Communication Protocols under Transparent Motives

Roberto Corrao* Yifan Dai†
MIT MIT

March 10, 2023

Abstract

We study optimal (information) mediation in sender-receiver communication games where the sender has transparent motives: she only cares about the receiver's actions and beliefs. We characterize the feasible distributions over the receiver's beliefs under mediation and the value of mediation. The sender achieves her optimal Bayesian persuasion value by mediation if and only if this value is attained by cheap talk. When the state is binary, mediation strictly improves on cheap talk if and only if the sender cannot do better than under cheap talk by always under- or over- stating the state.

*Department of Economics, MIT, rcorrao@mit.edu
†Department of Economics, MIT, yfdai@mit.edu

We are grateful to Ian Ball, Alessandro Bonatti, Simone Cerreia-Vioglio, Drew Fudenberg, Bob Gibbons, Haruki Kono, Giacomo Lanzani, Stephen Morris, Alexander Wolitzky, and Leeat Yariv for helpful comments and conversations, and the Gordon B. Pye Dissertation Fellowship for financial support.

1 Introduction

Consider a principal who faces a decision problem while being uncertain about some payoff-relevant state. The state is privately observed by an agent who can freely communicate with the principal to influence her decision and only cares about the ultimate decision the principal takes, that is, the agent has *transparent motives*. These situations are pervasive in economic settings such as (i) bilateral trade between a seller and a buyer; (ii) an organization with members pursuing contrasting agendas; (iii) a private firm trying to influence legislators and regulators. Intuitively, transparent motives hinder strategic communication because it is harder to sustain credible communication when the privately informed agent does not care about the true state and the principal's payoff.[1]

To improve communication between the two parties, we consider an *uninformed* third party who can mediate the transmission of information. This mediator cannot directly take the relevant decision in place of the receiver, but can credibly commit to information transmission policies, for example, because of their reputation. In the economic situations mentioned above, the role of the mediator can be taken respectively by (i) an advertiser that conveys credible information about the quality of the goods sold; (ii) a middle manager that is placed between an operative unit and the executive management of an organization; (iii) a professional lobbyist trying to influence the legislator pursuing the interest of a private firm. In all these cases, the mediator can use her commitment power to acquire relevant information in a credible way from the agent and transmit it to the principal. Moreover, the mediator's interest can be aligned with one of the two players or utterly independent of them (e.g., middle managers can have their own agendas).

We first analyze the optimal mediation problem. By the revelation principle, the mediator acts "as-if" selecting a *communication equilibrium* outcome of the sender-receiver game to maximize her expected payoff. We completely characterize the set of feasible distributions over the receiver's posterior beliefs under mediation in terms of simple moment conditions. This allows us to represent the optimal mediation problem as a linear program under moment constraints. Appealing to linear duality and the minimax theorem, we characterize the value of mediation as the lower envelope

[1] For example, in a one-dimensional setting where the agent's payoff is strictly increasing in the principal's action, no information transmission can be sustained in equilibrium, that is, the players are stuck in a babbling equilibrium.

of a class of distorted persuasion problems, where the distortion comes from the Lagrange multiplier of the truth-telling constraint. Moreover, we characterize optimal distributions over posteriors in terms of the complementary slackness properties of their supports.

We then compare the optimal outcome under mediation with other communication protocols such as Bayesian persuasion and pure cheap talk, providing insights into the role of mediators in communication. Our Theorem 1 shows that, when the sender and mediator are perfectly aligned, the sender can attain her optimal persuasion payoff under mediation if and only if this value can be attained under single-round cheap talk. Geometrically, this happens when the concave and quasi-concave envelopes of her value function coincide when evaluated at the prior. In other words, the comparison of the sender's optimal payoff under communication protocols with and without sender commitment is equivalent to the comparison of the concave and quasi-concave envelope of her value function.

Although mediation may not attain the optimal persuasion value, it can strictly improve communication by introducing additional randomness into the information structure. We characterize the value of this extra randomness in Proposition 2 and 3. When the state is binary, a mediator can credibly introduce extra randomness to strictly improve the sender's payoff only if the sender does not have an incentive to over/under-report the state compared to the optimal cheap talk value. We formalize the sender's relative misreport incentive by a weaker version of single crossing, which is called *mono-crossing*. When the babbling equilibrium is suboptimal under cheap talk, the sender attains her optimal mediation value under cheap talk if and only if her (shifted) value function is mono-crossing. In other words, the mediator's randomization adds value to communication if and only if the sender does not have a transparent tendency to misreport. When the babbling equilibrium is optimal under cheap talk, our duality results provide a characterization of babbling being optimal under mediation, which requires a distorted sender's utility being superdifferentiable at the prior.

We illustrate the geometric comparison of Bayesian persuasion, mediation, and cheap talk in Section 2 by a simple advertising model that compares the case where a seller directly communicates with a customer to the case where the seller hires an advertiser to mediate communication.

Literature review Our work adopts the belief-based approach, which has been widely used in the literature on communication games, for example, in Kamenica and Gentzkow (2011) and Lipnowski and Ravid (2020).[2] The former characterizes the sender's optimal payoff under persuasion as the concave envelope of the sender's value function, the second shows that the sender's best payoff under cheap talk with transparent motives is characterized by the quasi-concave envelope of her value function. We apply duality theory to characterize the set of optimal distributions over posterior beliefs under mediation. Similar techniques have been used in the Bayesian persuasion literature (Dworczak and Martini, 2019; Dworczak and Kolotilin, 2022; Kolotilin et al., 2022a; Arieli et al., 2023).

Our work is also closely connected to the vast literature on mediation, which uses very different tools, and focuses on particular cases such as the uniform-quadratic setting under the framework of Crawford and Sobel (1982) (e.g. Blume et al. (2007); Goltsman et al. (2009)) and finitely many actions (e.g. Salamanca (2021); Arieli et al. (2023)).[3] Many papers in this literature study the comparison between mediation, which is outcome equivalent to any communication equilibrium, and other specific forms of communication, such as noisy cheap talk (Blume et al., 2007), multiple-round cheap talk, and delegation (Goltsman et al., 2009). In this paper, we characterize the comparison between persuasion, mediation, and single-round cheap talk under sender's transparent motive.

Our work is related to recent papers studying Bayesian persuasion with limited commitment (Min, 2021; Lin and Liu, 2022; Lipnowski et al., 2022; Koessler and Skreta, 2021). Lin and Liu (2022) considers a Bayesian persuasion problem with limited commitment called "credible persuasion", which requires the sender to have no incentive to deviate to other experiments with the same marginal signal distribution. When the sender's payoff is state-independent, as in our main analysis, all Bayes-plausible policies are credible. Barros (2023) studies an information design problem where the sender commits to an experiment to acquire information and then cheap talks with the receiver. When the sender's payoff is state-independent, this design problem reduces to pure cheap talk.

Finally, this paper is connected to the literature on constrained information design

[2] See also Aumann and Maschler (1995); Aumann and Hart (2003).

[3] Arieli et al. (2022) studies a Bayesian persuasion problem with mediators that are not subject to truth-telling constraints.

(Doval and Smolin, 2021; Doval and Skreta, 2022). In particular, the truth-telling constraint has been considered in (Doval and Smolin, 2021; Doval and Skreta, 2022) as one of the potential additional constraints. These papers do not characterize the set of feasible distributions over posteriors and do not compare the solutions of the constrained problem with the solutions of the unconstrained one and with the cheap talk equilibria.

2 Illustrative Example

To illustrate the role of a mediator, consider a firm planning to commercialize a new product. The product's quality $\omega \in \Omega = \{0, 1\}$ is privately known by the firm, and a consumer has a prior $p \in (0, 0.55)$ on the quality being good ($\omega = 1$). We first consider the case when the firm can only communicate by cheap talk messages. After observing the message, the consumer updates her belief about the quality to $\mu \in [0, 1]$ and decides whether to purchase the good or take her outside option, which has value $\epsilon \sim G$, where G is the CDF of Beta(2,2). The price is normalized to $K > 0$ and the firm has reputation concerns, that is, the firm's indirect utility $\tilde{V}_S(\mu, \omega)$ given posterior μ and quality ω is

$$\tilde{V}_S(\mu, \omega) = KG(\mu) + \delta(\omega - \mu).$$

The linear term $\delta(\omega - \mu)$ captures the reputation effect, where $\delta > 0$ measures the positive effect of a surprisingly good product on the firm's future payoff. Conversely, when $\omega < \mu$, there is a negative reputation effect due to an unexpectedly bad product. As the state ω is known and the firm's payoff function is additively separable in $\tilde{V}_S(\mu, \omega)$, the firm acts to maximize

$$V_S(\mu) = KG(\mu) - \delta\mu.$$

With $K = 50/3$ and $\delta = 1.045K$,[4] this indirect utility V_S is a rotated S-shaped function as illustrated in Figure 1.[5] In particular, $V_S(0) = V_S(0.55) = 0, V_S(3/4) = 1,$

[4]One may also consider $K = 1, \delta = 1.045$. Our selection of parameters makes the calculation easier.

[5]Several recent papers in the persuasion literature focus on these S-shaped indirect utilities such as Kolotilin (2018); Kolotilin et al. (2022b); Arieli et al. (2023). Differently from those cases, here the indirect utility is non-monotone due to the reputation concern.

and $V_S(1) = -3/4$.

Applying the quasi-concave characterization in Lipnowski and Ravid (2020), we can see the best cheap talk equilibrium for the firm is to send messages $\mu = 0.55$ with probability $p/0.55$ and $\mu = 0$ with probability $1 - p/0.55$. Hence, the firm's optimal payoff under cheap talk is 0.

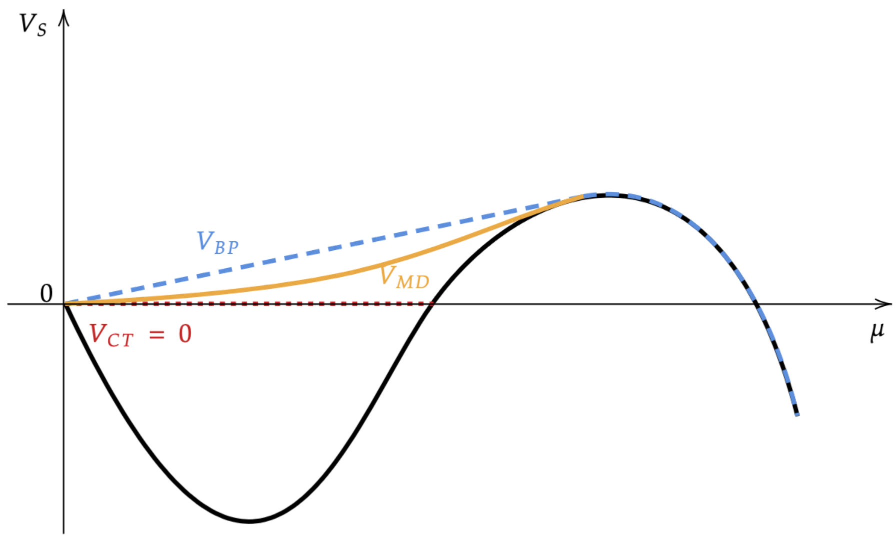

Figure 1: Comparison of Bayesian persuasion, mediation and cheap talk

Next, we illustrate that the firm can obtain a strictly higher payoff by hiring an advertiser (the mediator) who can credibly commit to revealing information about the quality of the good to consumers.[6] The advertiser does not have the expertise to assess the exact quality of the good and can only convey information the firm reports. To maintain credibility, the advertiser will design the information structure so that the firm is willing to report truthfully. The contract between the firm and the advertiser is fixed and binds the firm to pay the advertiser a fixed fraction of its revenue, so the advertiser's indirect payoff given posterior μ is $V_M(\mu) = kG(\mu)$ for some commission rate $k \in (0, K)$. In this case, the firm and advertiser's preferences are perfectly aligned (their expected payoffs are connected by positive affine transformations). And the advertiser can strictly increase the firm's expected payoff by introducing randomness to the message distribution conditional on the firm's quality report.

For instance, the advertiser can commit to the following signal structure:

[6]We assume the firm decides whether to hire a mediator before it learns the state ω, to avoid any additional signaling effects.

μ	0	3/4	1
$\omega = 1$	0	$\frac{9-9p}{21-25p}$	$\frac{12-16p}{21-25p}$
$\omega = 0$	$\frac{21-28p}{21-25p}$	$\frac{3p}{21-25p}$	0

where the first and second lines are the distribution over signals conditional on the firm's report $\omega = 1$ and 0, respectively. Direct calculation shows that the firm has no incentive to misreport, and the firm's expected payoff (before paying the advertiser) is $\frac{3p}{21-25p} > 0$. This payoff may not be the best payoff the mediator can secure but shows that with a small enough commission rate of k, the firm strictly benefits from hiring an advertiser to mediate communication. Note that this outcome is not a cheap talk equilibrium, as the firm would always have the incentive to report $\mu = 3/4$. In Section 5.1, we show similar constructions can provide a geometric characterization of conditions under which mediation leads to a strictly higher payoff than cheap talk.

Finally, if the mediator has the expertise to assess the quality of the goods without relying on the firm's reports, they design (and commit to) a test/information structure about the quality of the goods that will be revealed to the customer. The firm has a strict incentive to take this option, because it relaxes the truth-telling constraint (compared to mediation) and allows the firm to induce any Bayesian persuasion outcome. For instance, the firm can commit to sending messages $\mu = 3/4$ with probability $4p/3$ and $\mu = 0$ with probability $1 - 4p/3$. This information structure induces the optimal Bayesian persuasion outcome (one may verify this by concavification), and the optimal persuasion payoff $4p/3$ is greater than $\frac{3p}{21-25p}$, the payoff of the mediation plan we illustrated. We claim that this optimal persuasion payoff is indeed higher than the optimal mediation payoff, which will be illustrated in Section 5.1.

3 The Model

In this section, we introduce a model of mediated communication with transparent motives for the Sender. Let Ω be a finite set of states of the world with $|\Omega| = n$. The state is drawn according to a full-support common prior $p \in \Delta(\Omega)$. We first consider two agents: the Sender and the Receiver. The Sender is privately informed about the realization of the state $\omega \in \Omega$. The Receiver does not know the realized state but can take an action $a \in A$, where A is a compact metric space. The Sender can reveal information via some communication protocol, and the Receiver updates her belief once she obtains new information. The Sender's indirect utility given the Receiver's posterior belief is a continuous function $V_S : \Delta(\Omega) \to \mathbb{R}$.

Given this setting, there are two classical forms of communication between the Sender and the Receiver: Bayesian persuasion and cheap talk. In the former, the Sender can commit ex-ante to any Bayes-plausible information structure. That is, the Sender solves $\max_{\tau \in \mathcal{T}_{BP}(p)} \int V_S(\mu)\, \mathrm{d}\tau$, where $\mathcal{T}_{BP}(p) := \{\tau \in \Delta(\Delta(\Omega)) : \int \mu\, \mathrm{d}\tau = p\}$ is the set of all Bayes-plausible distribution over posteriors. In the latter, after having observed the state, the Sender sends a cheap talk message to the Receiver. As the Sender does not have commitment power, the feasible distributions over posteriors are those induced by PBEs. Hence, the set of feasible distribution is $\mathcal{T}_{CT}(p) := \{\tau \in \Delta(\Delta(\Omega)) : \int \mu\, \mathrm{d}\tau = p, V_S(\mu) = V_S(\mu') \ \forall \mu, \mu' \in \mathrm{supp}(\tau)\}$ by the characterization in Lipnowski and Ravid (2020). To compare with Bayesian persuasion, we assume the Sender can choose the equilibrium that maximizes her expected payoff.

Here, instead, we enlarge the standard model by considering a third-party Mediator whose continuous indirect utility is denoted by $V_M : \Delta(\Omega) \to \mathbb{R}$. The Mediator can commit to a communication mechanism but needs to gather information about the state from the Sender in an incentive-compatible way. By the transparent motive, the Sender's expected payoff should be the same conditional on every state report:

$$\int V_S(\mu)\, \mathrm{d}\tau_\omega(\mu) = \int V_S(\mu)\, \mathrm{d}\tau_{\omega'}(\mu) \qquad \forall \omega, \omega' \in \Omega, \tag{1}$$

where τ_ω denotes the conditional distribution over posteriors given the state ω. Therefore, the mediator's problem is

$$\sup_{\tau \in \mathcal{T}_{MD}(p)} \int V_M(\mu)\, \mathrm{d}\tau(\mu) \tag{P}$$

where $\mathcal{T}_{MD}(p) = \{\tau \in \Delta(\Delta(\Omega)) : \int \mu\, \mathrm{d}\tau = p, \tau \text{ satisfies } (1)\}$. Note that the sets of feasible distributions under these communication protocols are nested, i.e., $\mathcal{T}_{CT}(p) \subseteq \mathcal{T}_{MD}(p) \subseteq \mathcal{T}_{BP}(p)$. Therefore, mediation can be seen as an intermediate case between Bayesian persuasion and cheap talk.

We close this section by briefly discussing the assumptions on the indirect utility functions. In general, the indirect utility can be a Kakutani correspondence as in Lipnowski and Ravid (2020). We restrict our attention to continuous functions for simplicity. Our characterization of implementable distributions and the comparison of communication protocols can be generalized when the indirect utility is a correspondence, which will be illustrated in Appendix A. Note that the continuity assumption can remain valid even if the set of actions A is finite, for instance, our illustrative example has only two actions while V_S is continuous. In Section 6, we also discuss sufficient conditions on the primitives such that our assumption holds. For instance, if the Sender and the Receiver's payoff functions $u_S : A \to \mathbb{R}$, $u_R : A \times \Omega \to \mathbb{R}$ are continuous and the Receiver has a single best-response at any posterior μ, then the indirect utility V_S is a continuous function.

4 The Value of Mediation

In this section, we investigate the set of distributions over posterior $\tau \in \mathcal{T}_{MD}(p)$ that is implementable under mediation. Applying the implementation results, we characterize the value of mediation and the set of optimal distributions over posteriors under mediation by analyzing the dual problem of (P).

In the first subsection, we characterize $\mathcal{T}_{MD}(p)$ by simplifying the truth-telling constraint as a moment constraint in τ. An immediate consequence of this characterization is that optimal mediation value can be implemented with no more than $2n-1$ signals. We also provide an alternative characterization of the truth-telling condition as a zero covariance condition. This covariance condition illustrates the statistical difference between mediation and cheap talk, which requires full independence between the Sender's indirect utility and the Receiver's belief. We then illustrate that it is necessary for the mediator to introduce randomness to improve communication.

In the second subsection, we set up the dual problem of (P) and show there is no duality gap between the primal and the dual. Applying standard techniques in convex analysis, we show the value of mediation can be represented as the infimum of a family of distorted concavifications. Under a standard Slater's condition, we show dual attainment. Hence complementary slackness provides a characterization for the set of optimal solutions $\mathcal{T}^*_{MD}(p)$.

4.1 Implementation

The truth-telling constraint (1) involves conditional distributions τ_ω which is defined by the Radon-Nikodym derivative with respect to τ in accordance with the Bayes rule. Formally, for τ-almost every μ,

$$\frac{\mathrm{d}\tau_\omega}{\mathrm{d}\tau}(\mu) = \frac{\mu(\omega)}{p(\omega)} \qquad \forall \omega \in \Omega,$$

that is, the derivative of the conditional distribution given ω at μ is linear in the posterior.[7]

We use this observation to characterize the set of implementable distributions under mediation by rewriting the truth-telling constraint purely in terms of the unconditional distribution τ. With the Bayes-plausibility constraint, this rewritten constraint can be viewed as a zero covariance condition, which is the following proposition.

Proposition 1. *For every prior $p \in \Delta(\Omega)$, the following are equivalent:*

(i) *$\tau \in \mathcal{T}_{MD}(p)$, that is, τ satisfies $\int \mu \, \mathrm{d}\tau(\mu) = p$ and the truth-telling constraint (1);*

(ii) *τ satisfies $\int \mu d\tau(\mu) = p$ and $\int V_S(\mu)(\mu - p) \, d\tau(\mu) = \mathbf{0}$;*

(iii) *τ satisfies $\int \mu d\tau(\mu) = p$ and $\mathrm{Cov}_\tau(V_S, h) = 0$ for all affine continuous functions $h : \Delta(\Omega) \to \mathbb{R}$.*

An immediate observation from Proposition 1 is that $\mathcal{T}_{MD}(p)$ is compact because V_S is continuous. Hence, there exists $\tau^* \in \mathcal{T}_{MD}(p)$ that solves the mediator's problem (P). Another observation is that the transformed feasibility constraints are in the form of moment conditions à la Winkler (1988), which immediately implies that optimal mediation can be achieved with finitely many signals. We summarize the observations in the following corollary.

Corollary 1. *There exists a solution $\tau^* \in \mathcal{T}_{MD}(p)$ that attains the supremum value of (P), with $|\mathrm{supp}(\tau^*)| \leqslant 2n - 1$ and such that the vectors*

$$\left\{ \left((\mu(\omega))_{\omega \in \Omega \setminus \{\omega_0\}}, (V_S(\mu)(\mu(\omega) - p(\omega)))_{\omega \in \Omega \setminus \{\omega_0\}}, 1 \right)^T \right\}_{\mu \in \mathrm{supp}(\tau^*)}$$

are linearly independent, where $\omega_0 \in \Omega$ is an arbitrary state we fix.[8]

[7] Arieli et al. (2020) and Arieli et al. (2021) make a similar observation but do not use it to simplify a truthtelling constraint.

[8] We exclude one state from the linear independent condition because the components of posteriors μ add up to 1.

The second part of the statement will be useful for our comparison between mediation and cheap talk. Indeed, when there exists no distribution τ that is implementable under cheap talk while satisfying the independence restriction above, it must be that mediation is strictly better.

Next, we remark on one implementation property of cheap talk that can be immediately derived from the results in Lipnowski and Ravid (2020).

Remark. For every prior $p \in \Delta(\Omega)$, $\tau \in \mathcal{T}_{CT}(p)$ if and only if τ satisfies $\int \mu d\tau(\mu) = p$ and $\mathrm{Cov}_\tau(V_S, H) = 0$ for all continuous functions $H : \Delta(\Omega) \to \mathbb{R}$.

Therefore, the covariance condition in point (iii) of Proposition 1 is useful to compare the restrictions on implementable distributions imposed under mediation with the ones imposed under cheap talk. The only difference between implementable distributions under mediation and cheap talk is the set of random variables for which the zero-covariance condition must hold. Under cheap talk, the indirect payoff of the sender must be uncorrelated with any random variable that depends on the receiver's posterior, whereas, under mediation, the indirect payoff must be uncorrelated with expectations induced by the posterior.

Proposition 1 also allows us to formalize the idea that mediation can strictly improve on cheap talk only if randomization is needed. In other words, if a mediation plan postulates deterministic signals conditional on a report, then the induced distribution over posteriors would be implementable under cheap talk as well. Formally, we say that a distribution over posteriors τ is *deterministic* if $|\operatorname{supp} \tau_\omega| = 1$ for all $\omega \in \Omega$. Let $\mathcal{T}^\star_{MD}(p)$ and $\mathcal{T}^\star_{CT}(p)$ denote the set of deterministic implementable distributions under mediation and cheap talk respectively.

Corollary 2. *For every prior $p \in \Delta(\Omega)$, we have $\mathcal{T}^\star_{MD}(p) = \mathcal{T}^\star_{CT}(p)$.*

4.2 Optimal Mediation

The dual problem of (P) is to find two Lagrange multipliers $f, g \in \mathbb{R}^n$ that solves the following minimization problem:

$$\begin{aligned} &\inf_{f,g \in \mathbb{R}^n} \langle f, p \rangle \\ &\text{subject to:} \\ &\qquad \langle f, \mu \rangle \geqslant V_M(\mu) + \langle g, \mu - p \rangle V_S(\mu) \qquad \forall \mu \in \Delta(\Omega), \end{aligned} \tag{D}$$

where $\langle \cdot, \cdot \rangle$ stands for the standard inner product on $\mathbb{R}^n$ and we treat $\mu \in \Delta(\Omega)$ as vectors in the simplex Δ^{n-1}.

Given prior $p \in \Delta(\Omega)$, let $\widehat{V}(p)$ be the value of the primal (P) and $\bar{V}(p)$ be the value of its dual. We begin by showing weak duality, which can be viewed as a verification tool. Namely, for any feasible $\tau \in \mathcal{T}_{MD}(p)$, if there exists feasible f, g in (D) such that $\int V_M(\mu)\, \mathrm{d}\tau = \langle f, p \rangle$, then τ is optimal. Note that this result still holds under weaker technical assumptions, say, V_M, V_S can be upper semi-continuous.

Lemma 1 (Weak Duality)**.** *$\widehat{V}(p) \leqslant \bar{V}(p)$.*

The next result shows that (P) and (D) have the same value. The proof is based on the fact that there is no duality gap in the Bayesian persuasion problem (as shown in Dworczak and Kolotilin (2022)) and addresses the additional truth-telling constraint by using Sion's minimax Theorem.

Lemma 2 (No Duality Gap)**.** *$\widehat{V}(p) = \bar{V}(p)$.*

No duality gap has two useful implications. The first implication is that the value of mediation is the infimum of a family of distorted concavifications.

Corollary 3 (Value of Mediation)**.** *For any $g \in \mathbb{R}^n$, we define a distorted utility function $V_M^g(\mu) = V_M(\mu) + \langle g, \mu \rangle V_S(\mu)$. Then,*

$$\widehat{V}(p) = \inf_{g \in R^n : \langle g, p \rangle = 0} \operatorname{cav}(V_M^g)(p)$$

Another implication of no duality gap is that we can provide a verification result for a feasible $\tau \in \mathcal{T}_{MD}(p)$ to be optimal under mediation, which is also known as complementary slackness.

Corollary 4 (Complementary Slackness)**.** *For all $\tau \in \mathcal{T}_{MD}(p)$ and feasible f, g in (D), τ and f, g are optimal in (P) and (D) respectively if and only if*

$$\langle f, \mu \rangle = V_M(\mu) + \langle g, \mu - p \rangle V_S(\mu) \tag{2}$$

for all $\mu \in \operatorname{supp}(\tau)$.

To apply Corollary 4 as a verification tool, we need to guarantee dual attainment, that is, there exists feasible f, g that attain $\bar{V}(p)$.[9] To make full use of this verification

[9]If the infimum of (D) is not attained, then (2) is never satisfied. Hence, we cannot verify the optimality of $\tau \in \mathcal{T}_{MD}(p)$ by the complementary slackness if the dual is not attained.

approach, we further prove that there exists feasible f, g that attains $\bar{V}(p)$ under the following Slater's condition. To formally state the condition, we define $\mathcal{X} \subset \mathbb{R}^{2n+2}$ by $\mathcal{X} := \{x \in \mathbb{R}^{2n+2} : x_1 = \int V_M \, \mathrm{d}\tau, x_2 = \int V_S \, \mathrm{d}\tau, x_3 = \int \mu \, \mathrm{d}\tau, x_4 = \int \mu V_S \, \mathrm{d}\tau$ for some $\tau \in \Delta(\Delta(\Omega))\}$.

Slater's Condition. (P) satisfies Slater's condition at prior $p \in \Delta(\Omega)$ if there exists $x \in \operatorname{relint} \mathcal{X}$ that satisfies the mediation constraints, that is, $x_3 = p$ and $px_2 = x_4$.

Corollary 5 (Dual Attainment)**.** *If* (P) *satisfies Slater's condition at* $p \in \Delta(\Omega)$*, then there exists* $f^*, g^* \in \mathbb{R}^n$ *that attains the infimum value of the dual* $\bar{V}(p)$*.*

To see this, note that (P) can be transformed into a linear program in Euclidean space with equality constraints only. By Corollary 28.22 in Rockafellar (1970), if the Slater's condition is satisfied, then the dual (D) is attained.

Therefore, if (P) satisfies Slater's condition at $p \in \Delta(\Omega)$, we may characterize the set of optimal solutions as

$$\mathcal{T}^*_{MD}(p) = \{\tau \in \mathcal{T}_{MD}(p) : \exists f, g \text{ feasible in (D) s.t. (2) holds } \forall \mu \in \operatorname{supp}(\tau)\}. \quad (3)$$

5 Comparison of Communication Protocols

In this section, we compare the information designer's optimal value under Bayesian persuasion, mediation, and cheap talk, which are denoted as $\mathcal{V}^*_{BP}(V_M, p)$, $\mathcal{V}^*_{MD}(V_M, p)$, and $\mathcal{V}^*_{CT}(V_M, p)$, respectively. We start by analyzing the leading case when the Mediator and the Sender's preferences are perfectly aligned. The main economic insight we obtain is that mediation is as good as Bayesian persuasion if and only if cheap talk is as good as persuasion. Geometrically, this means the concave envelope and the quasi-concave envelope of the Sender's utility coincide at p. We also characterize geometrically when cheap talk is as good as mediation with a weaker version of single crossing property. In particular, this is the case if the Sender's utility V_S is concave, affine, convex, or weakly monotone. Finally, we relax the perfect alignment assumption and show that when the Mediator has a conflict of interest with the Sender, she may have a strict incentive to engage in the communication.

5.1 Perfectly Aligned Preference

As in the illustration example, the Mediator's preference is often aligned with the Sender's in many economic applications. We say their preferences are perfectly aligned when there exists a positive affine transformation between their expected payoffs. In this case, we may compare the Sender's optimal value among different communication protocols since the Mediator acts to maximize the Sender's payoff, and we suppress the argument V_M in the optimal value.

Note that under Bayesian persuasion and mediation, the feasible sets $\mathcal{T}_{BP}(p)$ and $\mathcal{T}_{MD}(p)$ are compact and convex, so by Bauer's maximum principle, the corresponding objective functions are maximized at an extreme point.[10] By Winkler (1988), we obtain a characterization of these extreme points. In particular, all extreme points have finite supports, so it is without loss of optimality to consider distributions with finite supports. We now compare the Sender's optimal value among different communication protocols using our characterization of optimal solutions and extreme points.

Our first result compares Bayesian persuasion and mediation and shows that mediation attains the optimal persuasion value if and only if this value can be attained under cheap talk. To see this, a mediator could potentially improve communication by randomizing over posteriors with different values while respecting the truth-telling constraint. The duality results for Bayesian persuasion shows that for any $\tau \in \mathcal{T}^*_{BP}(p)$, $V_S(\mu_i) = \langle f, \mu_i \rangle$ for some $f \in \mathbb{R}^n$ for all μ_i in the support of τ (Dworczak and Kolotilin, 2022). That is, $(\mu_i, V_S(\mu_i))_{\text{supp}(\tau)}$ lies on the same supporting hyperplane of the graph of $V_S(\mu)$ in $\Delta(\Omega) \times \mathbb{R}$. Jensen's inequality then shows the only hyperplane that satisfies truth-telling is the flat one when τ is nondegenerate.

Theorem 1. *If V_M, V_S are perfectly aligned, then*

$$\begin{aligned} \mathcal{V}^*_{BP}(p) = \mathcal{V}^*_{MD}(p) &\Longleftrightarrow \mathcal{V}^*_{BP}(p) = \mathcal{V}^*_{CT}(p) \\ &\Longleftrightarrow \text{cav}(V_S)(p) = \overline{\text{cav}}(V_S)(p), \end{aligned}$$

where $\text{cav}(V_S)$ and $\overline{\text{cav}}(V_S)$ stands for the concave and quasi-concave envelopes of V_S.

Theorem 1 provides a geometric comparison between the sender's optimal value under sender commitment and under any communication protocol that respects truth-

[10] The feasible set under cheap talk $\mathcal{T}_{CT}(p)$ is not convex.

telling. If the sender cannot achieve the optimal persuasion value using single-round cheap talk, then she cannot attain this value using any communication protocol without sender commitment (e.g. multiple-round cheap talk, noisy cheap talk), which can be replicated by a mediator by revelation principle. Connecting to the geometric characterization of the value of Bayesian persuasion and cheap talk (Kamenica and Gentzkow, 2011; Lipnowski and Ravid, 2020), we illustrate the communication value with and without sender commitment coincide if and only if the concave and quasi-concave envelope of V_S coincide at p.

An immediate consequence of Theorem 1 is that when the value of mediation and persuasion are the same, the optimal mediation value can be achieved by cheap talk. However, the converse of this statement is not true: It is possible that mediation and cheap talk have the same value while Bayesian persuasion has a strictly higher value. A quick example is to consider $V_S(\mu) = 4\mu(\mu - 1/2) + 1/4$ with $p = 1/4$. As illustrated in Figure 2, the concave envelope of this function is the line segment connecting $V_S(0)$ and $V_S(1)$, so the value at p is strictly higher than $V_S(0) = 1/4$. While the optimal mediation and cheap talk are attained by a 50-50 lottery supported on $\mu = 0$ and $1/2$, with value $V_S(0) = 1/4$. Indeed, the value of mediation coincides with the quasi-concave envelope of V_S by our next proposition.

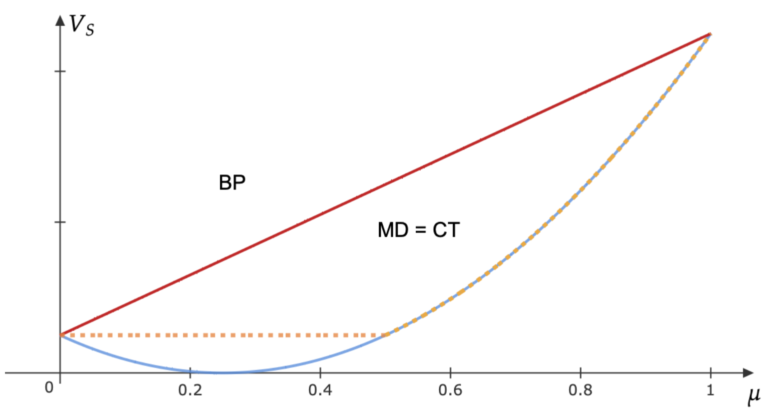

Figure 2

In this example, the Sender always has an incentive to over-report the state compared to the optimal cheap talk value. The truth-telling constraint then hinders the Mediator's ability to randomize over posteriors with different values. Proposition 3 formalizes this intuition with a weaker version of the single crossing condition, which is defined as follows.

We say a function $V : [0,1] \to \mathbb{R}$ is *mono-crossing from below (above)* if for any $\mu < \mu'$, $V(\mu) < (>)\, 0$ implies $V(\mu') \leqslant (\geqslant)\, 0$. And we say V is *mono-crossing* if it is mono-crossing either from below or from above.[11]

Proposition 2. *Suppose Ω is binary and V_M, V_S are perfectly aligned. If $V_S(\mu) - \mathcal{V}^*_{CT}(p)$ is mono-crossing, then $\mathcal{V}^*_{MD}(p) = \mathcal{V}^*_{CT}(p)$. Moreover, the converse is true if no disclosure is suboptimal under cheap talk.*

Proposition 2 shows that when the state is binary, for any communication protocol with truth-telling (equivalently a Mediator) to strictly improve the Sender's value compared to cheap talk, it is necessary that the Sender does not have a transparent tendency to misreport. This misreporting tendency is captured by the mono-crossing property. By Lipnowski and Ravid (2020), $V_S(\mu) \leqslant \mathcal{V}^*_{CT}(p)$ on at least one of $[0,p)$ or $(p,1]$. In the former case, the Sender always prefers to over-claim the state if her preference is mono-crossing from below. For an example and a non-example, recall the example plotted in Figure 2 and the illustrative example in Section 2. Figure 3 plots the shifted Sender's utility in these examples at prior $p = 1/4$.

As illustrated in the left panel, the shifted Sender's utility in our previous example is mono-crossing since its sign only changes at $\mu = 0.5$. So the Sender has a tendency to over-report the state. The (shifted) truth-telling constraint $\int (V_S(\mu) - \mathcal{V}^*_{CT}(p))(\mu - p)\, \mathrm{d}\tau = 0$ then requires that for any $p < \mu_1 < 0.5 < \mu_2$, a truthful τ must assign a higher weight on μ_1 than μ_2 due to this over-claim tendency. Therefore, all the Sender-preferred randomizations are ruled out.

In our illustrative example (right panel), the shifted Sender's utility is not mono-crossing since its sign changes at $\mu = 0.55$ and 0.95. We have seen in Section 2 there exists a $\tau \in \mathcal{T}_{MD}(p)$ with support $\{0, 3/4, 1\}$ that strictly increases Sender's payoff compared to cheap talk. In particular, when the Sender does not have a transparent tendency to over-report, the Mediator can assign a higher probability to $\mu = 3/4$ while maintaining truthful report by assigning a (relatively) lower probability to $\mu = 1$.

Moreover, we observe that the construction in the not mono-crossing case above is always feasible when $V_S(p) < \mathcal{V}^*_{CT}(p)$. If no disclosure is optimal under cheap talk,

[11] This property is called weak single crossing in Shannon (1995), which requires a real-valued function f on the product of two partially ordered sets $X \times T$ has mono-crossing differences, that is, $f(x',t) - f(x,t)$ is mono-crossing in t for any $x < x'$. This notion is a weaker version of the standard single crossing property, for instance, a constant zero function is not single crossing but mono-crossing.

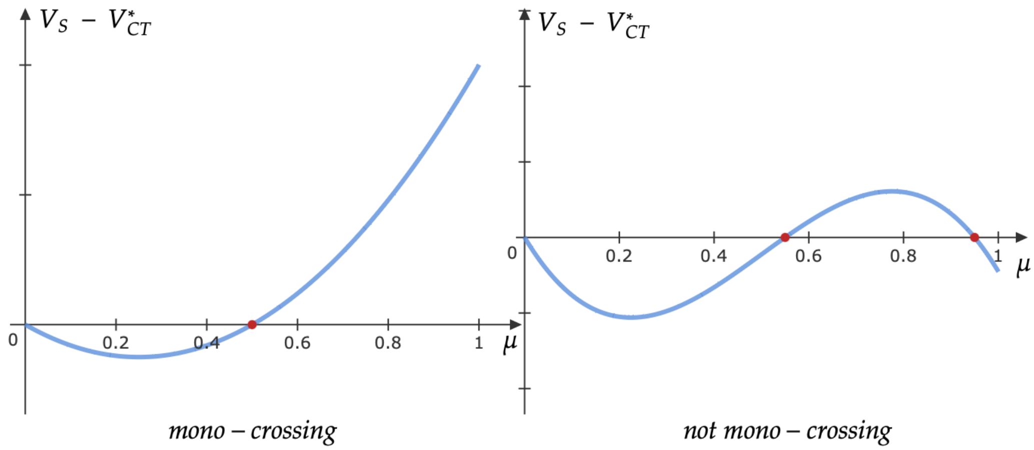

Figure 3

then the construction may fail to be Bayes-plausible. Therefore, we obtain a complete characterization of $\mathcal{V}_{MD}^*(p) = \mathcal{V}_{CT}^*(p)$ in the binary case by discussing two disjoint cases separately. If no disclosure is optimal under cheap talk, by our characterization (3) of $\mathcal{T}_{MD}^*(p)$ via duality, it is optimal under mediation if and only if there exists $g \in \mathbb{R}$ such that the distorted value function $(1 + g(\mu - p))V_S(\mu)$ is superdifferentiable at p. If no disclosure is suboptimal under cheap talk, then it suffices to check whether the shifted utility function $V_S(\mu) - \mathcal{V}_{CT}^*(p)$ is mono-crossing.

Proposition 3. *Suppose Ω is binary and V_M, V_S are perfectly aligned. $\mathcal{V}_{MD}^*(p) = \mathcal{V}_{CT}^*(p)$ if and only if either of the following is true:*

(i) *No disclosure is optimal under mediation;*

(ii) *No disclosure is suboptimal under cheap talk and $V_S(\mu) - \mathcal{V}_{CT}^*(p)$ is mono-crossing.*

We then apply Proposition 2 and 3 to study some natural Sender utility functions for which cheap talk and mediation are equivalent in value. It is straightforward that if V_S is concave or affine over $[0, 1]$, it is identical to its concave envelope and hence no disclosure is optimal. Note that any convex function on $[0, 1]$ satisfies the mono-crossing condition, so Proposition 2 implies that cheap talk attains the optimal mediation value. When V_S is weakly monotone, we observe that the feasible set under mediation and cheap talk are the same. That is, when the Sender always prefers to over/under-report the state, the Mediator cannot credibly introduce any extra randomness into communication.

Proposition 4. *Suppose Ω is binary and V_M, V_S are perfectly aligned.*

(i) *If V_S is concave or convex, then $\mathcal{V}^*_{MD}(p) = \mathcal{V}^*_{CT}(p)$ for all $p \in (0,1)$.*

(ii) *If V_S is (weakly) monotone, then $\mathcal{T}^*_{MD}(p) = \mathcal{T}^*_{CT}(p)$ for all $p \in (0,1)$.*

These sufficient conditions are useful when we compare the value of mediation and cheap talk in some simple settings. For instance, condition (ii) immediately shows that the seller's reputation concern in Section 2 is crucial for them to be willing to hire an advertiser. Also, (i) immediately implies the value of mediation and cheap talk are the same in the example in Figure 2.

5.2 Non-perfectly Aligned Preference

In many economic applications, it is common for the Sender and Mediator to have perfectly aligned preferences, especially when the Sender hires the Mediator and shares her profit as in our illustrative example. However, it is also natural to consider a Mediator (more generally information designer) that has conflicting interests with the Sender.

For instance, instead of considering an advertiser in our illustrative example, we can consider a market regulator who maximizes the aggregate welfare of both parties and does not have the ability or formal authority to verify the quality of the good. Compared to the advertiser, the regulator may have a stronger incentive to intervene in the communication between the seller and buyer. Unlike the perfectly aligned case, the Mediator can attain the optimal persuasion value while strictly improving on cheap talk. We illustrate this with our next proposition, which compares the Mediator's optimal value under Bayesian persuasion and mediation when the state is binary and preferences are not perfectly aligned.

Proposition 5. *Suppose Ω is binary. $\mathcal{V}^*_{BP}(V_M, p) = \mathcal{V}^*_{MD}(V_M, p)$ if and only if one of the following is true:*

(i) *V_M is superdifferentiable at p;*

(ii) *There exists $\tau \in \mathcal{T}_{CT}(p)$ with binary support $\{\mu_1, \mu_2\}$ such that the line connecting $(\mu_1, V_M(\mu_1))$ and $(\mu_2, V_M(\mu_2))$ lies weakly above the graph of $V_M(\mu)$;*

(iii) *There exists* $\tau \in \mathcal{T}_{MD}(p) \backslash \mathcal{T}_{CT}(p)$ *that supports on* $\mu_1 < \mu_2 < \mu_3$ *with* $p \neq \mu_2$ *such that the line connecting* $(\mu_1, V_M(\mu_1))$ *and* $(\mu_3, V_M(\mu_3))$ *lies weakly above the graph of* $V_M(\mu)$ *and coincides with* V_M *at* μ_2*.*

In particular, if $\mathcal{V}^*_{BP}(V_M, p) = \mathcal{V}^*_{MD}(V_M, p)$ *while the conditions (i) and (ii) do not hold, then* $\mathcal{V}^*_{MD}(V_M, p) > \mathcal{V}^*_{CT}(V_M, p)$*.*

This proposition shows that perfectly aligned preferences are necessary for our observation in Proposition 1. Indeed, for an arbitrary indirect utility of the Mediator, if the value of mediation is the same as in Bayesian persuasion (e.g., when condition (iii) holds), and neither (i) nor (ii) holds, then mediation is strictly better than cheap talk.

To prove Proposition 5, we first observe that mediation attains the optimal Bayesian persuasion value if and only if there exists an extreme point of $\mathcal{T}_{MD}(p)$ that maximizes the mediator's expected payoff among all the Bayes-plausible distributions of beliefs. As in Corollary 1, we can use Winkler (1988) to characterize these extreme points. In particular, as the state is binary, it suffices to check $\tau \in \mathcal{T}_{MD}(p)$ with $|\operatorname{supp}(\tau)| \leqslant 3$ and satisfies a linear-independence condition. Next, conditions (ii) and (iii) respectively establish that an extreme point with two or three points in the support satisfies the complementary slackness condition characterizing optimality under Bayesian persuasion (see e.g. Dworczak and Kolotilin (2022)).

Our next proposition characterizes when the Mediator has no strict incentive to mediate when the state is binary.

Proposition 6. *Suppose* Ω *is binary. Then* $\mathcal{V}^*_{MD}(V_M, p) = \mathcal{V}^*_{CT}(V_M, p)$ *if and only if either of following is true:*

(i) $\delta_p \in \mathcal{T}^*_{MD}(p)$

(ii) *There exists* $\tau \in \mathcal{T}_{CT}(p)$ *with binary support and* $f \in \mathbb{R}$ *such that*

$$V_M(\mu) - \sum \tau_i V_M(\mu_i) \leqslant \left[f \left(1 - \frac{V_S(\mu)}{V_S(\mu_1)} \right) + \frac{V_S(\mu)(1-\tau_1)(V_M(\mu_1) - V_M(\mu_2))}{\mu_1 - p} \right] (\mu - p).$$

for all $\mu \in [0, 1]$.

Since the set of feasible distributions under mediation and cheap talk are nested, $\mathcal{V}^*_{MD}(V_M, p) = \mathcal{V}^*_{CT}(V_M, p)$ if and only if there exists a cheap talk equilibrium inducing

$\tau \in \mathcal{T}_{CT}(p)$ that attains the optimal mediation value. Moreover, we apply the characterization of extreme points to show there exists such a τ which is also an extreme point of $\mathcal{T}_{BP}(p)$. As the state is binary, it suffices to check such τ with $|\operatorname{supp}(\tau)| \leqslant 2$. Condition (ii) characterizes when a $\tau \in \mathcal{T}_{CT}(p)$ with binary support satisfies the complementary slackness condition characterizing optimality under mediation (Corollary 4).

6 Extension and Discussion

In this section, we discuss our main assumptions and sketch some attempts to relax our assumptions. Throughout the paper, we adopt the belief-based approach and work with the indirect utilities V_S and V_M. As we discussed at the end of Section 3, it is possible to derive the indirect utilities by fixing an action space A for the receiver and continuous explicit payoff functions $u_S : A \to \mathbb{R}$, $u_R : A \times \Omega \to \mathbb{R}$ and $u_M : A \times \Omega \to \mathbb{R}$ for the Sender, Receiver, and Mediator, respectively. To construct the indirect utility functions, we calculate the Sender and the Mediator's payoff at each posterior μ, condition on the Receiver taking her best response to μ. This induces two indirect payoff correspondences $\mathbf{V}_S$ and $\mathbf{V}_M$ for the Sender and Mediator, respectively. In particular, when the Receiver has a single best response for every posterior, the Sender and Mediator's indirect payoff correspondences reduce to continuous indirect utility functions V_S, V_M, which is the main assumption of this paper. A natural example of this case is the linear persuasion case in Appendix B, namely, when $\Omega, A \subseteq \mathbb{R}^k$ and the Receiver's best-response is the posterior mean.

There are several approaches to relax our main assumption that V_S, V_M are continuous real-valued functions. The first is to work with the payoff correspondences $\mathbf{V}_S, \mathbf{V}_M$ directly. We can extend Proposition 1 in terms of these payoff correspondences, and the qualitative difference in implementation between mediation and cheap talk in terms of zero-covariance conditions remains to be true. Moreover, the comparison results in the perfectly aligned case extend to this general case (Theorem 1, Proposition 2 and 3), so the economic insights of our main analysis are without loss of generality.

The first approach is not very efficient in computing the optimal value of mediation, because the truth-telling condition expressed in terms of correspondence $\mathbf{V}_S$ makes it hard to analyze the mediation problem as an (infinite dimensional) linear program. In order to analyze the optimal value of mediation, we consider a richer space. The second approach we attempt is to consider joint distributions $\beta \in \Delta(A \times \Delta(\Omega))$ of the Receiver's actions and posterior beliefs. This allows us to express the Sender's truth-telling constraint and the Receiver's obedience constraint as linear constraints in β, and hence we can analyze the mediation problem as a linear program. This approach can also be applied to study delegation models, in which the Mediator has full control of the Receiver's action and can commit to a distribution of actions given the Sender's reports. We will explore these extensions in our future work and include delegation into our analysis.

References

Aliprantis, C. D. and K. C. Border (2006): *Infinite Dimensional Analysis: a Hitchhiker's Guide*, Berlin; London: Springer.

Arieli, I., Y. Babichenko, and F. Sandomirskiy (2022): "Bayesian persuasion with mediators," *arXiv preprint arXiv:2203.04285*.

Arieli, I., Y. Babichenko, F. Sandomirskiy, and O. Tamuz (2021): "Feasible joint posterior beliefs," *Journal of Political Economy*, 129, 2546–2594.

Arieli, I., Y. Babichenko, and R. Smorodinsky (2020): "Identifiable information structures," *Games and Economic Behavior*, 120, 16–27.

Arieli, I., Y. Babichenko, R. Smorodinsky, and T. Yamashita (2023): "Optimal persuasion via bi-pooling," *Theoretical Economics*, 18, 15–36.

Aumann, R. J. and S. Hart (2003): "Long cheap talk," *Econometrica*, 71, 1619–1660.

Aumann, R. J. and M. Maschler (1995): *Repeated games with incomplete information*, MIT press.

Barros, L. (2023): "Information Acquisition Design," .

Blume, A., O. J. Board, and K. Kawamura (2007): "Noisy talk," *Theoretical Economics*, 2, 395–440.

Chakraborty, A. and R. Harbaugh (2010): "Persuasion by cheap talk," *American Economic Review*, 100, 2361–82.

Crawford, V. P. and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.

Doval, L. and V. Skreta (2022): "Constrained information design: Toolkit," *arXiv preprint arXiv:1811.03588*.

Doval, L. and A. Smolin (2021): "Information payoffs: An interim perspective," *arXiv preprint arXiv:2109.03061*.

Dworczak, P. and A. Kolotilin (2022): "The Persuasion Duality," *SSRN Electronic Journal*.

Dworczak, P. and G. Martini (2019): "The simple economics of optimal persuasion," *Journal of Political Economy*, 127, 1993–2048.

Gentzkow, M. and E. Kamenica (2016): "A Rothschild-Stiglitz approach to Bayesian persuasion," *American Economic Review*, 106, 597–601.

Goltsman, M., J. Hörner, G. Pavlov, and F. Squintani (2009): "Mediation, arbitration and negotiation," *Journal of Economic Theory*, 144, 1397–1420.

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

Koessler, F. and V. Skreta (2021): "Information design by an informed designer," .

Kolotilin, A. (2018): "Optimal information disclosure: A linear programming approach," *Theoretical Economics*, 13, 607–635.

Kolotilin, A., R. Corrao, and A. Wolitzky (2022a): "Persuasion with Non-Linear Preferences," ArXiv: 2206.09164 [econ] version: 2.

Kolotilin, A., T. Mylovanov, and A. Zapechelnyuk (2022b): "Censorship as optimal persuasion," *Theoretical Economics*, 17, 561–585.

Lin, X. and C. Liu (2022): "Credible Persuasion," *arXiv preprint arXiv:2205.03495*.

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

Lipnowski, E., D. Ravid, and D. Shishkin (2022): "Persuasion via weak institutions," *Journal of Political Economy*, 130, 2705–2730.

Min, D. (2021): "Bayesian persuasion under partial commitment," *Economic Theory*, 72, 743–764.

Molchanov, I. S. (2005): *Theory of random sets*, vol. 19, Springer.

Rockafellar, R. T. (1970): *Convex analysis*, vol. 18, Princeton university press.

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

Shannon, C. (1995): "Weak and strong monotone comparative statics," *Economic Theory*, 5, 209–227.

Strassen, V. (1965): "The existence of probability measures with given marginals," *The Annals of Mathematical Statistics*, 36, 423–439.

Winkler, G. (1988): "Extreme points of moment sets," *Mathematics of Operations Research*, 13, 581–587.

A Correspondence-valued Indirect Utility

Let the action space A be a compact metric space and fix explicit payoff functions $u_S : A \to \mathbb{R}$, $u_R : A \times \Omega \to \mathbb{R}$, and $u_M : A \times \Omega \to \mathbb{R}$ for the sender, receiver, and mediator respectively. We assume each explicit payoff function is continuous in the receiver's action. With this primitive setting, we can define the best-response correspondence $\mathbf{A}_R : \Delta(\Omega) \rightrightarrows A$ as

$$\mathbf{A}_R(\mu) = \underset{a \in A}{\operatorname{argmax}} \sum_{\omega \in \Omega} u_R(a, \omega)\mu(\omega).$$

In turn this induces value correspondences $\mathbf{V}_S : \Delta(\Omega) \rightrightarrows \mathbb{R}$ and $\mathbf{V}_M : \Delta(\Omega) \rightrightarrows \mathbb{R}$ as follows

$$\mathbf{V}_S(\mu) = \operatorname{co}\left(u_S\left(\mathbf{A}_R(\mu)\right)\right) \quad \text{and} \quad \mathbf{V}_M(\mu) = \operatorname{co}\left(\sum_{\omega \in \Omega} u_M\left(\mathbf{A}_R(\mu), \omega\right)\mu(\omega)\right)$$

where the sum in the second expression denotes the Minkowski addition operation. In words, $\mathbf{V}_S(\mu)$ and $\mathbf{V}_M(\mu)$ respectively denote the (convex hull of the) sets of attainable payoff for the sender and the mediator given the receiver's posterior. By the maximum theorem, these correspondences are Kakutani correspondences, and the value envelope $v_S(\cdot) := V_S(\cdot)$ is upper semi-continuous.

When we assume that the receiver has a single best response for every posterior, the correspondence $\mathbf{A}_R = \mathbf{a}_R$ reduces to a continuous function, and the indirect payoffs reduce to $V_S(\mu) = u_S(\mathbf{a}_R(\mu))$ and $V_M(\mu) = \sum_{\omega \in \Omega} u_M\left(\mathbf{a}_R(\mu), \omega\right)\mu(\omega)$.[12] In this case, the indirect payoffs are continuous and all the results presented in the main analysis hold. This is for example the case when $\Omega, A \subseteq \mathbb{R}^k$ and $\mathbf{a}_R(\mu) = \mathbb{E}_\mu(\omega)$ which corresponds to the linear persuasion case.[13]

When the receiver has multiple best responses for some posterior, then we need to postulate a tie-breaking rule which usually is the one selecting the designer-preferred action. However, this is not an issue as long as we are interested in the space of all implementable distributions over posteriors. Indeed, we can extend Proposition 1

[12] Continuity of $\mathbf{a}_R$ follows by Berge's maximum theorem. See Kolotilin et al. (2022a) for sufficient conditions on the primitive payoff of the receiver for the best-response correspondence to be single-valued.

[13] This best response is induced, for example, when $u(a, \omega) = -||\omega - a||_2^2$. See Chakraborty and Harbaugh (2010) for several examples where the indirect payoff of the sender is well-defined, continuous, and only depends on conditional expectations.

as follows. First, following the standard definition of Aumann integral of correspondences (see Molchanov (2005)), we define

$$\int \mathbf{V}_S(\mu)\,(\mu - p)\,d\tau\,(\mu) = \left\{ \int V_S(\mu)\,(\mu - p)\,d\tau\,(\mu) \in \mathbb{R}^n : V_S \in \mathbf{V}_S \right\}$$

where $V_S \in \mathbf{V}_S$ are arbitrary measurable selections from $\mathbf{V}_S$.

Proposition 7. *For every prior $p \in \Delta(\Omega)$, the following are equivalent:*

(i) $\tau \in \mathcal{T}_{MD}(p)$;

(ii) τ *satisfies* $\int \mu d\tau\,(\mu) = p$ *and* $\mathbf{0} \in \int \mathbf{V}_S(\mu)\,(\mu - p)\,d\tau\,(\mu)$;

(iii) τ *satisfies* $\int \mu d\tau\,(\mu) = p$ *and there exists a measurable selection* $V_S \in \mathbf{V}_S$ *such that* $\mathrm{Cov}_\tau\,(V_S, h) = 0$ *for all affine continuous functions* $h : \Delta(\Omega) \to \mathbb{R}$.

The qualitative difference in implementation between mediation and cheap talk remains the same even when there are multiple best responses. Indeed, it is immediate to derive the following results from the results in Lipnowski and Ravid (2020).

Remark. For every prior $p \in \Delta(\Omega)$, $\tau \in \mathcal{T}_{CT}(p)$ if and only if τ satisfies $\int \mu d\tau\,(\mu) = p$ and there exists a measurable selection $V_S \in \mathbf{V}_S$ such that $\mathrm{Cov}_\tau\,(V_S, H) = 0$ for all continuous functions $H : \Delta(\Omega) \to \mathbb{R}$.

With Proposition 7, we may express the mediator's problem as follows

$$\begin{aligned} &\sup_{V_S,\tau} \int \mathbf{V}_M(\mu)\,\mathrm{d}\tau(\mu) \\ \text{subject to: } &\int \mu\,\mathrm{d}\tau = p \\ &\int V_S(\mu)(\mu - p)\,\mathrm{d}\tau = 0 \\ &V_S \text{ is a measurable selection from } \mathbf{V}_S. \end{aligned}$$

When the Sender and Mediator have aligned preference $V_M(\mu) = V_S(\mu)$ for any selection V_S that Mediator specifies, our comparison results in Section 5.1 still holds in this more general case. For instance, our Theorem 1 still holds.

Theorem 2. *Suppose Ω is finite and $\mathbf{V}_S$ is a Kakutani correspondence. If the Sender and Mediator are perfectly aligned, then*

$$\mathcal{V}^*_{BP}(p) = \mathcal{V}^*_{MD}(p) \Longleftrightarrow \mathcal{V}^*_{BP}(p) = \mathcal{V}^*_{CT}(p).$$

The proof of this result is essentially the same as Proposition 1. To solve the Bayesian persuasion problem with $\mathbf{V}_S$, it suffices to consider the upper envelope v_S of $\mathbf{V}_S$, which is upper semi-continuous. All the duality results for the Bayesian persuasion problem with indirect utility v_S are not affected even if we relax the continuity to semi-continuity (see Dworczak and Kolotilin (2022)). We then use the duality argument in Proposition 1 to show an optimal Bayesian persuasion solution τ is feasible under mediation (with selection v_S) if and only if v_S is constant over the support of τ.

Our geometric comparison of mediation and cheap talk in the binary case also extends when the Sender's indirect value is a correspondence.

Proposition 8. *Suppose Ω is binary, $\mathbf{V}_S$ is a Kakutani correspondence, and the Sender and Mediator are perfectly aligned. If for any measurable selection V_S from $\mathbf{V}_S$, $V_S - \mathcal{V}^*_{CT}(p)$ is mono-crossing, then $\mathcal{V}^*_{MD}(p) = \mathcal{V}^*_{CT}(p)$. Moreover, the converse is true if $v_S(p) < \mathcal{V}^*_{CT}(p)$.*

In other words, $\mathcal{V}^*_{MD}(p) = \mathcal{V}^*_{CT}(p)$ if and only if either of the following is true:

(i) *$v_S(p) = \mathcal{V}^*_{CT}(p)$ and $\delta_p \in \mathcal{T}^*_{MD}(p)$;*

(ii) *$v_S(p) < \mathcal{V}^*_{CT}(p)$ and for any measurable selection V_S from $\mathbf{V}_S$, $V_S - \mathcal{V}^*_{CT}(p)$ is mono-crossing.*

The only difference from Proposition 3 is the superdifferentiability condition (i) that is equivalent to no disclosure being optimal under mediation. With correspondence $\mathbf{V}_S$, we are not sure how to use duality results for mediation to verify the optimality of no disclosure.

B Linear Persuasion

In many economic settings (e.g. our illustrative example), the payoffs of the Sender and Mediator depend only on the posterior mean of the Receiver's belief. In this case, we focus on Euclidean state spaces $\Omega \subseteq \mathbb{R}^k$. Let $X := \operatorname{co}\Omega$ be the space of all possible posterior means. The Sender and Mediator have indirect utilities $v_s, v_m : X \to \mathbb{R}$. Note that this is a special case of our general model as we may write $V_S(\mu) = v_s(\mathbb{E}_\mu(\omega))$. Instead of considering distributions over posterior $\tau \in \Delta(\Delta(\Omega))$, the Mediator now focuses on the induced distributions of posterior means $q(x) \in \Delta(X)$. We say $q \in \Delta(X)$ is implementable under mediation if there exists $\tau \in \mathcal{T}_{MD}(p)$ that induces q. The mediator is maximizing her expected payoff among all implementable distributions over posterior means, which is denoted by $\mathcal{T}^m_{MD}(p)$.

If the Sender's payoff v_s only depends on the posterior mean of the Receiver's belief, then we can specialize Proposition 1 as follows. For any $q \in \Delta(X)$ and $v_s : X \to \mathbb{R}$, define the corresponding distorted distribution $q^s \in \Delta(X)$ by

$$\frac{\mathrm{d}q^s}{\mathrm{d}q}(x) = \frac{v_s(x)}{\int v_s(z)\,\mathrm{d}q(z)}.$$

And the set of implementable distributions over posterior means under mediation is characterized as follows.

Proposition 9. *Let $V_S(\mu) = v_s(\mathbb{E}_\mu(\omega))$. For any full support prior $p \in \Delta(\Omega)$, the following are equivalent:*

(i) *q is implementable under mediation;*

(ii) *There exists a dilation[14] $T : X \to \Delta(X)$ such that $Tq = Tq^s = p$;*

(iii) *There exists $\pi \in \Delta(\Omega \times X)$ such that $\operatorname{marg}_\Omega \pi = p$, $\operatorname{marg}_X \pi = q$, $\mathbb{E}_\pi[\omega|x] = x$ for all x, and $\operatorname{Cov}_\pi(v_s, g) = 0$ for all $g \in \mathbb{R}^\Omega$.*

Note that when there is no truth-telling constraint, by Strassen's Theorem,[15] condition (ii) reduces to the Bayes-plausibility condition in the linear persuasion literature, which is $q \preceq_{cvx} p$. With the truth-telling constraint, Strassen's Theorem implies both q and q^s are mean-preserving contractions of p.

[14] A map $T : X \to \Delta(X)$ is called a dilation if $x = \int y\,\mathrm{d}T_x(y)$ for all x, and the map $x \mapsto T_x(f)$ is measurable for all $f \in C(X)$. The product Tq is defined as by $Tq(S) = \int T_x(S)\,\mathrm{d}q(x)$ for all measurable $S \subseteq X$.

[15] Let X be a compact convex metrizable space and p, q are Borel probability measures on X. The Strassen's Theorem states that $q \preceq_{cvx} p$ if and only if there exists a dilation T such that $p = Tq$, see Strassen (1965); Aliprantis and Border (2006). This result has been widely applied in the linear persuasion literature, see Gentzkow and Kamenica (2016); Kolotilin (2018); Dworczak and Martini (2019).

Corollary 6. *Let $V_S(\mu) = v_s(\mathbb{E}_\mu(\omega))$. For any full support prior $p \in \Delta(\Omega)$, if $q \in \Delta(X)$ is induced by some $\tau \in \mathcal{T}_{MD}(p)$, then $q \preceq_{cvx} p$, $q^s \preceq_{cvx} p$, and*

$$\int v_s(x)(x - \overline{x}_p)\, dq(x) = 0,$$

where $\overline{x}_p$ is the prior expectation.

This corollary is useful for comparing the set of distributions over posteriors implementable under mediation with the ones corresponding to Bayesian persuasion and cheap talk, as the next corollary shows. When $V_S(\mu) = v_s(\mathbb{E}_\mu(\omega))$, we say that $q \in \Delta(X)$ is implementable under cheap talk if it is induced by $\tau \in \mathcal{T}_{CT}(p)$.

Corollary 7. *Let $\Omega \subseteq \mathbb{R}$ and $V_S(\mu) = v_s(\mathbb{E}_\mu(\omega))$ with v_s either increasing or decreasing. Then $q \in \Delta(X)$ is implementable under mediation if and only if it is implementable under cheap talk.*

This result extends our observation about implementation in Proposition 4 to the linear case. Namely, when the Sender always prefers to induce a higher/lower posterior mean, the Mediator cannot credibly introduce any extra randomness into communication.

C Proofs

C.1 Implementation

Proof of Proposition 1. We first show that (i) and (ii) are equivalent. Take any $\tau \in \mathcal{T}_{MD}(p)$, the truth-telling constraint (1) implies

$$\int V_S(\mu)\mu(\omega)\, \mathrm{d}\tau(\mu) = p(\omega) \cdot c$$

for some constant $c \in \mathbb{R}$ for all $\omega \in \Omega$. Taking summation across ω, we have $c = \int V_S(\mu)\, \mathrm{d}\tau(\mu)$, which implies (ii) holds. Conversely, for any $\tau \in \Delta(\Delta(\Omega))$ that satisfies (ii), we have

$$\int V_S(\mu)(\mu(\omega) - p(\omega))\, \mathrm{d}\tau(\mu) = 0$$

for all $\omega \in \Omega$. Therefore, $\int V_S(\mu)\frac{\mu(\omega)}{p(\omega)}\,\mathrm{d}\tau(\mu) = \int V_S(\mu)\,\mathrm{d}\tau(\mu)$ for all $\omega \in \Omega$ and the truth-telling condition holds.

Next, we show (ii) and (iii) are equivalent. Suppose τ satisfies condition (ii), then for any affine function $h(\mu) = \mathbf{a}^T\mu + b$ with $\mathbf{a} \in \mathbb{R}^n, b \in \mathbb{R}$, we have

$$\int V_S(\mu)(\mathbf{a}^T\mu+b)\,\mathrm{d}\tau = \mathbf{a}^T\int V_S(\mu)\mu\,\mathrm{d}\tau + b\int V_S(\mu)\,\mathrm{d}\tau = \left(\int V_S(\mu)\,\mathrm{d}\tau\right)\left(\int \mathbf{a}^T\mu + b\,\mathrm{d}\tau\right),$$

where the second equality is by Bayes-plausibility and truth-telling. Conversely, notice that the projection map $h_i(\mu) = \mu_i$ is affine for all $i = 1, \ldots, n$, so the zero covariance condition implies the truth-telling constraint. ■

Proof of Corollary 1. Note that the maps $\tau \mapsto \int \mu(\omega) - p(\omega)\,\mathrm{d}\tau$ and $\tau \mapsto \int V_S(\mu)(\mu(\omega) - p(\omega))\,\mathrm{d}\tau$ are continuous in τ for all ω because $\mu - p$ and $(\mu - p)V_S$ are continuous in μ. Hence the feasible region is a finite intersection of closed pre-images, which is closed. Note that $\Delta\Delta\Omega$ is compact under weak topology, hence the feasible region is compact. By Weierstrass Theorem, there exists τ^* that attains the maximum.

Note that (P) is a convex optimization problem, there exists a solution τ^* which is an extreme point of $\mathcal{T}_{MD}(p)$. Theorem 2.1 of Winkler (1988) then implies the linear independence condition, and the size of the support of τ^* is bounded by the number of moment constraints plus one. Note that the vector-valued moment constraints in condition (ii) of Proposition 1 have exactly $2(n-1)$ independent scalar-valued moment constraints, so any $\tau \in \mathrm{ext}(\mathcal{T}_{MD}(p))$ has $|\mathrm{supp}(\tau)| \leqslant 2(n-1)+1 = 2n-1$. ■

C.2 Linear Persuasion

Proof of Proposition 9. We first show (i) and (ii) are equivalent. Suppose $q \in \Delta(X)$ is implementable under mediation, then there exists $\tau \in \mathcal{T}_{MD}(p)$ that induces q, that is, $q(S) = \mathbb{P}_\tau[\mathbb{E}_\mu(\omega) \in S]$ for all measurable $S \subseteq X$. We construct a dilation $T : X \to \Delta(X)$ by $T_x = \mathbb{E}_\tau[\mu|\mathbb{E}_\mu(\omega) = x]$. By construction we have $x = \int y\,\mathrm{d}T_x(y)$ for all x and $\int T_x\,\mathrm{d}q(x) = \int \mu\,\mathrm{d}\tau(\mu) = p$. Note that $\int T_x v_s(x)\,\mathrm{d}q(x) = \int V_S(\mu)\mu\,\mathrm{d}\tau = p\int V_S\,\mathrm{d}\tau = p\int v_s\,\mathrm{d}q$, where the first and third equalities are obtained by iterated expectation and $V_S(\mu) = v_s(\mathbb{E}_\mu(\omega))$, and the second by truth-telling. Hence, the dilation constructed satisfies $Tq = Tq^s = p$.

Conversely, suppose there exists a dilation T such that $Tq = Tq^s = p$. Then let $\tau \in \Delta(\Delta(\Omega))$ be the induced measure by q, that is, $\tau(R) = \mathbb{P}_q[T_x \in R]$ for all measurable $R \subseteq \Delta(\Omega)$. By change of variable, we obtain $\int \mu \,\mathrm{d}\tau = \int T_x \,\mathrm{d}q = p$ and $\int V_S(\mu)\mu \,\mathrm{d}\tau = \int v_s(x) T_x \,\mathrm{d}q = p \cdot \int v_s \,\mathrm{d}q = p \int V_S \,\mathrm{d}\tau$, which implies $\tau \in \mathcal{T}_{MD}(p)$.

The equivalence between (ii) and (iii) is straightforward. Note that given a dilation T that satisfies (ii), we may construct $\pi \in \Delta(\Omega \times X)$ by $\pi(\cdot|x) = T_x$ with $\mathrm{marg}_X \pi = q$. The definition of dilation and $Tq = p$ ensures $\mathbb{E}_\pi[\omega|x] = x$ and $\mathrm{marg}_\Omega \pi = p$. For any $g \in \mathbb{R}^\Omega$, $\int_{\Omega \times X} v_s(x) g(\omega) \,\mathrm{d}\pi = \int_X v_s(x) \left(\int_\Omega g(\omega) \,\mathrm{d}T_x(\omega)\right) \mathrm{d}q(x) = (\int g \,\mathrm{d}p)(\int v_s(x) \,\mathrm{d}q)$, where the first equality is by iterated expectation and the second is by $Tq^s = p$. For the converse, similar argument shows that we can construct a dilation T that satisfies (ii) by $T_x = \pi(\cdot|x)$ given any π that satisfies (iii). ■

Proof of Corollary 6. Suppose q is induced by some $\tau \in \mathcal{T}_{MD}(p)$, then Proposition 9 implies there exists a dilation T such that $Tq = Tq^s = p$. The Strassen's Theorem then implies $q \preceq_{cvx} p$ and $q^s \preceq_{cvx} p$. In particular, observe that $q^s \preceq_{cvx} p$ implies $\int x \,\mathrm{d}q^s(x) = \int x \,\mathrm{d}p(x) = \bar{x}_p$, hence $\int v_s(x)(x - \bar{x}_p) \,\mathrm{d}q(x) = 0$. ■

Proof of Corollary 7. Take any $q \in \Delta(X)$ that is implementable under mediation. Corollary 6 implies that $\int_X x \,\mathrm{d}q(x) = \bar{x}_p$ and $\int_X v_s(x)(x - \bar{x}_p) \,\mathrm{d}q(x) = 0$, where $\bar{x}_p$ is the prior mean. We refer to these two equations as relaxed Bayes-plausibility and relaxed truth-telling constraint. Note that $X = [a, b]$ is an interval in $\mathbb{R}$. Substituting the relaxed Bayes-plausibility to the relaxed truth-telling constraint, we obtain

$$\int_a^b v_s(x)(x - \bar{x}_p) \,\mathrm{d}q(x) = \int_a^b \int_a^b v_s(x)(x - y) \,\mathrm{d}q(y) \,\mathrm{d}q(x) = \int_a^b \int_a^b v_s(x)(x - y) \,\mathrm{d}q(x) \,\mathrm{d}q(y),$$

where the second equality is by changing the integration order. We then split the double integral into two parts and swap the variables x, y in the second part:

$$\begin{aligned}
&\int_a^b \int_a^b v_s(x)(x-y)\,\mathrm{d}q(x)\,\mathrm{d}q(y)\\
=&\int_a^b \int_y^b v_s(x)(x-y)\,\mathrm{d}q(x)\,\mathrm{d}q(y)+\int_a^b \int_a^y v_s(x)(x-y)\,\mathrm{d}q(x)\,\mathrm{d}q(y)\\
=&\int_a^b \int_y^b v_s(x)(x-y)\,\mathrm{d}q(x)\,\mathrm{d}q(y)+\int_a^b \int_a^x v_s(y)(y-x)\,\mathrm{d}q(y)\,\mathrm{d}q(x)\\
=&\int_a^b \int_y^b v_s(x)(x-y)\,\mathrm{d}q(x)\,\mathrm{d}q(y)+\int_a^b \int_y^b v_s(y)(y-x)\,\mathrm{d}q(x)\,\mathrm{d}q(y)\\
=&\int_a^b \int_y^b (v_s(x)-v_s(y))(x-y)\,\mathrm{d}q(x)\,\mathrm{d}q(y)\geqslant 0,
\end{aligned}$$

where the second last equality is obtained by changing the integration order, and the last equality by combining the two integrals into one. We obtain the last inequality by the weak monotonicity. Hence, for q to satisfy the relaxed truth-telling constraint, it must be the case v_s is constant on $\operatorname{supp}(q)$, which implies q is implementable under cheap talk. ■

C.3 Duality

Proof of Lemma 1. For every feasible τ in the primal and any (f, g) feasible in the dual, we have

$$\begin{aligned}
\langle f, p\rangle &= \int \langle f, \mu\rangle\,\mathrm{d}\tau\\
&\geqslant \int V_M(\mu)+\langle g, \mu-p\rangle V_S(\mu)\,\mathrm{d}\tau\\
&= \int V_M(\mu)\,\mathrm{d}\tau+\langle g, \int(\mu-p)V_S(\mu)\,\mathrm{d}\tau\rangle\\
&= \int V_M(\mu)\,\mathrm{d}\tau,
\end{aligned}$$

where the first equality is by the Bayes-plausibility constraint, the second inequality is by the dual constraint, the third equality is by linearity, and the fourth equality is by the truth-telling constraint. ■

Proof of Lemma 2. First, we rewrite the dual value as follows

$$\begin{aligned}\bar{V}(p) = \inf_{\langle f,\mu\rangle \geqslant V_M(\mu)+\langle g,\mu-p\rangle V_S(\mu)} \langle f,p\rangle &= \inf_{g\in\mathbb{R}^n} \inf_{\langle f,\mu\rangle \geqslant V_M(\mu)+\langle g,\mu-p\rangle V_S(\mu)} \langle f,p\rangle \\ &= \inf_{g\in\mathbb{R}^n} \sup_{\tau\in\mathcal{T}_{BP}(p)} \int V_M(\mu) + \langle g,\mu-p\rangle V_S(\mu)\,\mathrm{d}\tau,\end{aligned}$$

where $\mathcal{T}_{BP}(p)$ is the set of Bayes-plausible τ, and the second equality is by the no duality gap results in Bayesian persuasion, using $V_M(\mu) + \langle g,\mu-p\rangle V_S(\mu)$ as the objective function. Since $\mathcal{T}_{BP}(p)$ is compact and $H(\tau,g) := \int V_M(\mu) + \langle g,\mu-p\rangle V_S(\mu)\,\mathrm{d}\tau$ is linear(affine) and continuous in $\tau(g)$, we may apply Sion's minimax theorem to exchange the sup and inf.

After exchanging the sup and inf, we have

$$\bar{V}(p) = \sup_{\tau\in\mathcal{T}_{BP}(p)} \inf_{g\in\mathbb{R}^n} \int V_M(\mu) + \langle g,\mu-p\rangle V_S(\mu)\,\mathrm{d}\tau.$$

Note that for every $\tau \in \mathcal{T}_{BP}(p)$ that is not feasible in the primal of mediation problem, that is, $\int V_S(\mu)(\mu-p)\,\mathrm{d}\tau \neq 0$, we have $\inf_{g\in\mathbb{R}^n} \int V_M(\mu) + \langle g,\mu-p\rangle V_S(\mu)\,\mathrm{d}\tau = -\infty$. Hence we may restrict our consideration to the feasible region $\mathcal{T}_{MD}(p)$ of the primal (P), which leads to

$$\bar{V}(p) = \sup_{\tau\in\mathcal{T}_{MD}(p)} \inf_{g\in\mathbb{R}^n} \int V_M(\mu) + \langle g,\mu-p\rangle V_S(\mu)\,\mathrm{d}\tau = \sup_{\tau\in\mathcal{T}_{MD}(p)} \int V_M(\mu)\,\mathrm{d}\tau = \widehat{V}(p),$$

as desired. ■

Proof of Corollary 3. Note that in the proof of no duality gap, we observe that

$$\bar{V}(p) = \inf_{g\in\mathbb{R}^n} \sup_{\tau\in\mathcal{T}_{BP}(p)} \int V_M(\mu) + \langle g,\mu-p\rangle V_S(\mu)\,\mathrm{d}\tau.$$

Since adding a constant vector to g does not affect the value of $\langle g,\mu-p\rangle$, we may restrict our attention to $g \in \mathbb{R}^n$ such that $\langle g,p\rangle = 0$. We know from Kamenica and

Gentzkow (2011) the value of Bayesian persuasion equals to the value of the concave envelope of objective at the prior, so we may express the value of mediation as

$$\bar{V}(p) = \inf_{g \in \mathbb{R}^n : \langle g, p \rangle = 0} \operatorname{cav}\left[V_M(\mu) + \langle g, \mu \rangle V_S(\mu)\right](p).$$

■

C.4 Comparison of Communication Protocols

We first invoke Theorem 2.1 of Winkler (1988) and characterize the set of extreme points as follows:

Lemma 3 (Extreme points)**.** *The extreme points* $\operatorname{ext}(\mathcal{T}_{BP}(p))$ *and* $\operatorname{ext}(\mathcal{T}_{MD}(p))$ *can be expressed as*

$$\operatorname{ext}(\mathcal{T}_{BP}(p)) = \bigcup_{m=1}^{n} \Big\{ \tau \in \mathcal{T}_{BP}(p) : \tau = \sum_{i=1}^{m} \tau_i \delta_{\mu_i}, \sum \tau_i = 1, \tau_i > 0,$$
$$\text{the vectors } (\mu_i^1, \ldots, \mu_i^{n-1}, 1)^T_{i=1,\ldots,m} \text{ are linearly independent} \Big\},$$

and

$$\operatorname{ext}(\mathcal{T}_{MD}(p)) = \bigcup_{m=1}^{2n-1} \Big\{ \tau \in \mathcal{T}_{MD}(p) : \tau = \sum_{i=1}^{m} \tau_i \delta_{\mu_i}, \sum \tau_i = 1, \tau_i > 0,$$
$$\text{the vectors } (\mu_i^1, \ldots, \mu_i^{n-1}, V_S(\mu_i)(\mu_i^1 - p^1), \ldots, V_S(\mu_i)(\mu_i^{n-1} - p^{n-1}), 1)^T_{i=1,\ldots,m}$$
$$\text{are linearly independent} \Big\}.$$

Applying the preceding lemma, we further characterize the relationship between the nested feasible sets and their extreme points.

Lemma 4.

$$\operatorname{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{MD}(p) = \operatorname{ext}(\mathcal{T}_{MD}(p)) \bigcap \mathcal{T}_{CT}(p) = \operatorname{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{CT}(p).$$

Proof. We first show that $\operatorname{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{MD}(p) \subseteq \mathcal{T}_{CT}(p)$. Suppose $\tau \in \operatorname{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{MD}(p)$, assume to the contrary that $\tau \notin \mathcal{T}_{CT}(p)$. Then, at least one of $\alpha_i := \tau_i \left[(\sum_i \tau_i V_S(\mu_i)) - V_S(\mu_i)\right]$

is non-zero. Let $v_i = (\mu_i^1, \ldots, \mu_i^{n-1}, 1)^T$, we have $\sum_i \alpha_i v_i = 0$, which contradicts the fact that $\{v_i\}_i$ is linearly independent.

Next, we show that $\text{ext}(\mathcal{T}_{MD}(p)) \bigcap \mathcal{T}_{CT}(p) \subseteq \text{ext}(\mathcal{T}_{BP}(p))$. This follows from the fact that for all $V \in \mathbb{R}$, $(\mu_i^1, \ldots, \mu_i^{n-1}, V(\mu_i^1 - p^1), \ldots, V(\mu_i^{n-1} - p^{n-1}), 1)^T_{i=1,\ldots,m}$ is linearly independent if and only if $(\mu_i^1, \ldots, \mu_i^{n-1}, 1)^T_{i=1,\ldots,m}$ is linearly independent.[16]

Now we prove the set equations. It is clear that $\text{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{MD}(p) \subseteq \text{ext}(\mathcal{T}_{MD}(p))$, and hence $\text{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{MD}(p) \subseteq \text{ext}(\mathcal{T}_{MD}(p)) \bigcap \mathcal{T}_{CT}(p)$. The reverse inclusion is implied by $\text{ext}(\mathcal{T}_{MD}(p)) \bigcap \mathcal{T}_{CT}(p) \subseteq \text{ext}(\mathcal{T}_{BP}(p))$ and hence the first equation holds. For the second equation, it suffices to show $\text{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{MD}(p) = \text{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{CT}(p)$, which is implied by $\text{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{MD}(p) \subseteq \mathcal{T}_{CT}(p)$. ■

Next, we prove two lemmas that characterize the pairwise comparison of the values of communication, which specialize to Proposition 1 and Proposition 3 in the perfectly aligned case.

Lemma 5 (Persuasion and Mediation). *$\mathcal{V}^*_{BP}(V_M, p) = \mathcal{V}^*_{MD}(V_M, p)$ if and only if*

$$\mathcal{T}^*_{BP}(V_M, p) \bigcap \text{ext}(\mathcal{T}_{MD}(p)) \neq \varnothing.$$

Proof. It is clear that $\mathcal{V}^*_{BP}(V_M, p) = \mathcal{V}^*_{MD}(V_M, p)$ if and only if $\mathcal{T}^*_{BP}(V_M, p) \bigcap \mathcal{T}_{MD}(p) \neq \varnothing$, so the if direction of the claim is immediate. For the only if direction, suppose there exists $\tau \in \mathcal{T}_{MD}(p)$ that is optimal in the persuasion problem. As $\mathcal{T}_{MD}(p)$ is compact and convex, by Krein-Milman, τ can be expressed as the barycenter of a probability measure ν on $\mathcal{T}_{MD}(p)$, which is supported by $\bar{\text{co}}\,\text{ext}(\mathcal{T}_{MD}(p))$.[17] Assume to the contrary that $\mathcal{T}^*_{BP}(V_M, p)) \bigcap \text{ext}(\mathcal{T}_{MD}(p)) = \varnothing$. Then $\int V_M \, d\tau^* < \mathcal{V}^*_{BP}(V_M, p)$ for all $\tau^* \in \text{ext}(\mathcal{T}_{MD}(p))$, which leads to $\int V_M \, d\tau = \int_{\text{ext}(\mathcal{T}_{MD}(p))} \int_{\Delta(\Omega)} V_M(\mu) \, d\tau^*(\mu) \, d\nu(\tau^*) < \mathcal{V}^*_{BP}(V_M, p)$, contradiction. ■

Lemma 6 (Mediation and Cheap talk). *$\mathcal{V}^*_{MD}(V_M, p) = \mathcal{V}^*_{CT}(V_M, p)$ if and only if*

$$\mathcal{T}^*_{MD}(V_M, p) \bigcap \text{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{CT}(p) \neq \varnothing.$$

[16] To see this, row operations won't affect the rank of the matrices.

[17] We might need to be careful with the closure here. For a rigorous treatment, see Winkler (1988) Section 3.

Proof. It is immediate that $\mathcal{V}^*_{MD}(V_M, p) = \mathcal{V}^*_{CT}(V_M, p)$ is equivalent to $\mathcal{T}^*_{MD}(V_M, p) \cap \mathcal{T}_{CT}(p) \neq \varnothing$. This establishes the if part of the statement. The converse would follow if $\mathcal{T}^*_{MD}(p) \cap \mathcal{T}_{CT}(p) \neq \varnothing$ implies that $\mathcal{T}^*_{MD}(p) \cap \text{ext}(\mathcal{T}_{BP}(p)) \cap \mathcal{T}_{CT}(p) \neq \varnothing$. Fix $\tau \in \mathcal{T}^*_{MD}(p) \cap \mathcal{T}_{CT}(p)$. By Krein-Milman, we can represent τ as the barycenter of a probability measure ν on $\mathcal{T}_{MD}(p)$, which is supported by $\bar{\text{co}}\,\text{ext}(\mathcal{T}_{MD}(p))$. Since $\tau \in \mathcal{T}_{CT}(p)$, we must have $\tau^* \in \mathcal{T}_{CT}(p)$ for all $\tau^* \in \text{supp}(\nu)$. Moreover, $\int V_M \, \mathrm{d}\tau = \int V_M \, \mathrm{d}\tau^*$ for all $\tau^* \in \text{supp}(\nu)$ (otherwise contradict the fact that $\tau \in \mathcal{T}^*_{MD}(V_M, p)$). Therefore, $\tau^* \in \mathcal{T}^*_{MD}(V_M, p) \bigcap \text{ext}(\mathcal{T}_{MD}(p)) \bigcap \mathcal{T}_{CT}(p)$. By Lemma 4, we have $\tau^* \in \mathcal{T}^*_{MD}(V_M, p) \bigcap \text{ext}(\mathcal{T}_{BP}(p)) \bigcap \mathcal{T}_{CT}(p)$ and hence the desired result. ■

Proof of Theorem 1. We only prove the first equivalence; the second equivalence is immediate by Kamenica and Gentzkow (2011) and Lipnowski and Ravid (2020). The if direction is immediate. For the only if direction, by Lemma 5, there exists $\tau = \sum_i \tau_i \delta_{\mu_i} \in \mathcal{T}^*_{BP}(p) \bigcap \text{ext}(\mathcal{T}_{MD}(p))$. By Corollary 1 of Dworczak and Kolotilin (2022), there exists $f \in \mathbb{R}^n$ such that

$$\begin{aligned} V_S(\mu_i) &= \langle f, \mu_i \rangle \qquad \forall i \\ V_S(\mu) &\leqslant \langle f, \mu \rangle \qquad \forall \mu \in \Delta(\Omega). \end{aligned}$$

Since $\tau \in \text{ext}(\mathcal{T}_{MD}(p))$, we have

$$\sum_i \tau_i V_S(\mu_i)(\mu_i - p) = 0.$$

Substituting the equations from complementary slackness, for every $k = 1, \ldots, n$, we have

$$\sum_i \tau_i \mu_i^k \langle f, \mu_i \rangle = p^k \sum_i \tau_i \langle f, \mu_i \rangle = p^k \langle f, p \rangle,$$

where the last equality is by Bayes-plausibility. Taking dot product with f, by bilinearity, we have

$$\langle f, p \rangle^2 = \langle f, p\langle f, p \rangle \rangle = \langle f, \sum_i \tau_i \mu_i \langle f, \mu_i \rangle \rangle = \sum_i \tau_i \langle f, \mu_i \langle f, \mu_i \rangle \rangle = \sum_i \tau_i \langle f, \mu_i \rangle^2.$$

By Jensen's inequality and Bayes-plausibility,

$$\sum_i \tau_i \langle f, \mu_i \rangle^2 \geqslant \left(\sum_i \tau_i \langle f, \mu_i \rangle \right)^2 = \langle f, p \rangle^2,$$

and the equality holds if and only if $\langle f, \mu_i \rangle = \langle f, \mu_j \rangle$ for all i, j. Therefore, $V_S(\mu_i) = \langle f, \mu_i \rangle = \langle f, p \rangle$ for all i, which implies that $\tau \in \mathcal{T}_{CT}(p)$. ■

Proof of Proposition 2. We first illustrate the sufficiency of mono-crossing. Suppose $V_S(\mu) - \mathcal{V}^*_{CT}(p)$ is mono-crossing, we want to show for any $\tau \in \mathcal{T}_{MD}(p)$, $\mathbb{E}_\tau[V_S] \leqslant \mathcal{V}^*_{CT}(p)$. Let $\mathcal{V}^*$ denotes $\mathcal{V}^*_{CT}(p)$. As it is the optimal cheap talk value, by Lipnowski and Ravid (2020), $V_S(\mu) \leqslant \mathcal{V}^*$ on at least one of $[0, p]$ or $[p, 1]$. WLOG, we assume the former is true. Then there are two cases: either $V_S(\mu) \leqslant \mathcal{V}^*$ on $[0, 1]$, or $V_S(\mu) - \mathcal{V}^*$ is mono-crossing from below and there exists $\tilde{\mu} > p$ with $V_S(\tilde{\mu}) > \mathcal{V}^*$. In the first case, the claim holds trivially.

In the second case, we let $\bar{\mu} := \inf\{\mu : V_S(\mu) > \mathcal{V}^*\} \geqslant p$. Since $V_S(\mu) - \mathcal{V}^*$ is mono-crossing from below, the integrand $(V_S(\mu) - \mathcal{V}^*)(\mu - p)$ of the truth-telling constraint is non-negative on $[0, p] \cup [\bar{\mu}, 1]$ and non-positive on $(p, \bar{\mu})$. Take any $\tau \in \text{ext}(\mathcal{T}_{MD}(p))$. If $\text{supp}(\tau)$ does not contain μ with $(V_S(\mu) - \mathcal{V}^*)(\mu - p) < 0$, then the truth-telling constraint implies $(V_S(\mu) - \mathcal{V}^*)(\mu - p) = 0$ almost surely, which implies $\mathbb{E}_\tau[V_S] \leqslant \mathcal{V}^*$. Suppose there exists $\mu_1 \in \text{supp}(\tau) \cap (p, \bar{\mu})$. As $V_S(\mu_1) < \mathcal{V}^*$, there must exist $\mu_2 \geqslant \bar{\mu}$ with $V_S(\mu_2) > \mathcal{V}^*$ in the support for τ to have an expected payoff higher than $\mathcal{V}^*$. To satisfy the Bayes-plausibility, τ must contain $\mu_0 < p$ in the support. As $|\text{supp}(\tau)| \leqslant 3$, these are all the posteriors in the support and $\tau = \sum \tau_i \delta_{\mu_i}$ for some $\tau_i \in (0, 1)$. The truth-telling constraint $\sum \tau_i (V_S(\mu_i) - \mathcal{V}^*)(\mu_i - p) = 0$ then implies $\tau_1(V_S(\mu_1) - \mathcal{V}^*) + \tau_2(V_S(\mu_2) - \mathcal{V}^*) \leqslant 0$, so $\sum \tau_i V_S(\mu_i) \leqslant \mathcal{V}^*$.

We prove the converse direction under the additional assumption that $V_S(p) < \mathcal{V}^*$. If $V_S(\mu) - \mathcal{V}^*$ is not mono-crossing, we can construct a feasible mediation plan that attains a value strictly higher than $\mathcal{V}^*$. Since $\mathcal{V}^*$ is attained by cheap talk, there exists $\mu_0 < p < \mu_1$ such that $V_S(\mu_0) = V_S(\mu_1) = \mathcal{V}^*$. We take μ_0, μ_1 that is the closest to p,[18] by continuity, $V_S < \mathcal{V}^*$ in (μ_0, μ_1). Since $V_S(\mu) - \mathcal{V}^*$ is not mono-crossing, there

[18]Formally, the "closest" here means $\mu_1 = \inf\{\mu > p : V_S(\mu) = \mathcal{V}^*\}$ and μ_0 be a supremum. These two points are well-defined as $V_S(p) < \mathcal{V}^*$.

exists $\tilde{\mu}$ such that $V_S(\tilde{\mu}) > \mathcal{V}^*$. WLOG, we assume $\tilde{\mu} > \mu_1$, then for all $\mu < p$, $V_S(\mu) \leqslant \mathcal{V}^*$ (otherwise $\mathcal{V}^*$ is not the optimal cheap talk value). Finally, there must exist a pair (μ_2, μ_3) with $\mu_1 < \mu_2 < \mu_3$ such that $V_S(\mu_2) > \mathcal{V}^* > V_S(\mu_3)$. Suppose not, then for all $\mu_2 > \mu_1$ with $V_S(\mu_2) > \mathcal{V}^*$, we have $V_S(\mu_3) \geqslant \mathcal{V}^*$ for all $\mu_3 > \mu_2$, which contradicts the fact $V_S(\mu) - \mathcal{V}^*$ is not mono-crossing from below. Finally, direct calculation shows that $\tau = \tau_0\delta_{\mu_0} + \tau_2\delta_{\mu_2} + \tau_3\delta_{\mu_3}$ leads to expected value strictly higher than $\mathcal{V}^*$, where

$$\tau_0 = \frac{(V_2 - V_3)(\mu_2 - p)(\mu_3 - p)}{V_2(\mu_2 - p)(\mu_3 - \mu_0) + V_3(\mu_3 - p)(\mu_0 - \mu_2)}, \tau_2 = \frac{V_3(\mu_0 - p)(\mu_3 - p)}{V_2(\mu_2 - p)(\mu_3 - \mu_0) + V_3(\mu_3 - p)(\mu_0 - \mu_2)}, \tau_3 = 1 - \tau_0 - \tau_2,$$

with $V_2 = V_S(\mu_2)$ and $V_3 = V_S(\mu_3)$. ■

Proof of Proposition 3. Since $V_S(p) \leqslant \mathcal{V}^*_{CT}(p)$, we can consider the disjoint cases $V_S(p) = \mathcal{V}^*_{CT}(p)$ and $V_S(p) < \mathcal{V}^*_{CT}(p)$ separately. When $V_S(p) = \mathcal{V}^*_{CT}(p)$, then $\mathcal{V}^*_{MD}(p) = \mathcal{V}^*_{CT}(p)$ if and only if no disclosure is optimal under mediation, which is equivalent to condition (i). In particular, when the Slater condition holds, we may use our duality result to verify whether no disclosure is optimal, which is equivalent to saying there exists $g \in \mathbb{R}$ such that the distorted utility $(1 + g(\mu - p))V_S(\mu)$ is superdifferentiable at p. When $V_S(p) < \mathcal{V}^*_{CT}(p)$, condition (ii) is implied by Proposition 2. ■

Proof of Proposition 4. Condition (i) is straightforward when V_S is concave or affine. if V_S is convex, then either 0 or 1 attains its maximum value. WLOG, assume $V_S(0) \leqslant V_S(1)$, and by continuity there exists a largest $\tilde{p} \in [0, 1]$ such that $V_S(0) = V_S(\tilde{p})$. If $0 < p < \tilde{p}$, then $\mathcal{V}^*_{CT}(p) = V_S(0)$. Otherwise, $\mathcal{V}^*_{CT}(p) = V_S(p)$. In either case, $V_S(\mu) - \mathcal{V}^*_{CT}(p)$ is mono-crossing from below, and the conclusion follows from Proposition 2.

Condition (ii) can be shown using the same argument in the proof of Corollary 7. ■

Proof of Proposition 5. By Lemma 5, we need to characterize the condition when $\mathcal{T}^*_{BP}(V_M, p) \bigcap \text{ext}(\mathcal{T}_{MD}(p)) \neq \varnothing$. Note that $\text{ext}(\mathcal{T}_{MD}(p))$ only contains τ with $|\text{supp}(\tau)| \leqslant 3$. By complementary slackness, $\tau = \sum_i \tau_i \delta_{\mu_i} \in \mathcal{T}^*_{BP}(V_M, p) \bigcap \text{ext}(\mathcal{T}_{MD}(p))$ if and only if there exists $f_1, f_0 \in \mathbb{R}$ such that

$$\begin{cases} V_M(\mu_i) = f_1\mu_i + f_0 & \forall i \\ V_M(\mu) \leqslant f_1\mu + f_0 & \forall \mu \in [0,1] \\ \sum_i \tau_i\mu_i = p & \\ \sum_i \tau_i V_S(\mu_i)(\mu_i - p) = 0 & \\ (\mu_i, V_S(\mu_i)(\mu_i - p), 1)_i \text{ is linearly independent.} & \end{cases}$$

If $|\text{supp}(\tau)| = 1$ then it must be $\tau = \delta_p$. If $|\text{supp}(\tau)| = 2$, the complementary slackness requires the line connecting $(\mu_1, V_M(\mu_1))$ and $(\mu_2, V_M(\mu_2))$ lies weakly above the graph of $V_M(\mu)$. And the truth-telling constraint requires $V_S(\mu_1) = V_S(\mu_2)$. The linear independence condition is automatic if $\mu_1 \neq \mu_2$.

If the support of τ contains three distinct points $\mu_1 < \mu_2 < \mu_3$ with $\mu_2 = \alpha\mu_1 + (1-\alpha)\mu_3$. Again, the complementary slackness requires the line connecting $(\mu_1, V_M(\mu_1))$ and $(\mu_3, V_M(\mu_3))$ lies weakly above the graph of $V_M(\mu)$. In addition, this line coincides with V_M at μ_2. The linear independence condition is equivalent to

$$\alpha V_S(\mu_1)(\mu_1 - p) + (1-\alpha)V_S(\mu_3)(\mu_3 - p) \neq V_S(\mu_2)(\mu_2 - p).$$

Clearly, any $\tau \in \mathcal{T}_{CT}(p)$ does not satisfy this condition. Moreover, combining the linear independence condition with the Bayes plausibility and truth-telling constraint, it must be the case $\mu_2 \neq p$. Assume to the contrary that $\mu_2 = p$, then Bayes plausibility implies that $\tau_1 = \alpha(1-\tau_2)$ and $\tau_3 = (1-\alpha)(1-\tau_2)$, and

$$\begin{aligned} \sum_i \tau_i V_S(\mu_i)(\mu_i - p) &= (1-\tau_2)\left(\alpha V_S(\mu_1)(\mu_1 - p) + (1-\alpha)V_S(\mu_3)(\mu_3 - p)\right) \\ &\neq (1-\tau_2)V_S(\mu_2)(\mu_2 - p) = 0, \end{aligned}$$

contradicts the truth-telling constraint.

Finally, suppose $\mathcal{V}^*_{BP}(p) = \mathcal{V}^*_{MD}(p)$, while (i) and (ii) do not hold, then any $\tau \in \mathcal{T}_{CT}(p)$ is not in $\mathcal{T}^*_{BP}(V_M, p)$. Note that for any $\tau \in \mathcal{T}_{CT}(p)$ that is optimal under Bayesian persuasion and mediation, we may find a $\tau' \in \mathcal{T}_{CT}(p)$ with $|\text{supp}(\tau')| \leqslant 2$

that has the same value by Lemma 6. Therefore, if (i) and (ii) do not hold, no cheap talk equilibrium is optimal under Bayesian persuasion. ■

Proof of Proposition 6. By Lemma 6, it suffices to check if there exists $\tau \in \mathcal{T}_{CT}(p)$ with $|\operatorname{supp}(\tau)| \leqslant 2$ attains $\mathcal{V}^*_{MD}(p)$. By Corollary 4, such a $\tau \in \mathcal{T}^*_{MD}(p)$ if and only if there exists $(f_1, f_0, g) \in \mathbb{R}^3$ such that

$$\begin{aligned} V_M(\mu_i) + Vg(\mu_i - p) &= f_1\mu_i + f_0 \qquad i = 1, 2 \\ V_M(\mu) + V_S(\mu)g(\mu - p) &\leqslant f_1\mu + f_0 \qquad \forall \mu \in [0, 1], \end{aligned}$$

where $V = V_S(\mu_i)$. By Bayes-plausibility, we have

$$f_1 p + f_0 = \sum_i \tau_i(f_1\mu_i + f_0) = \sum_i \tau_i(V_M(\mu_i) + Vg(\mu_i - p)) = \sum \tau_i V_M(\mu_i).$$

As $|\operatorname{supp}(\tau)| = 2$, $\mu_i \neq p$ for $i = 1, 2$. Substituting the preceding equation into the complementary slackness condition at $i = 1, 2$, we have

$$(1 - \tau_i)(V_M(\mu_i) - V_M(\mu_{-i})) = (f_1 - Vg)(\mu_i - p),$$

which implies that

$$g = \frac{1}{V}\left(f_1 - \frac{(1 - \tau_i)(V_M(\mu_i) - V_M(\mu_{-i}))}{\mu_i - p}\right)$$

for $i = 1, 2$. By Bayes-plausibility the two equations are linearly dependent, so g is pinned down by f_1 and substitute τ into any of these two equations. Substitute the (g, f_1, f_0) into the inequality constraints, we obtain the desired condition in the statement. ■