incentive compatibility constraints (1) if and only if*

$$m_S(\alpha_{-j}) \leq W'_{-j,k}(\cdot; \alpha_{-j}, \alpha_k) \leq m_S(\alpha_k). \quad (\text{IC})$$

Proof. Sender's payoff from inducing α_{-j} at belief μ_k is $\mathbb{E}_{\alpha_{-j}}[u_S(a, \mu_{-j})] + m_S(\alpha_{-j})(\mu_k - \mu_{-j})$. Incentive compatibility requires that this payoff be lower than $\mathbb{E}_{\alpha_k}[u_S(a, \mu_k)]$, which is sender's payoff from inducing α_k at belief μ_k . This is equivalent to $m_S(\alpha_{-j}) \leq W'_{-j,k}(\cdot; \alpha_{-j}, \alpha_k)$. The second inequality in (IC) follows similarly from the requirement that the sender has no incentive to induce α_k when his private belief is μ_{-j} . \square

Lemma 2 suggests a way to find the optimal information structure. For each binary random posterior $\{\mu_{-j}, \mu_k\}$, we first check condition (IC) for all pairs $(\alpha_{-j}, \alpha_k) \in \Delta A_R(\mu_{-j}) \times \Delta A_R(\mu_k)$, and select the incentive compatible pair with the highest value of $W_{-j,k}(\mu_0; \alpha_{-j}, \alpha_k)$. Optimizing over j and k would then give the highest achievable payoff $W^*(\mu_0)$ for the sender.

The difficulty is that there are infinitely many pairs (α_{-j}, α_k) . We identify the most relevant pairs that will guarantee a solution by searching over such pairs. For a random posterior P with support $\{\mu_{-j}, \mu_k\}$, there are three types of receiver's best response we need to consider.

Pure strategy (PP). Suppose the receiver takes a pure action after each message. There are four possible PP pairs because the receiver's best response at each boundary belief typically contains two elements. Only one pair of actions is necessary for the search. Let \bar{a}_{-j} be the sender-preferred action in $A_R(\mu_{-j})$ at belief μ_{-j} , and \underline{a}_{-j} be the remaining action (less preferred by the sender) in $A_R(\mu_{-j})$. Similarly, let \bar{a}_k be the sender-preferred action in $A_R(\mu_k)$ at belief μ_k , and \underline{a}_k be the remaining action in $A_R(\mu_k)$. If the sender is indifferent between actions in $A_R(\mu_{-j})$ at belief μ_{-j} , then we let $\bar{a}_{-j} = a_{-j+1}$; and if the sender is indifferent between actions in $A_R(\mu_k)$ at belief μ_k , we choose $\bar{a}_k = a_{k-1}$.¹¹

If inequality (IC) holds for $(\alpha_{-j}, \alpha_k) = (\bar{a}_{-j}, \bar{a}_k)$, we say that the random posterior P is "IC-PP" and we define $W_{-j,k}^{PP} := W_{-j,k}(\mu_0; \bar{a}_{-j}, \bar{a}_k)$.

¹¹ We break indifference in this way because then the random posterior with support $\{\mu_{-j}, \mu_k\}$ is the most informative information structure that can induce a_{-j+1} and a_{k-1} if the sender reports truthfully.




Figure 3: Incentive compatibility for pure strategy.

The random posterior P with support $\{\mu_{-j}, \mu_k\}$ in Figure 3 is IC-PP. To see this, note that the inequalities in (IC) are geometrically equivalent to the following: the slope of the left black piece $m_S(\bar{a}_{-j})$ is smaller than the slope of the middle orange piece $W'_{-j,k}(\cdot; \bar{a}_{-j}, \bar{a}_k)$, which is smaller than the slope of the right black piece $m_S(\bar{a}_k)$. By this sequence of inequalities, when $u_S(\bar{a}_{-j}, \cdot)$ (the left black piece) is extended to μ_k , its value is below $u_S(\bar{a}_k, \mu_k)$ (the black dot on the right). This indicates that the sender would not misreport μ_{-j} when his true belief is μ_k . Similarly, he has no incentive to misreport μ_k when his true belief is μ_{-j} .

One-sided randomization (PM or MP). Suppose the receiver takes a mixed strategy after one of the messages. Consider the case of PM (the MP case is symmetric), and consider the pair $(\alpha_{-j}, \alpha_k) = (\bar{a}_{-j}, \alpha_k^{PM})$, where α_k^{PM} puts weight γ_k on \bar{a}_k and weight $1 - \gamma_k$ on \underline{a}_k . The value of γ_k is determined by the requirement that, at belief μ_{-j} , the sender is indifferent between his preferred action \bar{a}_{-j} and the mixed action α_k^{PM} ,

$$u_S(\bar{a}_{-j}, \mu_{-j}) = \gamma_k u_S(\bar{a}_k, \mu_{-j}) + (1 - \gamma_k) u_S(\underline{a}_k, \mu_{-j}). \quad (3)$$

Notice that the indifference condition (3) determines the highest probability that the receiver can take the sender-preferred action \bar{a}_k at belief μ_k without violating the sender's incentive compatibility at belief μ_{-j} . By construction, the pair $(\bar{a}_{-j}, \alpha_k^{PM})$ satisfies

$$W'_{-j,k}(\cdot; \bar{a}_{-j}, \alpha_k^{PM}) = m_S(\alpha_k^{PM}),$$

i.e., the second inequality in (IC) holds with equality. If it also satisfies the first inequal-




Figure 4: Relaxing incentive constraints by randomization.

ity in (IC), and if α_k^{PM} is a valid mixed action,¹² we say that the random posterior P is “IC-PM,” and we define $W_{-j,k}^{PM} := W_{-j,k}(\mu_0; \bar{a}_{-j}, \alpha_k^{PM})$.¹³

The left panel of Figure 4 illustrates the geometric construction of one-sided randomization. First, draw an affine line connecting the left black dot $u_S(\bar{a}_{-j}, \mu_{-j})$ and the green dot on the right (that is the intersection point between the extended curves of $u_S(\bar{a}_k, \cdot)$ and $u_S(\underline{a}_k, \cdot)$). If this affine line (colored in orange) intersects with the sender’s value correspondence $v(\cdot)$ at belief μ_k , then the blue dot at that intersection represents the mixed action α_k^{PM} . In addition, we observe that the slope of the left black piece is smaller than the slope of the orange affine piece. Therefore, the value of this orange affine line at the prior belief μ_0 is the sender’s expected utility from the one-sided randomization that we identify.

Double randomization (MM). This involves the receiver taking a mixed strategy after each message. Let α_{-j}^{MM} be a mixed action that puts weight γ_{-j} on \bar{a}_{-j} and weight $1 - \gamma_{-j}$ on \underline{a}_{-j} . Let α_k^{MM} be a mixed action that puts weight γ_k on \bar{a}_k and weight $1 - \gamma_k$ on \underline{a}_k . The weights γ_{-j} and γ_k are chosen in such way that the sender is indifferent between α_{-j}^{MM} and α_k^{MM} both at belief μ_{-j} and at belief μ_k :

$$\mathbb{E}_{\alpha_{-j}^{MM}}[u_S(a, \mu_{-j})] = \mathbb{E}_{\alpha_k^{MM}}[u_S(a, \mu_{-j})], \quad \mathbb{E}_{\alpha_{-j}^{MM}}[u_S(a, \mu_k)] = \mathbb{E}_{\alpha_k^{MM}}[u_S(a, \mu_k)]. \quad (4)$$

¹² The value of γ_k that satisfies equation (3) may be outside $[0, 1]$, in which case α_k^{PM} is not a probability distribution.

¹³ If α_k^{PM} is not a valid probability distribution, we let $\mathbb{E}_{\alpha_k} [u_S(a, \mu_k)] := \gamma_k u_S(\bar{a}_k, \mu_{-j}) + (1 - \gamma_k) u_S(\underline{a}_k, \mu_{-j})$, given that γ_k is the solution to equation (3). The corresponding value of $W_{-j,k}^{MP}$ is defined accordingly. We adopt a similar convention for the cases of PM and MM.

The two indifference conditions (4) determine the highest probability that the receiver can take the sender-preferred action at both beliefs μ_{-j} and μ_k without violating the sender's incentive compatibility at both beliefs. As by construction, $(\alpha_{-j}^{MM}, \alpha_k^{MM})$ satisfies $m_S(\alpha_{-j}^{MM}) = W'_{-j,k}(\cdot; \alpha_{-j}^{MM}, \alpha_k^{MM}) = m_S(\alpha_k^{MM})$, and so the incentive constraints (IC) hold. If the value of (γ_{-j}, γ_k) that solves these two equations lies inside $[0, 1]^2$, then both α_{-j}^{MM} and α_k^{MM} are valid mixed actions.¹⁴ Then we say that the random posterior P is "IC-MM," and we define $W_{-j,k}^{MM} = W_{-j,k}(\mu_0; \alpha_{-j}^{MM}, \alpha_k^{MM})$.

The right panel of Figure 4 illustrates the construction of double randomization. We first draw an orange affine line connecting the intersection points of the two left black pieces and the two right black pieces. The sender's expected utilities $u_S(\alpha_{-j}^{MM}, \cdot)$, $u_S(\alpha_k^{MM}, \cdot)$, and his expected payoff from such double randomization $W_{-j,k}(\mu_0; \alpha_{-j}^{MM}, \alpha_k^{MM})$ all coincide on this (orange) affine line. In addition, if this (orange) affine line intersects with the sender's value correspondence $v(\cdot)$ at both μ_{-j} and μ_k (i.e., the two blue dots in Figure 4(b) lies in his value correspondence), then α_{-j}^{MM} and α_k^{MM} are valid mixed actions and the random posterior $\{\mu_{-j}, \mu_k\}$ is IC-MM.

Now we introduce an algorithm that yields the highest achievable payoff $W^*(\mu_0)$, together with an implied optimal random posterior P^* .

Algorithm 1:

1. For every pair $(-j, k) \in \{-J-1, \dots, -1\} \times \{1, \dots, K+1\}$, compute $W_{-j,k}(\mu_0; \bar{a}_{-j}, \bar{a}_k)$ and rank these values from highest to lowest.¹⁵ Starting from the pair with the highest value, verify whether it is IC-PP or not. Stop the first time an IC-PP pair is found. Assign $W^1 = W_{-j,k}^{PP}$ for such pair and let the set of $(-j, k)$ pairs with $W_{-j,k}^{PP}$ strictly higher than W^1 be S_1 . If there does not exist an IC-PP pair, assign $W^1 = \bar{v}(\mu_0)$ and let $S_1 = \{-J-1, \dots, -1\} \times \{1, \dots, K+1\}$,
2. For every pair $(-j, k)$ in S_1 :
 - (a) Compute $W_{-j,k}(\mu_0; \bar{a}_{-j}, \alpha_k^{PM})$ and re-rank these values from highest to lowest. Starting with the pair with the highest value, verify whether it is IC-PM or not. Stop the first time when an IC-PM pair is found. Assign $W^{(a)} = W_{-j,k}^{PM}$ for such pair and let the set of $(-j, k)$ pairs with $W_{-j,k}^{PM}$ strictly higher than $W^{(a)}$ be $S^{(a)}$. If none of them is IC-PM, assign $W^{(a)} = \bar{v}(\mu_0)$ and $S^{(a)} = S_1$

¹⁴ Whenever γ_j and γ_k are not uniquely pinned down by the indifference conditions (e.g., when all four actions share the same marginal incentives), we pick the values that maximize the sender's ex ante payoff.

¹⁵ It is not important how we break ties.

- (b) Go through a symmetric procedure in the case of MP. Assign $W^{(b)} = W_{-j,k}^{MP}$ the first time an IC-MP pair is found and let the set of $(-j, k)$ pairs with $W_{-j,k}^{MP}$ strictly higher than $W^{(b)}$ be $S^{(b)}$. If none of them is IC-PM, assign $W^{(b)} = \bar{v}(\mu_0)$ and $S^{(b)} = S_1$.
- (c) Let $W^2 = \max\{W^{(a)}, W^{(b)}\}$. Let $S_2 = S^{(a)} \cup S^{(b)}$.
3. For every pair $(-j, k)$ in S_2 , compute $W_{-j,k}(\mu_0; \alpha_{-j}^{MM}, \alpha_k^{MM})$ and re-rank these values from the highest to lowest. Starting with the pair with the highest value, verify whether it is IC-MM or not. Stop the first time an IC-MM pair is found and assign $W^3 = W_{-j,k}^{MM}$ for such pair. If none of them is IC-MM, assign $W^3 = \bar{v}(\mu_0)$.
4. Assign $W^*(\mu_0) = \max\{W^1, W^2, W^3\}$. The random posterior with support $\{\mu_{-j}, \mu_k\}$ corresponding to the $(-j, k)$ pair that yields $W^*(\mu_0)$ is optimal.

Theorem 1. *Algorithm 1 determines the highest achievable payoff for the sender.*

Theorem 1 relies on the fact that for any pair of beliefs (μ_{-j}, μ_k) , it is sufficient to consider only four possibilities, namely IC-PP, IC-PM, IC-MP, and IC-MM. In principle the set of incentive-compatible action pairs (α_{-j}, α_k) at these beliefs is infinite. However if we represent α_{-j} by the weight $\gamma_{-j} \in [0, 1]$ it puts on the sender-preferred action \bar{a}_{-j} , and represent α_k by the weight $\gamma_k \in [0, 1]$ it puts on the sender-preferred action \bar{a}_k , the incentive-compatibility constraints (IC) are linear in (γ_{-j}, γ_k) . Since the sender's payoff is also linear (and increasing) in (γ_{-j}, γ_k) , the corresponding action pair that achieves the highest payoff for the sender must be at one of the extreme points of this set. In the case of IC-PP, $(1, 1)$ is an extreme point and it must give the highest payoff to the sender. In the case of IC-PM, $(1, \gamma_k^{PM})$ is an extreme point of the same, with γ_k^{PM} given by the solution to equation (3); and similarly for $(\gamma_{-j}^{MP}, 1)$ for IC-MP. If an action pair is IC-MM, the solution to equation (4) gives another extreme point of the set. The sufficiency of Algorithm 1 indicates that only these four extreme points are exposed. Therefore, the number of steps of the procedure is bounded above by $|A|^2 + 4|A|$. Nevertheless, the procedure we describe in Algorithm 1 guarantees a faster search without checking all possibilities across all $(-j, k)$.

To find the highest equilibrium payoff across different prior beliefs, we would in principle run Algorithm 1 for every prior belief μ_0 . However, it is unnecessary given the linearity of the problem. We only need to apply the algorithm again when the prior belief crosses a boundary belief, i.e., when the set of $(-j, k)$ satisfying Bayesian plausibility changes.

The construction behind this algorithm generalizes Lipnowski and Ravid (2020) under the case of a binary state with arbitrary preferences. When the sender has state-independent preferences (transparent motives), the marginal incentive $m_S(\alpha)$ is equal to 0 for every mixed action α (including pure action). The incentive compatibility requirement (IC) in Lemma 1 would then require $W'_{-j,k}(\cdot; \alpha_{-j}, \alpha_k) = 0$ for any action pair. This implies that to find the sender's highest achievable payoff, we can search for the highest piecewise step functions such that every endpoint of a piece is inside the sender's value correspondence. This leads to the quasiconcave envelope of $\bar{v}(\cdot)$. In our setup, the fact that $m_S(\alpha_{-j})$ is in general different from $m_S(\alpha_k)$ means that $W'_{-j,k}(\cdot; \alpha_{-j}, \alpha_k)$ is not restricted to be equal to 0. The sender in our setup can achieve a payoff greater than or less than the quasiconcave envelope of $\bar{v}(\cdot)$.

The use of randomization to relax incentive compatibility constraints is emphasized in Lipnowski and Ravid (2020). Nevertheless, double randomization is never optimal under transparent motives. To see this, if the sender is recommending mixed actions α_{-j} and α_k at beliefs μ_{-j} and μ_k , he could strictly raise his payoff by putting more weight on \bar{a}_{-j} and \bar{a}_k in these mixed actions, provided that the new pair of mixed actions are still incentive compatible. Such deviation is always feasible as long as marginal incentives $m_S(\cdot)$ are equal for all actions. In our model with general preferences, such deviation may not be feasible, and therefore double randomization can remain a candidate as part of optimal information design.

Algorithm 1 determines the sender-preferred equilibrium of this model. Another natural question is to study the receiver-preferred equilibrium. Interestingly, such equilibrium is much easier to compute because the receiver's indirect value function is convex. The process begins with the lowest boundary belief μ_{-J-1} . Identify the largest μ_k where $k \in \{1, \dots, K+1\}$ such that $\{\mu_{-J-1}, \mu_k\}$ can form an incentive compatible pair of posterior beliefs. Denote this index by k_1 and compute the receiver's equilibrium payoff $W^R(-J-1) = W^R_{-J-1, k_1}(\mu_0; a_{-J-1}, a_{k_1})$, where $W^R_{-j,k}(\cdot)$ is computed similarly to $W_{-j,k}(\cdot)$ in equation (2) but the sender's utility function $u_S(\cdot)$ is replaced by the receiver's utility function $u_R(\cdot)$. If no such k_1 exists, meaning no incentive-compatible μ_k can be found, the receiver's payoff is set to $W^R(-J-1) = u_R(a_0, \mu_0)$. We proceed to the next step if $k_1 \neq K+1$. Otherwise, the receiver's highest equilibrium payoff is $W^R(-J-1)$. The next step considers the second highest boundary belief μ_{-J} . Identify the largest μ_k from the reduced set $k \in \{k_1+1, \dots, K+1\}$ that remains incentive compatible with μ_{-J} . Let the corresponding index be k_2 . Compute $W^R(-J) = W^R_{-J, k_2}(\mu_0; a_{-J-1}, a_{k_2})$; and if no

k_2 can be found, we set the receiver's payoff to $W^R(-J) = u_R(a_0, \mu_0)$. Similarly, we proceed further if $k_2 \neq K + 1$. Otherwise, the receiver's highest equilibrium payoff is $\max\{W^R(-J-1), W^R(-J)\}$. This procedure continues maximally for $J-1$ steps, progressing through the sequence of boundary beliefs $\mu_{-J-1}, \mu_{-J}, \dots, \mu_{-1}$. The highest equilibrium payoff the receiver can achieve is the maximum value across all computed payoffs, $\max\{W^R(-J-1), \dots, W^R(-1)\}$.

4. Positive Information Transmission

In this section, we explore the existence of informative information transmission in our model. Algorithm 1 in Section 3 determines the sender's maximum payoff $W^*(\mu_0)$ under an optimal information structure. We say that there is *positive information transmission* if $W^*(\mu_0) > \bar{v}(\mu_0)$ for some prior belief $\mu_0 \in [0, 1]$. Otherwise, if $W^*(\mu_0) = \bar{v}(\mu_0)$ for all $\mu_0 \in [0, 1]$, no information can be transmitted on path.

In Kamenica and Gentzkow (2011) or Lipnowski and Ravid (2020), information transmission is positive if and only if the sender's indirect value function $\bar{v}(\cdot)$ is not concave or not quasiconcave, respectively.¹⁶ In contrast, in our model, whether information transmission is positive depends less on the concavity properties of $\bar{v}(\cdot)$ and more on the structure of marginal incentives $m_S(\cdot)$. There is no straightforward way to characterize the necessary and sufficient condition for positive information transmission. In the following, we provide some economically meaningful conditions that will settle this question.

We introduce the following concepts that relate to the conflict of interest between sender and receiver.

Definition 1. Sender and receiver have *opposite marginal incentives* if, for any $a', a'' \in A$,

$$m_R(a') < m_R(a'') \iff m_S(a') > m_S(a'').$$

They have *aligned marginal incentives* if, for any $a', a'' \in A$,

$$m_R(a') < m_R(a'') \iff m_S(a') < m_S(a'').$$

¹⁶ In a model with discrete action space, the sender's value function $\bar{v}(\cdot)$ is (generically) discontinuous at beliefs for which the receiver is indifferent between different actions. Since a discontinuous function is not concave, information transmission is always positive according to our definition when there is full commitment.

The notion of opposite or aligned marginal incentives has little to do with comparing the level (or the ranking) of utilities attached to different actions at a given belief by the receiver and by the sender. For example, sender and receiver may have identical preference ranking over actions in A if they know the true state is, say, state 0; yet they may still have opposite marginal incentives according to Definition 1.

Our definition is related to supermodularity or submodularity between action and state. With a binary state space, it is without loss of generality to assume that the receiver preferences are supermodular in (a, θ) (because we order actions in such a way that higher actions are chosen at higher beliefs). According to this convention, if $u_S(\cdot, \cdot)$ is strictly submodular, then sender and receiver have opposite marginal incentives. If $u_S(\cdot, \cdot)$ is strictly supermodular, they have aligned marginal incentives.

4.1. Opposite marginal incentives

Proposition 2. *If sender and receiver have opposite marginal incentives, then no information can be transmitted.*

Proof. Consider an arbitrary prior belief $\mu_0 \in (0, 1)$. Take any pair of boundary beliefs such that $\mu_{-j} < \mu_0 < \mu_k$, and take any arbitrary receiver's best responses $\alpha_{-j} \in \Delta A_R(\mu_{-j})$ and $\alpha_k \in \Delta A_R(\mu_k)$, with $\alpha_{-j} \neq \alpha_k$. Our convention of ordering actions implies that $m_R(\alpha_{-j}) < m_R(\alpha_k)$, and hence $m_S(\alpha_{-j}) > m_S(\alpha_k)$. By Lemma 1, this pair of actions (α_{-j}, α_k) cannot be incentive compatible. This means that there does not exist an incentive compatible (binary) random posterior that can induce different actions at different interim beliefs. In other words, in any equilibrium, the receiver's action does not depend on the sender's messages. That is, no information can ever be transmitted. \square

Proposition 2 is valid regardless of whether sender and receiver have the same or different rankings over the set of actions in the two states. As long as their marginal incentives are opposite, information cannot be transmitted.¹⁷ Figure 5 shows one such example. The sender's value function $\bar{v}(\cdot)$ in this figure is obviously not concave. Nevertheless, because the left-derivative of $\bar{v}(\mu)$ is decreasing in μ , Proposition 2 implies that,

¹⁷ The logic behind this result is similar to Lin and Liu (2024) in their model of credible persuasion. With opposite marginal incentives, the sender's overall net gain from swapping messages in both states is always positive. Therefore they cannot generate incentive compatibility in their model.




Figure 5: No information transmission.

despite the sender's power to commit to an information structure, this cannot improve his payoff from a babbling equilibrium for any prior belief.

4.2. Aligned marginal incentives

Now, we turn to the case where sender and receiver have aligned marginal incentives.

Definition 2. An action $a' \in A$ blocks $a'' \in A$ if

$$u_S(a', \mu'') > u_S(a'', \mu'') \quad \text{for all } \mu'' \in \{\mu : a'' \in A_R(\mu)\}.$$

If a' does not block a'' and a'' does not block a' , then the incentive compatibility constraints (1) can be satisfied and there is an IC-PP information structure at some initial belief that will induce these two actions. Convexity of $\max\{u_S(a', \cdot), u_S(a'', \cdot)\}$ implies that such information structure would strictly improve the sender's payoff if the initial belief μ_0 is, say, in the interior of $\{\mu : \mu \in A_R(a')\}$. Thus the existence of a pair of pure actions that does not block one another is a sufficient condition for positive information transmission. This condition, however, is far from necessary.

To state the necessary and sufficient condition for positive information transmission under aligned marginal incentives, we order the actions in such a way that $a_1 < \dots < a_n$. Let the interval of beliefs for which action a_i (for $i = 1, \dots, n$) is optimal for the receiver be $[\underline{I}_i, \bar{I}_i]$, with $\underline{I}_1 = 0$, $\bar{I}_n = 1$, and $\bar{I}_i = \underline{I}_{i+1}$.

Proposition 3. Suppose sender and receiver have aligned marginal incentives. Then there is positive information transmission if and only if the following two conditions hold:

- (1) there exists $i < j$ such that a_i does not block a higher action a_j ; and

(2) there exists $i' < j'$ such that $a_{j'}$ does not block a lower action $a_{i'}$.

Proof. We prove necessity first. Suppose condition (1) fails. Then, for each $i = 1, \dots, n-1$, a_i blocks all higher actions $a_j > a_i$. This implies that, for an arbitrary a_j and any $\mu \in [\underline{I}_j, \bar{I}_j]$, we have $\mathbb{E}_{\alpha_i}[u_S(a, \mu)] > \mathbb{E}_{\alpha_j}[u_S(a, \mu)]$ for all optimal action $\alpha_i \in \Delta\{a_{i-1}, a_i\}$ and $\alpha_j \in \Delta\{a_{j-1}, a_j\}$ with $i < j$ (including the case $i = j - 1$). Hence, there does not exist an incentive compatible pair of (mixed) actions. The argument for the necessity of condition (2) is symmetric.

For sufficiency, consider a pair of actions $a_i < a_j$ such that a_i does not block a_j . Condition (1) ensures that such a pair exists. There are three possibilities:

- (a) a_j does not block a_i . Since neither action blocks the other, the pair a_i and a_j are IC-PP and positive information transmission is possible.
- (b) a_j blocks a_i , and a_j does not block some lower action. Then both the set $NBL(j) := \{k < j : a_j \text{ does not block } a_k\}$ and the set $BL(j) := \{k < j : a_j \text{ blocks } a_k\}$ are nonempty (because i belongs to $BL(j)$), and they form a partition of the set $\{1, \dots, j - 1\}$. Therefore there exists $k^* < j$ such that either (i) $k^* \in NBL(j)$ and $k^* - 1 \in BL(j)$, or (ii) $k^* \in BL(j)$ and $k^* - 1 \in NBL(j)$. In case (i), if a_{k^*} does not block a_j , then a_{k^*} and a_j form an IC-PP pair. Otherwise, if a_{k^*} blocks a_j , then $u_S(a_{k^*}, \bar{I}_j) > u_S(a_j, \bar{I}_j) > u_S(a_{k^*-1}, \bar{I}_j)$. Thus, given aligned marginal incentives, there exists a mixed action in $\Delta\{a_{k^*-1}, a_{k^*}\}$ such that it is IC-MP with action a_j . Case (ii) is similar.
- (c) a_j blocks a_i , and a_j blocks all lower actions. In this case, condition (2) requires that there exists $l \neq j$ (and $l > 1$) such that a_l does not block all lower actions. Let k be the smallest integer such that a_l does not block a_k . In other words, $k = \min NBL(l)$. There are two subcases:

- (i) $k < j$. If a_k does not block a_l , then a_k and a_l form an IC-PP pair, and positive information transmission is possible. If a_k blocks a_l , then the set $BH(k) := \{m > k : a_k \text{ blocks } a_m\}$ is nonempty because $l \in BH(k)$. Moreover, since a_j blocks all lower actions under case (c), together with aligned marginal incentives, we have a_k does not block a_j . Thus, the set $NBH(k) := \{m > k : a_k \text{ does not block } a_m\}$ is nonempty because $j \in NBH(k)$. Since $BH(k)$ and $NBH(k)$ form a partition of $\{k + 1, \dots, n\}$, there exist two adjacent actions in $\{a_{k+1}, \dots, a_n\}$ such that one action belongs to $BH(k)$ and the other belongs to $NBH(k)$. Thus, with a similar argument as in (b), there is

positive information transmission.

- (ii) $k \geq j$. The fact that $\min NBL(l) = k \geq j > i$ implies that $i \in BL(l)$. Thus both $NBL(l)$ and $BL(l)$ are nonempty. With a similar argument as in (b), there is positive information transmission. \square

Conditions (1) and (2) of Proposition 3 for positive information transmission under aligned marginal incentives are not very restrictive. They only require one action that does not block a higher action and one that does not block a lower action, where the pairs of actions involved need not be related.

Consider, for example, the special case where the sender's preferences are *ordinally state-independent*, in the sense that sender has identical rankings over actions in both state 0 and state 1. In this case, no action pairs are IC-PP. However, out of the $n!$ possible ordinally state-independent rankings over n actions, only two (the rankings where a_i is above a_j if and only if $i > j$ or $i < j$) violate conditions (1) and (2). The remaining $n! - 2$ rankings allow for positive information transmission under aligned marginal incentives. As shown in the proof of Proposition 3, this is because IC-PM or IC-MP pairs of actions can be constructed even when no action pair is IC-PP.

4.3. Neither opposite nor aligned marginal incentives

The definitions of opposite and aligned marginal incentives require the same ordering of marginal incentives for all actions, i.e., with opposite (aligned) marginal incentives, the slope of the sender's piecewise indirect value function is decreasing (increasing). Nevertheless, having a monotone ordering of marginal incentives for all actions is restrictive, and is not necessary for generating credibility. The role of the receiver's randomization is to smooth the sender's marginal incentives and thereby generate the correct ordering needed for incentive compatibility. We illustrate this with an example.

An example. A policy maker is seeking public support to change the status quo from no carbon tax (a_0) to a policy of either a small carbon tax (a_1) or high carbon tax (a_2). There are two states: low risk of climate change (state 0), and high risk of climate change (state 1). In state 0, the public's preference is $a_0 \succ a_1 \succ a_2$. In state 1, the public's preference is reversed: $a_2 \succ a_1 \succ a_0$. At prior belief μ_0 , the public prefers to remain at the status quo. The policy maker's preference is $a_1 \succ a_0 \succ a_2$ in state 0, and is $a_1 \succ a_2 \succ a_0$ in state 1, indicating that the policy maker and the public have aligned interests for extreme actions a_0 and a_2 , which implies $m_S(a_0) < m_S(a_2)$. Low




Figure 6: Neither opposite nor aligned marginal incentives. Opposite marginal incentives between a_0 and a_1 ; aligned marginal incentives between a_0 and a_2 .

carbon tax (a_1) is the most preferred action for the policy maker because he finds it most politically expedient. However, he also considers that small taxation has a limited impact on greenhouse gas emissions and may foster a false sense of accomplishment, potentially postponing the public's pursuit of alternative measures, which would be especially harmful if climate risk is high. This consideration suggests that $m_s(a_1) < m_s(a_0)$. Thus, the policy maker and the public exhibit opposite marginal incentives for a_0 and a_1 but aligned marginal incentives for a_0 and a_2 . The policy maker can hire a research group to conduct specific research to learn about the state. However, the outcomes of this research are unverifiable and, therefore, considered as cheap talk messages to the public.

Figure 6 describes the policy maker's indirect value function in this scenario. He achieves the highest equilibrium payoff by designing the optimal random posterior $\{0, \mu_2\}$, and the receiver randomizes between the two taxation magnitudes a_1 and a_2 upon receiving message μ_2 . This randomization between the sender's most-preferred action a_1 and the least-preferred action a_2 at belief 0 helps mitigate the sender's incentive to misreport a higher belief μ_2 when his true private belief is 0. In addition, this randomization smooths the sender's marginal incentive, and thereby restores a proper ordering of marginal incentives (i.e., $m_s(a_0) < W'_{-1,2}(\cdot; a_0, \alpha_2^{PM}) = m_s(\alpha_2^{PM})$) to satisfy the incentive compatibility at belief μ_2 . Overall, with state-dependent preferences, the receiver's randomization has two roles: one is to average the utilities, and the other is to average the marginal incentives.

4.4. One-sided common interest

In many situations, the sender and the receiver may have common interests in one state but conflicting interests in another state. By this, we mean that the receiver's optimal action in one state is also the sender's most-preferred action in that state (their rankings over other actions in that state can be different).

Definition 3. Sender and receiver have *common interest in one state* if, for $\theta = 0$ or $\theta = 1$,

$$u_S(a, \theta) \geq u_S(a', \theta) \quad \text{for all } a \in A_R(\theta) \text{ and all } a' \in A.$$

With common-interest in one state, we can disentangle the sender's trade-off between acquiring more information and alleviating the conflicts of interest. On the information side, the sender may want to reveal more information about the common interest state—instead of pooling the common-interest state with the other state—so that he can make the correct recommendation more often. On the side of conflicts of interest, since sender and receiver prefer the same action under the common-interest state, revealing it can further increase the sender's ex-post payoff in that common-interest state and thereby on average increase the sender's ex-ante payoff. The proposition below confirms this intuition. It shows that under some mild conditions, the optimal information structure generates a conclusive signal on the common-interest state.

Proposition 4. *Let the common-interest state be state 0, and let the optimal action corresponding to that state be a_{-j} . If there exists an action $a_k \in \{A_R(\mu) : \mu \in (\mu_0, 1]\}$ such that a_{-j} does not block a_k , then $0 \in \text{supp } P^*$.*

Proof. Since a_{-j} does not block a_k , we have $u_S(a_k, \mu_{k+1}) \geq u_S(a_{-j}, \mu_{k+1})$. Let \bar{a}_{k+1} be the sender-preferred action in $A_R(\mu_{k+1})$. Then $u_S(\bar{a}_{k+1}, \mu_{k+1}) \geq u_S(a_k, \mu_{k+1}) \geq u_S(a_{-j}, \mu_{k+1})$. From the definition of common interest in state 0, $u_S(a_{-j}, 0) \geq u_S(\bar{a}_{k+1}, 0)$. Therefore, the random posterior P with support $\{0, \mu_{k+1}\}$ is IC-PP if the receiver optimally chooses between a_{-j} and \bar{a}_{k+1} . Consider another incentive compatible random posterior P' with support $\{\mu_{-j}, \mu_{k'}\}$ that induces $\alpha_{-j} \in A_R(\mu_{-j})$ and $\alpha_{k'} \in A_R(\mu_{k'})$. There are two possibilities.

Case (1) $\mu_{k'} = \mu_{k+1}$. Since P' is incentive compatible, the payoff from this informa-

tion structure is

$$\begin{aligned}
W_{-j,k'}(\mu_0; \alpha_{-j}, \alpha_{k'}) &= \mathbb{E}_{P'} \left[\max \left\{ \mathbb{E}_{\alpha_{-j}} [u_S(a, \mu)], \mathbb{E}_{\alpha_{k'}} [u_S(a, \mu)] \right\} \right] \\
&\leq \mathbb{E}_P \left[\max \left\{ \mathbb{E}_{\alpha_{-j}} [u_S(a, \mu)], \mathbb{E}_{\alpha_{k'}} [u_S(a, \mu)] \right\} \right] \\
&\leq \mathbb{E}_P \left[\max \left\{ u_S(a_{-j}, \mu), \mathbb{E}_{\alpha_{k'}} [u_S(a, \mu)] \right\} \right] \\
&\leq \mathbb{E}_P \left[\max \left\{ u_S(a_{-j}, \mu), u_S(\bar{a}_{k+1}, \mu) \right\} \right] \\
&= W_{-j,k+1}(\mu_0; a_{-j}, \bar{a}_{k+1}).
\end{aligned}$$

The first inequality follows from the fact that P is a mean-preserving spread of P' ; therefore there is positive information value when the receiver's action space is fixed: belief μ_{k+1} is realized more often under P and the sender can correctly recommend $\alpha_{k'}$ instead of α_{-j} at μ_{k+1} . The second inequality follows from common-interest at state 0, $\mathbb{E}_{\alpha_{-j}} [u_S(a, 0)] \leq u_S(a_{-j}, 0)$ —revealing state 0 increases the sender's payoff as the receiver will take a more favorable action when state 0 is realized. The third inequality comes from $\mathbb{E}_{\alpha_{k'}} [u_S(a, \mu_{k+1})] \leq u_S(\bar{a}_{k+1}, \mu_{k+1})$. The last equality comes from the fact that the random posterior with support $\{0, \mu_{k+1}\}$ is IC-PP for a_{-j} and \bar{a}_{k+1} .

Case (2) $\mu_{k'} \neq \mu_{k+1}$. If the information structure $\{0, \mu_{k'}\}$ that induces a_{-j} and $\alpha_{k'}$ is incentive compatible, then the same argument provided in case (1) shows that this information structure will give a higher payoff to the sender than does P' . So we only need to consider the case that $\{0, \mu_{k'}\}$ is not incentive compatible. In this case, because a_{-j} is the sender's most-preferred action in state 0, incentive compatibility can fail only when $u_S(a_{-j}, \mu_{k'}) > \mathbb{E}_{\alpha_{k'}} [u_S(a, \mu_{k'})]$ (i.e., the sender prefers a_{-j} to $\alpha_{k'}$ at belief $\mu_{k'}$). Moreover, since the sender prefers $\alpha_{k'}$ to α_{-j} at belief $\mu_{k'}$ (incentive compatibility), by transitivity he prefers a_{-j} to α_{-j} at belief $\mu_{k'}$. He also prefers a_{-j} to α_{-j} at belief 0. Because preferences are linear in beliefs, this implies that he prefers a_{-j} to α_{-j} at belief μ_{-j} . Therefore,

$$\begin{aligned}
\mathbb{E}_{P'} \left[\max \left\{ \mathbb{E}_{\alpha_{-j}} [u_S(a, \mu)], \mathbb{E}_{\alpha_{k'}} [u_S(a, \mu)] \right\} \right] &< \mathbb{E}_{P'} [u_S(a_{-j}, \mu)] \\
&= u_S(a_{-j}, \mu_0) \\
&\leq \mathbb{E}_P \left[\max \left\{ u_S(a_{-j}, \mu), u_S(\bar{a}_{k+1}, \mu) \right\} \right].
\end{aligned}$$

The first inequality follows from the fact that a_{-j} is strictly better than α_{-j} and $\alpha_{k'}$ at belief μ_{-j} and belief $\mu_{k'}$, respectively. The last inequality follows from the fact that the information structure P is incentive compatible for a_{-j} and \bar{a}_{k+1} . Therefore there is a

positive information value as the sender can correctly recommend \bar{a}_{k+1} instead of a_{-j} at belief μ_{k+1} . \square

Proposition 4 implies that as long as $u_S(a_{-j}, \mu) \leq \bar{v}(\mu)$ for some $\mu > \mu_0$, the support of the optimal random posterior contains 0. If it also contains 1, then the optimal experiment reveals perfect information. If it does not contain 1, the optimal experiment will generate a conclusive signal of the common-interest state. In other words, the underlying Blackwell experiment corresponding to this optimal random posterior will produce a signal that reveals the common-interest state 0 with probability strictly less than 1 when the true state is 0, and never produces a signal that would suggest the state is 0 when the true state is 1. This means that the ex-ante probability that the receiver takes action a_{-j} under the optimal information structure cannot exceed $1 - \mu_0$.

5. Informativeness Compared to Bayesian Persuasion

In general, the optimal experiment when the sender has no commitment power can be more or less informative than (or not Blackwell-comparable to) the optimal experiment chosen when the sender can commit to truthfully revealing the outcome of the experiment. For example, when sender and receiver have opposite marginal incentives, Proposition 2 shows that the optimal experiment in our setup is a totally uninformative experiment, while the optimal experiment with full commitment is typically non-degenerate, as the concave envelope of $\bar{v}(\cdot)$ does not coincide $\bar{v}(\cdot)$ itself.

For an example in which the optimal experiment in our setup is more informative than that in a model with full commitment, consider the case where the sender strictly prefers some action a_n to all other actions in both states. Let a_n be the receiver's best response when the belief is in the interval $[\underline{I}_n, \bar{I}_n]$, and let the prior belief μ_0 be lower than \underline{I}_n . The lesson we learn from Kamenica and Gentzkow (2011) is that if the optimal experiment with full commitment induces a_n and some other action, it maximizes the ex-ante probability that a_n will be taken by inducing the smallest posterior belief \underline{I}_n that is just enough to induce the receiver to choose a_n . When the sender lacks commitment power in communication, inducing the pure action a_n is not incentive compatible. However, it may be incentive compatible to induce a mixture action between a_n and a_{n+1} at belief \bar{I}_n . Because \bar{I}_n is farther from μ_0 than \underline{I}_n is from μ_0 , the resulting experiment is more informative than the optimal experiment under full commitment.¹⁸ The next

¹⁸ Lipnowski et al. (2022) show that lowering the sender's credibility can result in a more informative




Figure 7: Optimal experiment with and without commitment.

proposition specifies the precise conditions for an analogous argument to be valid.

Proposition 5. *Assume that sender and receiver have aligned marginal incentives and $|A| \geq 3$. If an action a_n ($n \neq -J, K$) is strictly preferred to all other actions in both states, then the optimal information structure in our model is strictly more informative than the optimal experiment under full commitment for some prior belief.*

The proof of this proposition is in the Appendix. Consider Figure 7. The sender strictly prefers a_n to all other actions in both states. The optimal experiment under full commitment at prior belief μ_0 has support $\{l_{n'}, l_n\}$. The optimal experiment under full commitment at belief μ'_0 has support $\{l_n, l_{n''}\}$.

In the left panel (case a), the sender with belief $l_{n'}$ strictly prefers a_n over $a_{n'}$ over $a_{n''}$. It implies that there exists a randomization α_n^{PM} between a_n and $a_{n''}$ such that the experiment with support $\{l_{n'}, l_n\}$ is IC-PM. Recall that with aligned marginal incentives, $m_S(\alpha_n^{PM}) > m_S(a_n)$. Therefore the expected payoff from such IC-PM experiment (the orange triangle) is strictly higher than $u_S(a_{n'}, l_{n'}) + m_S(a_n)(\mu_0 - l_{n'})$ (lying on the gray dashed line). Moreover, with aligned marginal incentives, any incentive compatible experiment that induces $a_{n'}$ and some action smaller than a_n can only lead to an expected payoff strictly below $u_S(a_{n'}, l_{n'}) + m_S(a_n)(\mu_0 - l_{n'})$. For example, in Figure 7(a), the experiment with support $\{l_{n'}, l_n\}$ is IC-PP for $a_{n'}$ and $a_{n'+1}$. However, the sender's expected payoff from it is bounded by the gray dashed line because the slope of the sender's expected payoff is smaller than the marginal incentives of $a_{n'+1}$ (implied by Lemma 2)

equilibrium, and label this phenomenon as “productive mistrust.”

which is smaller than $m_s(a_n)$. Thus, in this case, under the prior belief μ_0 , the optimal experiment in our model is more informative than that under full commitment.

It is possible that the sender with belief $\underline{I}_{n'}$ strictly prefers all actions higher than a_n over $a_{n'}$, so that we cannot find an incentive compatible experiment that is more informative than $\{\underline{I}_{n'}, \underline{I}_n\}$. This happens in the right panel (case b). However, given the assumption of aligned marginal incentives, there must exist in this case two actions (weakly) smaller than a_n such that the sender with belief $\bar{I}_{n''}$ prefers one over $a_{n''}$ over the other one. In Figure 7(b), type- $\bar{I}_{n''}$ sender prefers a_n over $a_{n''}$ over $a_{n'}$. With similar reasoning as in case (a), under the prior belief μ'_0 , the optimal experiment in our model is more informative than that under full commitment.

6. Canonical Cheap Talk

In this section, we discuss the connection between our model and the canonical cheap talk model under binary state space and finite action space. We introduce a modification of Algorithm 1 to find the highest equilibrium payoff for the sender in a canonical cheap talk game.

In the canonical cheap talk game, the sender is initially perfectly informed about the true state. His reporting strategy is essential for generating credible information transmission. It is obvious that for every equilibrium in the canonical cheap talk game, there is a corresponding game in our model inducing the same equilibrium outcome—namely, the sender commits to the information structure that is his reporting strategy in the canonical cheap talk game, and then truthfully reports his private information outcomes.

Conversely, pick a truth-telling equilibrium in our game. To ensure that its outcome is also feasible in the canonical cheap talk model, we need to verify an additional constraint: while fixing the receiver's decision rule to be the same as in the truth-telling equilibrium, the sender cannot gain from deviating to a more informative information structure. If this constraint fails, using the information structure in our game as the reporting strategy is not incentive compatible for a sender who knows the true state.¹⁹

To illustrate the logic behind this constraint, Figure 8(a) provides an example of one-sided randomization (PM). The random posterior $\{\mu_{-j}, \mu_k\}$ is IC-PM that the sender is

¹⁹ Specifically, this constraint requires that, given a fixed truth-telling equilibrium in our model, each type $\theta \in \{0, 1\}$ is indifferent among any actions that is induced under any posterior beliefs $\mu \neq 1 - \theta$.

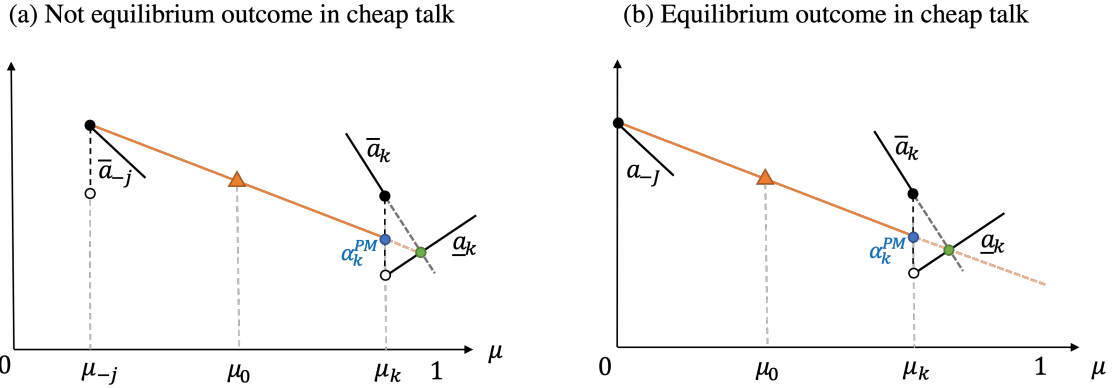


Figure 8: Comparing with canonical cheap talk

indifferent between mixed action α_k^{PM} and pure action \bar{a}_{-j} at belief μ_{-j} , and strictly prefers α_k^{PM} to \bar{a}_{-j} at belief μ_k . Because both μ_{-j} and μ_k are interior and preferences are linear in beliefs, this in turn implies that the sender strictly prefers α_k^{PM} to \bar{a}_{-j} at belief 1 and strictly prefers \bar{a}_{-j} to α_k^{PM} at belief 0. The equilibrium outcome induced by this one-sided randomization cannot be sustained as an equilibrium outcome in the canonical cheap talk game. To produce an outcome that induces interior beliefs μ_{-j} and μ_k , the sender must adopt a reporting strategy that recommends both actions with positive probabilities in each state, but this is not incentive compatible for a sender who knows the true state. To put it slightly differently, the information structure with support $\{\mu_{-j}, \mu_k\}$ cannot be sustained as an equilibrium outcome if the sender cannot commit to this experiment, because he can gain from deviating to learn more about the state.

Figure 8(b) modifies the example to show when one-sided randomization can be supported as an equilibrium outcome in the canonical cheap talk game. The only difference is that, in Figure 8(b), the pure action $\bar{a}_{-j} = a_j$ is taken at degenerate belief 0. To produce the random posterior $\{0, \mu_k\}$ under cheap talk, the sender must adopt a reporting strategy which recommends the mixed action α_k^{PM} only in state 1, and recommends both actions with positive probability in state 0. Since the sender is indifferent between these two actions in state 0, such a reporting strategy is indeed incentive compatible and will produce an interior belief μ_k for the receiver upon getting the recommendation to choose α_k^{PM} .

This example demonstrates that an IC-PM or IC-MP outcome in our model cannot be supported as an equilibrium outcome of the cheap talk game if the pure action is

chosen at an interior belief and the sender has a strict preference at the other belief,²⁰ but it is an equilibrium outcome of the cheap talk game if the pure action is chosen at degenerate beliefs (0 or 1). In other words, when looking for pure action in an IC-PM or IC-MP equilibrium of the canonical cheap talk game, we only need to consider action a_{-j} or a_k respectively. Similarly, when looking for an IC-PP equilibrium in the canonical cheap talk game, we only need to consider the actions a_{-j} and a_k . On the other hand, every IC-MM outcome in our model can be supported as an equilibrium outcome in the canonical cheap talk game. Because by construction, the sender is indifferent between a pair of IC-MM actions at beliefs μ_{-j}, μ_k , and therefore he is also indifferent between them at beliefs 0 and 1.

This suggests that the sender-optimal equilibrium in the cheap talk model can be obtained from a straightforward modification of Algorithm 1.

Algorithm 2:

1. For the pair $(-j-1, K+1)$, compute $W_{-j-1, K+1}(\mu_0; a_{-j}, a_k)$. If this pair is IC-PP let $W^1 = W_{-j-1, K+1}(\mu_0; a_{-j}, a_k)$; otherwise let $W^1 = \bar{v}(\mu_0)$.
2. For every pair $(-j-1, k)$ where $k \in \{1, \dots, K+1\}$, compute $W_{-j-1, k}(\mu_0; a_{-j}, \alpha_k^{PM})$ and rank these values from highest to lowest. Starting with the pair with the highest value, verify whether it is IC-PM or not. Stop the first time when an IC-PM pair is found. Assign $W^{(a)} = W_{-j-1, k}^{PM}$ for such pair. If none of them is IC-PM, assign $W^{(a)} = \bar{v}(\mu_0)$. Symmetrically, for every pair $(-j, K+1)$ where $j \in \{-j-1, \dots, -1\}$, go through a similar procedure in the case for MP. Assign $W^{(b)} = W_{-j, K+1}^{MP}$ the first time an IC-MP is found. If none of them is IC-MP, assign $W^{(b)} = \bar{v}(\mu_0)$. Set $W^2 = \max\{W^{(a)}, W^{(b)}\}$.
3. For every pair $(-j, k) \in \{-j, \dots, -1\} \times \{1, \dots, K\}$, compute $W_{-j, k}(\mu_0; \alpha_{-j}^{MM}, \alpha_k^{MM})$ and rank these values from the highest to lowest. Starting with the pair with the highest value, verify whether it is IC-MM or not. Stop the first time an IC-MM pair is found and assign $W^3 = W_{-j, k}^{MM}$ for such pair. If none of them is IC-MM, assign $W^3 = \bar{v}(\mu_0)$.
4. Assign $W^c(\mu_0) = \max\{W^1, W^2, W^3\}$. The random posterior with support $\{\mu_{-j}, \mu_k\}$ corresponds to the pair that yields $W^c(\mu_0)$ is the sender's reporting strategy in the optimal equilibrium.

²⁰ If the sender is indifferent at the other belief that induces the mixed action, then such one-sided randomization can be an equilibrium outcome in a cheap talk game. However, it is covered by IC-MM.

Corollary 1. *Algorithm 2 determines the highest equilibrium payoff for the sender in a canonical cheap talk game (with binary states and finite actions).*

It is immediate that sender’s commitment on the information structure is valuable if and only if $W^*(\mu_0) > W^c(\mu_0)$, i.e., when the two algorithms produce different outcomes.

7. Discussion

The model in this paper has a close relation with a particular scheme of mediated communication, in which a mediator maximizes the ex-ante welfare of an informed sender (Salamanca, 2021). Specifically, a perfectly informed sender sends a message about his private information to a mediator. The mediator then communicates a message to the receiver according to a noisy reporting rule that the mediator commits to at the beginning of the game. After receiving the message from the mediator, the receiver takes an action. If we consider the mediator’s reporting rule as a mapping from the sender’s private information to a distribution of action recommendations, this rule can be interpreted as an information structure. The incentive constraints for this mediated communication game are imposed at the *ex ante* stage, which require every type of sender who perfectly knows the state to report his private information truthfully before observing the message sent by the mediator.

In contrast, the sender in our model is uninformed when he commits to an information structure, and then reports his private information to the receiver after observing the outcome of the experiment. Therefore, our model requires *interim-stage* incentive constraints, such that the sender with an interim belief derived from the observed outcome prefers to report his private information truthfully. In spite of this difference, if our sender and the mediator in Salamanca (2021) commit to the same information structure in equilibrium, then both models will yield the same equilibrium outcomes.

Interestingly, under binary state space, the highest equilibrium payoff $W^*(\mu_0)$ that the sender can achieve in our model is always weakly lower than the maximum ex-ante welfare of the sender (i.e., evaluated at μ_0 before the sender becomes perfectly informed) in Salamanca (2021) for any μ_0 .

To see this, suppose the optimal random posterior in our sender-receiver game has support $\{\mu', \mu''\}$ and the receiver optimally chooses $a' \in A_R(\mu')$ and $a'' \in A_R(\mu'')$ at the respective beliefs (the same argument will go through if the receiver takes mixed strat-

egy). Without loss of generality, let $\mu' < \mu''$. Then incentive compatibility constraints (1) in our model implies that the following also holds:

$$u_S(a', 0) \geq u_S(a'', 0), \quad u_S(a'', 1) \geq u_S(a', 1).$$

This means that it is incentive compatible for an informed sender to truthfully report his private information (belief 0 or 1) to the mediator, whenever the mediator commits to a reporting rule that recommends a' more often if the sender reports 0 and recommends a'' more often if the sender reports 1. Therefore, the sender's incentive constraints in the mediated communication game are satisfied if the mediator commits to the same information structure as the underlying experiment that induces our optimal random posterior. In other words, interim incentive compatibility in our model is more stringent than the incentive compatibility restrictions required by the mediator model, and therefore our model delivers a (weakly) lower expected payoff for the sender than that achievable in Salamanca (2021).

Appendix

Proof of Lemma 1. Given P , σ_S , and a message $m \in M$, the receiver forms a posterior belief $\hat{\mu}^m$, where

$$\hat{\mu}^m(\theta) = \sum_{\mu \in \text{supp}(P)} \frac{P(\mu)\sigma_S(m|\mu)}{\sum_{\mu \in \text{supp}(P)} P(\mu)\sigma_S(m|\mu)} \mu(\theta),$$

for $\theta \in \Theta$. We use $\hat{P} \in \Delta(\Delta\Theta)$ to denote the distribution of the receiver's posterior beliefs, with $\hat{P}(\hat{\mu}^m) = \sum_{\mu \in \text{supp}(P)} P(\mu)\sigma_S(m|\mu)$. Then player i 's expected utility can be simplified to:

$$U_i(\sigma_S, \sigma_R, P) = \sum_{m \in M, \theta \in \Theta, a \in A} \hat{P}(\hat{\mu}^m) \hat{\mu}^m(\theta) \sigma_R(a|m) u_i(a, \theta).$$

Thus each player's expected utility only depends on the joint distribution of the receiver's posterior belief and the action. If we let the sender directly commit to the random posterior \hat{P} , and construct a (truth-telling) reporting strategy $\hat{\sigma}_S$ such that for all $\mu \in \text{supp}(\hat{P})$, $\hat{\sigma}_S(m|\hat{\mu}^m) = 1$, then player i 's expected utility further simplifies to:

$$U_i(\sigma_S, \sigma_R, P) = U_i(\hat{\sigma}_S, \sigma_R, \hat{P}).$$

Moreover, (σ_S, σ_R, P) being an equilibrium strategy profile implies $(\hat{\sigma}_S, \sigma_R, \hat{P})$ is an equilibrium strategy profile. Since reporting $m \in M$ is a best response to σ_R for every sender type of $\{\mu \in \text{supp}(P) : \sigma_S(m|\mu) > 0\}$, reporting m is also a best response for sender type $\hat{\mu}^m$, as $\hat{\mu}^m$ is a convex combination of $\{\mu \in \text{supp}(P) : \sigma_S(m|\mu) > 0\}$.

To prove the second part, suppose a random posterior P with $|\text{supp}(P)| > |\Theta|$ that can lead to a truth-telling equilibrium is optimal. By Carathéodory's Theorem and Krein-Milman Theorem, $\mu_0 = \mathbb{E}_P[\mu]$ can be written as a convex combination of $|\Theta|$ elements of $\text{supp}(P)$, denoted as $P' \in \Delta(\Delta\Theta)$ with $|\text{supp}(P')| = |\Theta|$ and $\text{supp}(P') \subset \text{supp}(P)$. Let the receiver preserve σ_R , then P' can lead to a truth-telling equilibrium. Let $c : \text{co}(\text{supp}(P)) \rightarrow \mathbb{R}$ be the smallest concave function such that $c(\mu) \geq \sum_{\theta, a} \mu(\theta) \sigma_R(a|\mu) u_S(a, \theta)$ at all $\mu \in \text{supp}(P)$. The random posterior P' can perform equally well as P because c must be affine on $\text{co}(\text{supp}(P))$. \square

Proof of Theorem 1.

Claim 1. *If the optimal random posterior has support $\{\mu_{-j}, \mu_k\}$, and the receiver uses mixed strategies $\alpha_{-j} \in \Delta A_R(\mu_{-j})$ and $\alpha_k \in \Delta A_R(\mu_k)$ (with full support) in the sender-preferred equilibrium, then at least one of the following is true:*

- (i) $\alpha_{-j} = \alpha_{-j}^{MM}$ and $\alpha_k = \alpha_k^{MM}$;
- (ii) $W^*(\mu_0) \in \{W_{-j,k}^{PM}, W_{-j,k}^{MP}, W_{-j,k}^{PP}\}$.

Proof of Claim 1. Since the receiver uses mixed strategies at each belief, $\mu_{-j} \neq 0$ and $\mu_k \neq 1$. Let $A_R(\mu_{-j}) = \{a_{-j}, a_{-j+1}\}$ and $A_R(\mu_k) = \{a_{k-1}, a_k\}$. Since α_{-j} and α_k are incentive compatible for the sender, from (1),

$$\mathbb{E}_{\alpha_{-j}}[u_S(a, \mu_{-j})] \geq \mathbb{E}_{\alpha_k}[u_S(a, \mu_{-j})], \quad \mathbb{E}_{\alpha_k}[u_S(a, \mu_k)] \geq \mathbb{E}_{\alpha_{-j}}[u_S(a, \mu_k)]. \quad (5)$$

Suppose $u_S(a_{-j}, \mu_{-j}) \neq u_S(a_{-j+1}, \mu_{-j})$ and $u_S(a_{k-1}, \mu_k) \neq u_S(a_k, \mu_k)$. Suppose further that both inequalities hold strictly and α_{-j} and α_k have full support. Then the sender can achieve a strictly higher expected payoff if the receiver deviates from α_{-j} by assigning a slightly larger probability on the sender-preferred action $\bar{a}_{-j} \in A_R(\mu_{-j})$. As long as the increase in probability of choosing \bar{a}_{-j} is small enough, such deviation would raise the payoff from truth-telling at belief μ_{-j} without violating the truth-telling constraint at belief μ_k . Suppose the first inequality in (5) holds as an equality and the second inequality in (5) holds strictly. Then the sender can achieve a strictly higher expected payoff under the same argument. If the first inequality in (5) holds strictly and the second inequality in (5) holds as an equality, a symmetric argument will apply. Therefore, the optimality of α_{-j} and α_k implies that both inequalities in (5) hold as an equality.

If $m_S(a_{-j})$, $m_S(a_{-j+1})$, $m_S(a_{k-1})$ and $m_S(a_k)$ are not all equal, then there is a unique $(\alpha_{-j}^{MM}, \alpha_k^{MM}) \in \Delta A_R(\mu_{-j}) \times \Delta A_R(\mu_k)$ such that the constraints (5) hold as equalities. Next, if $m_S(a_{-j}) = m_S(a_{-j+1}) = m_S(a_{k-1}) = m_S(a_k)$, then there exist infinitely many solutions. The optimality of α_{-j} and α_k will then imply that the receiver takes either \bar{a}_{-j} at belief μ_{-j} and α_k^{MM} at belief μ_k , or \bar{a}_k at μ_k and α_{-j}^{MM} at belief μ_{-j} . This contradicts the premise that both α_{-j} and α_k have full support. Moreover, in this case $W^*(\mu_0) \in \{W_{-j,k}^{PM}, W_{-j,k}^{MP}\}$.

If at least one of $u_S(a_{-j}, \mu_{-j}) \neq u_S(a_{-j+1}, \mu_{-j})$ and $u_S(a_{k-1}, \mu_k) \neq u_S(a_k, \mu_k)$ does not hold, we can slightly alter the above argument to show $W^*(\mu_0) \in \{W_{-j,k}^{PM}, W_{-j,k}^{MP}, W_{-j,k}^{PP}\}$. \square

Claim 2. *If the optimal random posterior has support $\{\mu_{-j}, \mu_k\}$, and the receivers takes*

pure action $a' \in A_R(\mu_{-j})$ and mixed action $\alpha_k \in \Delta A_R(\mu_k)$ (with full support) in the sender-preferred equilibrium, then at least one of the following is true:

- (i) $a' = \bar{a}_{-j}$ and $\alpha_k = \alpha_k^{PM}$;
- (ii) $W^*(\mu_0) \in \{W_{-j,k}^{MM}, W_{-j,k}^{MP}, W_{-j,k}^{PP}\}$.

Proof of Claim 2. Since the receiver uses mixed strategies at belief μ_k , $\mu_k \neq 1$ and $A_R(\mu_k) = \{a_{k-1}, a_k\}$. Since a' and α_k are incentive compatible for the sender, from (1),

$$u_S(a', \mu_{-j}) \geq \mathbb{E}_{\alpha_k}[u_S(a, \mu_{-j})], \quad \mathbb{E}_{\alpha_k}[u_S(a, \mu_k)] \geq u_S(a', \mu_k). \quad (6)$$

Suppose that $u_S(a_{k-1}, \mu_k) \neq u_S(a_k, \mu_k)$. We first show that the first inequality in (6) must hold as an equation. Suppose to the contrary that this inequality holds strictly. Then the sender can achieve a strictly higher payoff if the receiver deviates from α_k by assigning a slightly larger probability on the sender-preferred action $\bar{a}_k \in A_R(\mu_k)$. As long as the increase in probability of choosing \bar{a}_k is small enough, such deviation would raise the sender's payoff from truth-telling at belief μ_k without violating (6) at belief μ_{-j} , leading to a contradiction.

Now, suppose the second inequality in (6) also holds as an equation. Given the result established above, the sender is indifferent between a' and α_k both at belief μ_{-j} and at belief μ_k . This case then reduces to the double-randomization case. Therefore, $W^*(\mu_0) = W_{-j,k}^{MM}$, and part (ii) of this claim is satisfied.

Next, suppose the second inequality in (6) holds strictly. There are two possibilities: (1) the sender obtains different payoffs from a_{-j} and a_{-j+1} at belief μ_{-j} , (2) the sender obtains the same payoff from a_{-j} and a_{-j+1} at belief μ_{-j} .

Case (1). Suppose a' is not the sender-preferred action \bar{a}_{-j} at belief μ_{-j} . Then the sender can achieve a strictly higher payoff by inducing the receiver to deviate from a' to a mixed strategy α_{-j} that assigns a small positive probability on \bar{a}_{-j} , without violating the incentive constraints. This shows that a' must be equal to \bar{a}_{-j} . Because $a' = \bar{a}_{-j}$, and the first condition of (6) as an equation implies that $\alpha_k = \alpha_k^{PM}$, and part (i) of this claim is satisfied.

Case (2). When the sender obtains the same payoff from a_{-j} and a_{-j+1} at belief μ_{-j} , the convention we adopt is $\bar{a}_{-j} = a_{-j+1}$. Suppose $a' = a_{-j} \neq \bar{a}_{-j}$. Then the optimality of $\{\mu_{-j}, \mu_k\}$ implies that the sender is indifferent between α_k^{PM} and a_{-j} at belief μ_k . Otherwise, the random posterior with support $\{\mu_{-j-1}, \mu_k\}$ can perform strictly better.

Therefore, the second inequality in (6) cannot hold strictly under the convention we adopt. This contradiction implies that we must have $a' = a_{-j+1} = \bar{a}_{-j}$. The first condition of (6) as an equation then implies that $\alpha_k = \alpha_k^{PM}$, and part (i) of this claim is satisfied.

Suppose $u_S(a_{k-1}, \mu_k) = u_S(a_k, \mu_k)$. If the first inequality in (6) holds with equality, then we can use the same argument as above. Otherwise, if the first inequality in (6) holds strictly, then we can slightly alter the above argument to show that $W^*(\mu_0) \in \{W_{-j,k}^{MM}, W_{-j,k}^{MP}, W_{-j,k}^{PP}\}$. \square

Claim 3. *If the optimal random posterior has a support $\{\mu_{-j}, \mu_k\}$, and the receiver uses only pure strategies $a' \in A_R(\mu_{-j})$ and $a'' \in A_R(\mu_k)$ in the sender-preferred equilibrium, then either one of the following is true:*

- (i) $a' = \bar{a}_{-j}$ and $a'' = \bar{a}_k$;
- (ii) $W^*(\mu_0) \in \{W_{-j,k}^{PM}, W_{-j,k}^{MP}, W_{-j,k}^{MM}\}$.

Proof of Claim 3. If $\mu_{-j} = 0$ and $\mu_k = 1$, then $A_R(\mu_{-j}) = \bar{a}_{-j}$, $A_R(\mu_k) = \bar{a}_k$, and part (i) is satisfied.

If $\mu_{-j} \neq 0$ and $\mu_k \neq 1$, then $A_R(\mu_{-j}) = \{a_{-j}, a_{-j+1}\}$, $A_R(\mu_k) = \{a_{k-1}, a_k\}$, and there are four possibilities. (1) The sender obtains different payoffs from a_{-j} and a_{-j+1} at belief μ_{-j} , and different payoffs from a_{k-1} and a_k at belief μ_k . (2) The sender obtains same payoff from a_{-j} and a_{-j+1} at belief μ_{-j} , but different payoffs from a_{k-1} and a_k at belief μ_k . (3) The sender obtains different payoffs from a_{-j} and a_{-j+1} at belief μ_{-j} , and same payoff from a_{k-1} and a_k at belief μ_k . (4) The sender obtains same payoff from a_{-j} and a_{-j+1} at belief μ_{-j} , and same payoff from a_{k-1} and a_k at belief μ_k .

Case (1). By incentive compatibility,

$$u_S(a', \mu_{-j}) \geq u_S(a'', \mu_{-j}), \quad u_S(a'', \mu_k) \geq u_S(a', \mu_k). \quad (7)$$

Suppose both inequalities hold strictly. The optimality of (a', a'') implies $a' = \bar{a}_{-j}$ and $a'' = \bar{a}_k$, with a reasoning similar to that in the proof of Claim 1. Thus, part (i) is satisfied. Suppose the first inequality holds as equality and the second inequality holds strictly. Then the optimality of (a', a'') implies $a' = \bar{a}_{-j}$, with a reasoning similar to that in the proof of Claim 2. Moreover, when $a' = \bar{a}_{-j}$ and the first inequality holds as equality, we have $a'' = \alpha_k^{PM}$. Therefore, $W^*(\mu_0) = W_{-j,k}^{PM}$, and part (ii) is satisfied. Similarly, if

the first inequality hold strictly and the second inequality holds as an equality, then $W^*(\mu_0) = W_{-j,k}^{MP}$ and part (ii) is also satisfied. Finally, suppose that both inequalities hold as an equality. Then if $m_S(a_{-j})$, $m_S(a_{-j+1})$, $m_S(a_{k-1})$ and $m_S(a_k)$ are not all equal, $a' = \alpha_{-j}^{MM}$ and $a'' = \alpha_k^{MM}$. Therefore, $W^*(\mu_0) = W_{-j,k}^{MM}$. If $m_S(a_{-j})$, $m_S(a_{-j+1})$, $m_S(a_{k-1})$ and $m_S(a_k)$ are equal, then $W^*(\mu_0) \in \{W_{-j,k}^{PM}, W_{-j,k}^{MP}\}$. Part (ii) is again satisfied.

Case (2a). If $a' = a_{-j} \neq \bar{a}_{-j}$, then the optimality of $\{\mu_{-j}, \mu_k\}$ implies that the sender is indifferent between a'' and a_{-j} at belief μ_k ; otherwise the random posterior with support $\{\mu_{-j-1}, \mu_k\}$ is strictly better. Moreover, if the sender is indifferent between a_{-j} and a'' at belief μ_{-j} , then $W^*(\mu_0) = W_{-j,k}^{MM}$. On the other hand, if the sender strictly prefers a_{-j} over a'' at belief μ_{-j} , then the optimality of a'' implies $a'' = \bar{a}_k$. That is, $W^*(\mu_0) = W_{-j,k}^{MP}$. In both sub-cases, part (ii) is satisfied. Case (2b). If $a' \neq a_{-j}$, then $a' = \bar{a}_{-j}$. Then if the first inequality in (7) hold strictly, the optimality of (\bar{a}_{-j}, a'') implies that $a'' = \bar{a}_k$, and part (i) is satisfied. On the other hand, if the first inequality in (7) hold as equality, then $W^*(\mu_0) = W_{-j,k}^{PM}$, and part (ii) is satisfied.

Case (3) is symmetric to case (2), and if we apply all arguments above, case (4) implies $W^*(\mu_0) \in \{W_{-j,k}^{PP}, W_{-j,k}^{PM}, W_{-j,k}^{MP}, W_{-j,k}^{MM}\}$.

Finally, if either $\mu_{-j} = 0$ or $\mu_k = 1$, then with a similar reasoning we can conclude $W^*(\mu_0) \in \{W_{-j,k}^{PP}, W_{-j,k}^{PM}, W_{-j,k}^{MP}, W_{-j,k}^{MM}\}$. \square

Claims 1–3 (together with an analogous claim for the case of MP) imply that it is sufficient to only focus on $W_{-j,k}^{PP}$, $W_{-j,k}^{PM}$, $W_{-j,k}^{MP}$, and $W_{-j,k}^{MM}$.

Lastly, in step 2 of the algorithm, we only consider $(-j, k) \in S_1$. To see the sufficiency of it, suppose a pair of $(-j, k)$ outside S_1 is IC-PP, or IC-PM, or IC-MP, or IC-MM. Then,

$$\max \{W_{-j,k}^{PP}, W_{-j,k}^{PM}, W_{-j,k}^{MP}, W_{-j,k}^{MM}\} \leq W_{-j,k}(\mu_0; \bar{a}_{-j}, \bar{a}_k) \leq W^1.$$

The first inequality follows from the fact that the sender's value correspondence $v(\mu_{-j})$ is lower than $u_S(\bar{a}_{-j}, \mu_{-j})$, and $v(\mu_k)$ is lower than $u_S(\bar{a}_k, \mu_k)$. The second inequality comes from the construction of S_1 . So even if the random posterior with support $\{\mu_{-j}, \mu_k\}$ outside S_1 is incentive compatible, it will never improve on the outcome from step 1. Moreover, in step 3 of the algorithm, we only consider $(-j, k) \in S_2$. To see the sufficiency of it, suppose a pair of $(-j, k)$ outside S_2 is IC-MM. Then,

$$W_{-j,k}^{MM} \leq \max \{W_{-j,k}(\mu_0; \bar{a}_{-j}, \alpha_k^{PM}), W_{-j,k}(\mu_0; \alpha_{-j}^{PM}, \bar{a}_k)\} \leq \max \{W^{(a)}, W^{(b)}\} = W^2.$$

The second inequality is from our construction of S_2 . To see the first inequality, notice that under IC-MM, the sender is indifferent between α_{-j}^{MM} and $\alpha_k^{MM} := (\gamma_k, 1 - \gamma_k)$ at belief μ_{-j} ,

$$\mathbb{E}_{\alpha_{-j}}[u_S(a, \mu_{-j})] = \gamma_k u_S(\bar{a}_k, \mu_{-j}) + (1 - \gamma_k) u_S(\underline{a}_k, \mu_{-j}).$$

From our construction of $\alpha_k^{PM} := (\gamma'_k, 1 - \gamma'_k)$,

$$u_S(\bar{a}_{-j}, \mu_{-j}) = \gamma'_k u_S(\bar{a}_k, \mu_{-j}) + (1 - \gamma'_k) u_S(\underline{a}_k, \mu_{-j}).$$

Therefore $\mathbb{E}_{\alpha_{-j}}[u_S(a, \mu_{-j})] \leq u_S(\bar{a}_{-j}, \mu_{-j})$ implies $\gamma_k \leq \gamma'_k$, which further implies that $\mathbb{E}_{\alpha_k^{MM}}[u_S(a, \mu_k)] \leq \mathbb{E}_{\alpha_k^{PM}}[u_S(a, \mu_k)]$. We thus have $W_{-j,k}^{MM} \leq W_{-j,k}(\mu_0; \bar{a}_{-j}, \alpha_k^{PM})$. For a similar reason, $W_{-j,k}^{MM} \leq W_{-j,k}(\mu_0; \alpha_{-j}^{PM}, \bar{a}_k)$. Notice that this argument does not require $\{\mu_{-j}, \mu_k\}$ to be IC-PM or IC-MP. \square

Proof of Proposition 5. With aligned marginal incentives, the concavification result in Kamenica and Gentzkow (2011) implies that there exists an $n' < n$ such that with prior belief $\mu_0 \in (I_{n'}, \bar{I}_{n'})$, the optimal experiment under full commitment has support $\{I_{n'}, I_n\}$. Similarly, there exists an $n'' > n$ such that with a different prior belief $\mu'_0 \in (I_{n''}, \bar{I}_{n''})$, the optimal experiment under full commitment has support $\{\bar{I}_n, \bar{I}_{n''}\}$. We want to show that, in our model, the optimal experiment is strictly more informative than that under full commitment either when the prior is μ_0 or when the prior is μ'_0 . There are only two cases. (a) There exists a $k \geq n$ such that $u_S(a_k, I_{n'}) \geq u_S(a_{n'}, I_{n'}) \geq u_S(a_{k+1}, I_{n'})$. (b) For every $k \geq n$, $u_S(a_k, I_{n'}) > u_S(a_{n'}, I_{n'})$. Notice that if there exists an experiment with support $\{I_{n'}, \mu\}$ where $\mu \in (\bar{I}_{n'}, I_n]$ that is incentive compatible, then the sender's expected payoff from such experiment is smaller than $u_S(a_{n'}, I_{n'}) + m_S(a_n)(\mu_0 - I_{n'})$ (this is implied by Lemma 2). Similarly, if there exists an experiment with support $\{\mu, \bar{I}_{n''}\}$ where $\mu \in (\bar{I}_n, \bar{I}_{n''})$ that is incentive compatible, then the sender's expected payoff under μ'_0 from such experiment is smaller than $u_S(a_{n''}, \bar{I}_{n''}) - m_S(a_n)(\bar{I}_{n''} - \mu'_0)$.

In case (a), the experiment with support $\{I_{n'}, \bar{I}_k\}$ is IC-PM given aligned marginal incentives. Moreover, sender's expected payoff from such experiment is greater than $u_S(a_{n'}, I_{n'}) + m_S(a_n)(\mu_0 - I_{n'})$. Because the receiver randomizes between a_k and a_{k+1} at belief \bar{I}_k and thereby the marginal incentive from such randomization α_k is greater than $m_S(a_n)$. Also, from the construction of IC-PM, the sender's expected payoff equals $u_S(a_{n'}, I_{n'}) + m_S(\alpha_k)(\mu_0 - I_{n'})$. Therefore, under μ_0 , there exists an experiment with support $\{I_{n'}, \bar{I}_k\}$ that is better than any IC experiment with support $\{I_{n'}, \mu\}$ where $\mu \in$

$(\bar{I}_{n'}, \underline{I}_n]$.

Moreover, there cannot exist an IC experiment $\{\mu, \mu'\}$ with $\mu < \underline{I}_{n'}$ and $\bar{I}_{n'} < \mu' < \underline{I}_n$ that yields the sender an expected payoff higher than $u_S(a_{n'}, \underline{I}_{n'}) + m_S(a_n)(\mu_0 - \underline{I}_{n'})$. To see this point, note that Lemma 2 implies that the slope of sender's expected payoff is smaller than $m_S(a')$ where $a' \in A_R(\mu')$, which in turn is smaller than $m_S(a_n)$ from aligned marginal incentives. This implies the sender's expected payoff from such an experiment, if the prior belief is $\underline{I}_{n'}$, would be higher than $u_S(a_{n'}, \underline{I}_{n'})$, which contradicts the fact that $u_S(a_{n'}, \underline{I}_{n'})$ lies on the concave envelope of $\bar{v}(\cdot)$.

Since $\bar{I}_k > \underline{I}_n$, the optimal experiment in our model under μ_0 , which has support $\{\underline{I}_{n'}, \bar{I}_k\}$, is strictly more informative than the (full commitment) experiment with support $\{\underline{I}_{n'}, \underline{I}_n\}$.

In case (b), since $u_S(a_{n''}, \underline{I}_{n'}) > u_S(a_{n'}, \underline{I}_{n'})$, we have $u_S(a_{n''}, \bar{I}_{n''}) > u_S(a_{n'}, \bar{I}_{n''})$ under the assumption of aligned marginal incentive. Then there must exist an a_k with $n' < k \leq n$ such that $u_S(a_k, \bar{I}_{n''}) \geq u_S(a_{n''}, \bar{I}_{n''}) \geq u_S(a_{k-1}, \bar{I}_{n''})$. Therefore, there exists an IC-MP experiment with support $\{\underline{I}_k, \bar{I}_{n''}\}$ such that the receiver randomizes between a_k and a_{k-1} at belief \underline{I}_k and such randomization α_k has a marginal incentive $m_S(\alpha_k)$ smaller than $m_S(a_n)$. Under μ'_0 , this experiment generates the sender an expected payoff higher than $u_S(a_{n''}, \bar{I}_{n''}) - m_S(a_n)(\bar{I}_{n''} - \mu'_0)$. Moreover, such experiment is better than any IC experiment with support $\{\mu, \bar{I}_{n''}\}$ where $\mu \in [\bar{I}_n, \bar{I}_{n''})$. Because $\underline{I}_k < \underline{I}_n$, the optimal experiment in our model under μ'_0 , which has support $\{\underline{I}_k, \bar{I}_{n''}\}$, is strictly more informative than the (full commitment) experiment with support $\{\bar{I}_n, \bar{I}_{n''}\}$. \square

References

- Alonso, Ricardo and Odilon Câmara**, “Organizing Data Analytics,” *Management Science*, May 2024, 70 (5), 3123–3143.
- Argenziano, Rossella, Sergei Severinov, and Francesco Squintani**, “Strategic Information Acquisition and Transmission,” *American Economic Journal: Microeconomics*, August 2016, 8 (3), 119–155.
- Arieli, Itai, Yakov Babichenko, Rann Smorodinsky, and Takuro Yamashita**, “Optimal Persuasion via Bi-pooling,” *Theoretical Economics*, 2023, 18 (1), 15–36.
- Babichenko, Yakov, Inbal Talgam-Cohen, Haifeng Xu, and Konstantin Zabarnyi**, “Algorithmic Cheap Talk,” 2023.
- Barros, Lucas**, “Information Acquisition Design,” *Working Paper*, 2022.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris**, “First-Price Auctions with General Information Structures: Implications for Bidding and Revenue,” *Econometrica*, January 2017, 85 (1), 107–143.
- Crawford, Vincent P and Joel Sobel**, “Strategic Information Transmission,” *Econometrica*, November 1982, 50 (6), 1431–1451.
- Deimen, Inga and Dezső Szalay**, “Delegated Expertise, Authority, and Communication,” *American Economic Review*, April 2019, 109 (4), 1349–1374.
- Denti, Tommaso, Massimo Marinacci, and Aldo Rustichini**, “Experimental Cost of Information,” *American Economic Review*, September 2022, 112 (9), 3106–3123.
- Dughmi, Shaddin and Haifeng Xu**, “Algorithmic Bayesian Persuasion,” *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, 2016, pp. 412–425.
- Felgenhauer, Mike and Elisabeth Schulte**, “Strategic Private Experimentation,” *American Economic Journal: Microeconomics*, November 2014, 6 (4), 74–105.
- Guo, Yingni and Eran Shmaya**, “Costly Miscalibration,” *Theoretical Economics*, May 2021, 16 (2), 477–506.

- Ivanov, Maxim**, “Informational Control and Organizational Design,” *Journal of Economic Theory*, March 2010, 145 (2), 721–751.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian Persuasion,” *American Economic Review*, October 2011, 101 (6), 2590–2615.
- Kleiner, Andreas, Benny Moldovanu, and Philipp Strack**, “Extreme Points and Majorization: Economic Applications,” *Econometrica*, 2021, 89 (4), 1557–1593.
- Krähmer, Daniel**, “Information Design and Strategic Communication,” *American Economic Review: Insights*, March 2021, 3 (1), 51–66.
- Kreutzkamp, Sophie**, “Endogenous Information Acquisition in Cheap-Talk Games,” *Working paper*, 2023.
- Lin, Xiao and Ce Liu**, “Credible Persuasion,” *Journal of Political Economy*, July 2024, 132 (7), 2228–2273.
- Lipnowski, Elliot**, “Equivalence of Cheap Talk and Bayesian Persuasion in a Finite Continuous Model,” *Working Paper*, 2020.
- and **Doron Ravid**, “Cheap Talk with Transparent Motives,” *Econometrica*, July 2020, 88 (4), 1631–1660.
- , —, and **Denis Shishkin**, “Persuasion via Weak Institutions,” *Journal of Political Economy*, October 2022, 130 (10), 2705–2730.
- Lou, Yichuan**, “Sender-Optimal Learning and Credible Communication,” *Working paper*, 2023.
- Nguyen, Anh and Teck Yong Tan**, “Bayesian Persuasion with Costly Messages,” *Journal of Economic Theory*, April 2021, 193, 105212.
- Pei, Harry Di**, “Communication with Endogenous Information Acquisition,” *Journal of Economic Theory*, December 2015, 160, 132–149.
- Roesler, Anne-Katrin and Balász Szentes**, “Buyer-Optimal Learning and Monopoly Pricing,” *American Economic Review*, July 2017, 107 (7), 2072–2080.
- Salamanca, Andrés**, “The Value of Mediated Communication,” *Journal of Economic Theory*, March 2021, 192, 105191.