The Bounds of Mediated Communication

Roberto Corrao* Yifan Dai†

MIT MIT

July 2024

Abstract

We study the bounds of mediated communication in sender-receiver games in which the sender's payoff is state-independent. We show that the feasible distributions of beliefs under mediation are those that induce zero correlation, but not necessarily independence, between the sender's payoff and the receiver's belief. Mediation attains the upper bound on the sender's value, i.e., the Bayesian persuasion value, if and only if this value collapses to the lower bound, i.e., the cheap talk value. When the sender's value function is strictly quasiconvex and a full-dimensionality condition holds, mediation lies strictly below and above these two bounds. More generally, mediation is strictly valuable when the sender has *countervailing incentives* in the space of the receiver's belief. We apply our results to asymmetric-information settings such as bilateral trade and lobbying and explicitly construct mediation policies such that the informed and uninformed parties are better off than under unmediated communication.

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

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

We are particularly indebted to Drew Fudenberg, Stephen Morris, and Alexander Wolitzky for many suggestions that substantially improved the paper. We are grateful to Charles Angelucci, Ian Ball, Alessandro Bonatti, Simone Cerreia-Vioglio, Laura Doval, Bob Gibbons, Andrew Koh, Marina Halac, Nicole Immorlica, Haruki Kono, Giacomo Lanzani, Xiao Lin, Elliot Lipnowski, Joel Flynn, Teddy Mekonnen, Doron Ravid, Sivakorn Sanguanmoo, Alex Smolin, Juuso Toikka, Giacomo Weber, and Leeat Yariv for helpful comments and conversations. Roberto Corrao gratefully acknowledges the Gordon B. Pye Dissertation Fellowship for financial support. An earlier version of this paper was circulated under the title "Mediated Communication with Transparent Motives."

1 Introduction

Consider a receiver who faces a decision problem under uncertainty and a sender who is privately informed about the state. The sender can communicate with the receiver before they take an action, and their final payoff only depends on the receiver's action, that is, the sender has *transparent motives* in the sense of Chakraborty and Harbaugh (2010) and Lipnowski and Ravid (2020). These situations are pervasive in economics: a seller has superior information about the quality of a good and always wants to maximize the probability of selling it to buyers. One of two extreme assumptions is usually considered: 1) The sender can commit ex-ante to any information policy, such as an experiment that conveys verifiable information to the receiver, or 2) The sender cannot commit to any experiment, their private information is not verifiable (i.e., it is soft), but they can freely send messages to the receiver. The first case has been extensively analyzed in recent years and corresponds to the *Bayesian persuasion* model of Kamenica and Gentzkow (2011). The second case corresponds to a game of strategic information transmission or *cheap talk* as introduced in Crawford and Sobel (1982). It is well known that with commitment, the sender can often achieve a strictly higher payoff than the one obtained by conveying no information. Perhaps more surprisingly, Chakraborty and Harbaugh (2010) and Lipnowski and Ravid (2020) showed that the sender can also achieve a strictly higher payoff under cheap talk than without communication.

This paper revisits the intermediate case of *mediated communication* introduced in Myerson (1982). We expand the set of players by considering a third-party mediator who cannot take the relevant decision in place of the receiver and is uninformed about the state; hence they must resort to information willingly shared by the sender. However, the mediator can commit to any *communication mechanism* that collects reports from the sender and sends messages to the receiver. In the buyer-seller example above, the mediator can represent an advertising agency or a financial intermediary with a prominent reputation that collects reports from the seller and conveys credible information to the buyers.

We focus on the case where the mediator's preference is aligned with the sender's, so they act to maximize the sender's payoff. Clearly, the sender-optimal values across the three protocols considered are weakly ordered because the space of feasible information policies becomes smaller from persuasion to mediation and from mediation to cheap talk: $BP \geq MD \geq CT$.[1] We decompose the gap between Bayesian persuasion and cheap talk as follows:

$$\underbrace{BP - CT}_{\text{Value of Commitment}} = \underbrace{BP - MD}_{\text{Value of Elicitation}} + \underbrace{MD - CT}_{\text{Value of Mediation}}.$$

[1] Here, BP, MD, and CT respectively denote the sender-optimal values attained under Bayesian persuasion, mediation, and cheap talk.

The gap $BP - CT$ represents the *value of commitment* for the sender. The first component of this gap is $BP - MD$ which captures the *value of elicitation*. In both persuasion and mediation, there is an entity with commitment power, the sender and the mediator, respectively. However, the mediator is not directly informed about the state and has to *elicit* this information in an incentive-compatible way. Differently, the gap $MD - CT$ captures the *value of mediation* because it corresponds to the additional value that an uninformed third party with commitment can secure to the sender when the latter has no commitment power. Our results provide sufficient and necessary conditions such that the values of elicitation and mediation are strictly positive.

Outline of the results By the revelation principle, the mediator acts "as-if" selecting a *communication equilibrium* outcome of the sender-receiver game. However, differently from Myerson (1982) but similarly to the recent literature on Bayesian persuasion and cheap talk, we adopt a *belief-based approach* to mediation. We show that the feasible distributions of receiver's beliefs are those that induce zero correlation, but not necessarily independence, between the sender's payoff and the receiver's belief. This condition translates the truth-telling constraint of the sender from the space of mechanisms to the space of beliefs. We can then represent the optimal mediation problem as a linear program under moment constraints in the belief space: the standard Bayes plausibility constraint and the zero-correlation constraint.

Exploiting this rewriting of the mediation problem, we show that the sender can attain the optimal persuasion payoff under mediation if and only if this value can be attained under cheap talk. Therefore, we show that when elicitation is valueless, so is commitment. Given that the value of commitment is often strictly positive, this implies that an uninformed mediator cannot usually guarantee the same value that the sender would achieve if they could commit in the first place.

Next, we introduce a key concept for cheap talk: the *full-dimensionality condition*. This condition holds when the cheap talk value at the prior can still be attained when the prior probability of an arbitrary state is slightly decreased. For example, it is satisfied for almost every prior when the receiver's action set is finite and, at every binary prior such that the babbling equilibrium is not sender optimal.

Under the full-dimensionality condition, we characterize the cases where elicitation and mediation are strictly valuable, that is, $BP > MD$ and $MD > CT$, respectively. We prove these results by first providing distinct sufficient and necessary conditions for the values of elicitation and mediation to be strictly positive without any additional assumption and then show that under full dimensionality these conditions are the same. Elicitation is strictly valuable if and only if there exists a belief $\mu \in \Delta(\Omega)$ of the receiver such that the maximum cheap talk value at μ is strictly higher than the maximum cheap talk value at the prior p.[2] Mediation is strictly valuable if and only if there exist two beliefs $\mu_+, \mu_- \in \Delta(\Omega)$ of the receiver that are colinear with the prior p and such that the maximum cheap talk value at p lies strictly between the maximum cheap talk value at μ_+ and the minimum cheap talk value at μ_-. In particular, we construct an improving mediation plan by randomizing over distributions of beliefs that include cheap talk equilibria at μ_+ and μ_- respectively. In this case, the optimal mediation plan must be random.

All the aforementioned conditions admit geometric characterizations in terms of the quasiconcave and quasiconvex envelopes of the sender's value function. We use these to show that when the sender's value function is strictly quasiconvex and the full-dimensionality condition holds, then either $BP > MD > CT$ or all three communication protocols attain the same value equal to the global maximum of the sender's payoff.

In general, we find that mediation has a strictly positive value when the sender has *countervailing incentives* in the space of the receiver's beliefs, that is, when the sender would like to induce more optimistic beliefs for some realized messages and more pessimistic beliefs for some others. More formally, this translates to the failure of a weak form of single-crossing. For multidimensional environments with strictly quasiconvex utility for the sender, countervailing incentives are captured by the non-monotonicity of the restriction of the sender's utility to the edges of the simplex.

We revisit the think tank example in Lipnowski and Ravid (2020) by assuming that the think tank acts as a mediator between an interest group (the sender) and the lawmaker (the receiver). Countervailing incentives arise because the interest group strictly prefers the lawmaker to approve one of several new policies as opposed to retaining the status quo. Similarly, we apply our results to study advertising agencies or financial intermediaries that operate as mediators between sellers and buyers. Countervailing incentives can arise because of reputation concerns of the seller or because of non-monotone preferences over risky prospects (e.g., mean-variance) of the receiver. For these examples, both elicitation and mediation are usually strictly valuable, thereby rationalizing the ubiquitous presence of intermediaries in these markets.

Next, we analyze a class of sender-receiver games that we call *acceptance games* where the receiver chooses whether to accept or not a risky prospect whose value depends only on the posterior expectations of a finite number of functions of the state and the value of their outside option is private information of the receiver. The sender's payoff coincides with the probability of acceptance and, under standard assumptions, it is quasiconvex and consistent

[2] Here, Ω denotes the finite state space and $p \in \Delta(\Omega)$ denotes the common prior.

with countervailing incentives, which allows us to conclude that mediation is strictly valuable. In addition, when the distribution of the outside option is log-concave, the receiver is also strictly better off under sender-preferred mediation than under any cheap-talk equilibrium.

Finally, we discuss some additional implications of our results. For long cheap talk (see Aumann and Hart (2003)) and repeated games with asymmetric information (see Hart (1985)), our results characterize the environments where the sender's payoff under the best correlated equilibrium is strictly higher than that under the best Nash equilibrium.

1.1 Illustrative Example

We illustrate the geometric comparison of Bayesian persuasion, mediation, and cheap talk with a simple advertising model that compares the case where a seller directly communicates with a buyer to the case where the seller hires an advertising agency to mediate communication.

A seller plans to commercialize a new product with quality $\omega \in \Omega = \{0, 1\}$ that they privately know. The buyer has prior belief $p \in (0, 0.55)$ that the quality is good ($\omega = 1$). Consider first the case when the seller can only communicate by cheap talk messages. For every message, the buyer updates their belief $\mu \in [0, 1]$ and decides whether to purchase the good or take their outside option with quality $\varepsilon \in [0, 1]$. The buyer is privately informed about the outside option, but the seller knows only that the distribution of ε is G. We assume that G has an unimodal density g.[3] Moreover, we assume that the buyer will eventually observe the quality of the product regardless of their decision.

The market is competitive, and we normalize the price of the good and the outside option to 1. Thus, given posterior μ, the buyer purchases the good if and only if $\varepsilon \leq \mu$, for a total mass of $G(\mu)$. The seller's utility depends on the total demand for the good and on reputation concerns:

$$\tilde{V}(\mu, \omega) = (1 - \delta)G(\mu) + \delta(\omega - \mu),$$

where $\delta > 0$. Specifically, when the realized quality exceeds the expectation, that is $\omega > \mu$, there is a positive effect of a surprisingly good product on the seller's future payoff. Conversely, when $\omega < \mu$, there is a negative reputation effect due to an unexpectedly bad product. Here, δ measures the impact of these reputation concerns on the seller's payoff.

As the state ω is privately known and the seller's payoff function is additively separable in $\tilde{V}(\mu, \omega)$, the seller acts to maximize $V(\mu) = (1 - \delta)G(\mu) - \delta\mu$. Under our assumptions

[3] This means that G is strictly convex up to some point $\hat{\varepsilon}$ and concave beyond that point, that is, G is S-shaped. Several recent papers in the persuasion literature focus on a similar class of indirect utility functions called S-shaped functions (Kolotilin, 2018; Kolotilin et al., 2022; Arieli et al., 2023).

on G, this indirect utility V is a rotated S-shaped function as illustrated in Figure 1.[4] An intermediate level of reputation concern induces countervailing incentives for the sender. For example, in Figure 1, for posteriors μ just before 3/4, the sender would like the buyer to be more optimistic about the product quality, whereas, for posteriors above 3/4, the seller would like the buyer to be more pessimistic.

By Lipnowski and Ravid (2020), the seller-optimal cheap talk value at any prior is given by the quasiconcave envelope of V at that prior, which is the dotted red line in Figure 1 (a). For example, the seller-preferred cheap talk equilibrium at p randomizes between $\mu = 0.55$ and $\mu = 0$ with probability $p/0.55$, and yields expected payoff of 0.

Figure 1: Illustrative Example

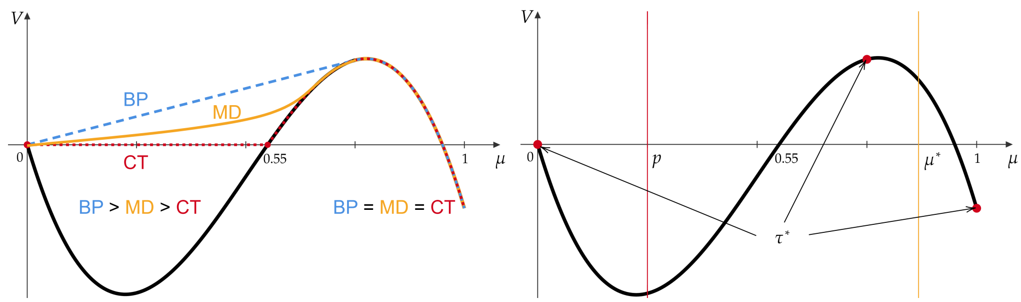

(a) Comparison of Bayesian persuasion, mediation, and cheap talk

(b) Construction of strictly improving mediation plan

The colored lines in (a) represent the seller's optimal payoff from Bayesian persuasion (blue dashed), mediation (yellow solid), and cheap talk (red dashed). The discussion here focuses on the case $p \in (0, 0.55)$, where the three lines do not coincide.

Next, we show that the seller 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 the buyer. The advertiser does not have the expertise to assess the exact quality of the good and can only convey information the seller reports. The contract between the seller and the advertiser is fixed and binds the seller to pay the advertiser a fixed fraction of its revenue, so the advertiser maximizes the seller's expected payoff.[5] For example, the advertiser can strictly increase the seller's expected payoff by introducing randomness to the message distribution conditional on the seller's quality report. This randomness conditional on the seller's quality reports can be interpreted as the use of inessential visual effects or vague language in the advertising campaign.

[4]Specifically, Figure 1 plots the indirect utility V induced by the Beta(2,2) distribution and a weight $\delta = \frac{209}{409}$.

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

We construct a distribution of beliefs that is feasible for the advertiser and that yields a strict improvement for the seller with respect to direct communication. First fix $\xi \in (0,1)$ such that $\xi \cdot V(3/4) \cdot (3/4 - p) + (1-\xi) \cdot V(1) \cdot (1-p) = 0$.[6] With this, fix the belief $\mu^* = \xi \cdot 3/4 + (1-\xi) \cdot 1$, highlighted by the yellow line in Figure 1 (b), and observe that there exists $\alpha > 1$ such that $\alpha p + (1-\alpha)\mu^* = 0$. Now, consider the distribution of beliefs

$$\tau^* = \{(0; 1/\alpha), (3/4; (\alpha-1)\xi/\alpha), (1; (\alpha-1)(1-\xi)/\alpha)\}.$$

The three points in the support of this distribution are highlighted by the red dots in Figure 1 (b). This distribution does not correspond to a cheap talk equilibrium, as the seller would always have the incentive to induce $\mu = 3/4$ at every state. However, τ^* averages to p and induces zero correlation between the buyers' beliefs μ and the seller's payoff $V(\mu)$. In Theorem 1 below, we show that this is necessary and sufficient for τ^* to be implementable under mediation. Finally, one can verify that the seller's expected payoff under this distribution of beliefs is

$$\tfrac{\alpha-1}{\alpha}(\xi \cdot V(3/4) + (1-\xi) \cdot V(1)) > 0$$

yielding a strict improvement. This shows that, with a small enough commission rate, the seller strictly benefits from hiring an advertiser to mediate communication.[7]

The buyer is strictly better off under the mediation plan we constructed than under the sender-optimal cheap talk equilibrium. Note that the buyer's indirect utility $V_R(\mu) = \mu G(\mu) + \int_\mu^1 \varepsilon \, dG(\varepsilon)$ is strictly convex, and the induced distributions of posteriors are supported on $\{0, 3/4, 1\}$ under mediation and $\{0, 0.55\}$ under cheap talk. Hence, the distribution of beliefs under mediation is a mean-preserving spread of that under cheap talk, which leads to a strictly higher buyer payoff.

1.2 Literature review

Our work uses the "belief-based approach," a widely adopted methodology in Bayesian persuasion (Kamenica and Gentzkow (2011)) and cheap talk (Lipnowski and Ravid (2020)), to study mediated communication (Myerson (1982) and Forges (1986)).[8] Recent works on this topic study the comparison between mediation and other specific forms of communication in the uniform-quadratic case of Crawford and Sobel (1982). Blume et al. (2007) focuses on contrasting noisy cheap talk with cheap talk, while Goltsman et al. (2009) compares

[6] This coefficient exists because $V(1) < 0 < V(3/4)$.

[7] In Section 5, we characterize when similar constructions that randomize among posteriors with values strictly above/below the cheap talk value lead to a strictly higher payoff than cheap talk.

[8] Aumann and Maschler (1995) and Aumann and Hart (2003) first adopted the belief-based approach to respectively study zero-sum repeated games with asymmetric information and long cheap talk.

mediation, (long) cheap talk, and delegation. Differently, we characterize the comparison between persuasion, mediation, and cheap talk under state-independent preferences for the sender, but without additional parametric assumptions.

The most related paper in the mediation literature is Salamanca (2021), where mediated communication for *finite* games is analyzed using a recommendation approach similar to the original one in Myerson (1982). Our analysis differs from the one in Salamanca (2021) for several reasons. First, the two models are not nested since we focus on the transparent-motive case but we allow an arbitrary action space for the receiver. Second, our analysis is entirely carried out with a belief-based approach as opposed to the recommendation approach they use, which allows us to readily derive the same "virtual-utility" representation of the sender-optimal value of mediation and to more directly compare mediated communication with persuasion and cheap talk. Third, our results are different: While Salamanca (2021) focuses on deriving strong duality for the recommendation-based mediation problem, we use a more direct perturbation approach that allows us to completely characterize when elicitation and mediation are valuable for finite games at almost all prior beliefs.[9] Moreover, we provide several sufficient conditions such that our characterization extends to infinite-action games.

Some work in the mediation literature allows for transfers between the informed party and the intermediary. This considerably expands the set of implementable outcomes. For example, Corrao (2023) considers an optimal mediation problem with transfers where the mediator maximizes their revenue from payments from the informed party. He shows that, with binary state, every distribution of the receiver's beliefs is implementable. This is in sharp contrast with the zero-correlation restriction imposed by the truthtelling constraint in our setting.

Finally, our work is related to recent papers studying Bayesian persuasion with limited commitment or additional constraints (Lin and Liu, 2024; Lipnowski et al., 2022; Koessler and Skreta, 2023; Doval and Skreta, 2024). Like mediation, the communication protocols studied in these works can be seen as intermediate cases between Bayesian persuasion and (single-round) cheap talk. The transparent-motive assumption sometimes makes these intermediate cases attain one of the two bounds given by persuasion and cheap talk. For example, the credible information structures in Lin and Liu (2024) are the same ones that are feasible under persuasion, when the sender has transparent motives. Under the same assumption, Lipnowski and Ravid (2020) show that the sender's optimal payoff in the long cheap talk model of Aumann and Hart (2003) is the same as the one of single-round cheap talk. Instead, we show that the optimal sender's value under mediation can be strictly between the two

[9] Salamanca (2021) provides a binary-state example under transparent motives where the strict inequalities $BP > MD > CT$ hold, but does not characterize when these inequalities are strict.

bounds and we characterize when this is the case in several settings.

2 The Model

Consider three players: a sender, a receiver, and a mediator. Let Ω be a finite state space with $|\Omega| = n$.[10] The state $\omega \in \Omega$ is drawn according to a full-support common prior $p \in \Delta(\Omega)$, and the realization of ω is the sender's private information.[11] The receiver does not know ω and takes a payoff-relevant action $a \in A$, where A is a compact metric space. We assume the sender has a *state-independent* utility function $u_S : A \to \mathbb{R}$, that is, they have *transparent motives*, and the receiver has utility $u_R : \Omega \times A \to \mathbb{R}$. Both utility functions are continuous.

The sender and receiver communicate through the mediator, who commits to a communication mechanism $\sigma : R \to \Delta(M)$ without knowing ω, where R is the reporting space for the sender and M is the space of messages for the receiver. After observing ω, the sender sends a report $r \in R$ to the mediator. Given the report, the mediator draws a random message $m \in M$ according to σ and sends it to the receiver, who then takes an action $a \in A$. We consider the communication game induced by σ and focus on its Bayes-Nash equilibria, also known as the *communication equilibria* (see Myerson (1982) and Forges (1986)).[12] We assume that the mediator is perfectly aligned with the sender and selects a mechanism and an equilibrium to maximize the sender's expected utility.

Any mechanism σ and a communication equilibrium induce an outcome distribution $\pi \in \Delta(\Omega \times A)$. Applying the Revelation Principle (Myerson, 1982; Forges, 1986), it is without loss to consider outcome distributions induced by direct incentive-compatible mechanisms, that is, a communication equilibrium where the mediator asks the sender for a state report in $R = \Omega$, provides an action recommendation in $M = A$ to the receiver, and the sender truthfully reports the state while the receiver follows the action recommendation.

Fact. Any outcome distribution $\pi \in \Delta(\Omega \times A)$ is induced by some communication equilibrium if and only if it satisfies:

(i) Consistency: $\operatorname{marg}_\Omega \pi = p$

(ii) Obedience: For π-almost all $a \in A$, $\mathbb{E}_{\pi^a}[u_R(\omega, a)] = \max_{a' \in A} \mathbb{E}_{\pi^a}[u_R(\omega, a')]$, where $\pi^a \in \Delta(\Omega)$ is a version of the conditional probability given $a \in A$;

[10] Appendix J extends Theorem 1, Theorem 2, Point 1 of Theorem 3, and Theorem 4 to the case where Ω is an arbitrary compact metric space.

[11] We identify $\Delta(\Omega)$ with the standard $(n-1)$–simplex in $\mathbb{R}^n$ and we endow it with the relative topology induced by the Euclidean topology. For any set $D \subseteq \Delta(\Omega)$, let $\operatorname{int} D$ denote its interior.

[12] Formally, the sender's strategy is $\rho : \Omega \to \Delta(R)$ and the receiver's strategy is $\alpha : M \to \Delta(A)$. (ρ, α) forms an equilibrium if and only if $\mathbb{E}_p[\mathbb{E}_\sigma[u_S(\alpha(m))|\rho(\omega)]] \geq \mathbb{E}_p[\mathbb{E}_\sigma[u_S(\alpha(m))|\tilde{\rho}(\omega)]]$ and $\mathbb{E}_p[\mathbb{E}_\sigma[u_R(\omega, \alpha(m))|\rho(\omega)]] \geq \mathbb{E}_p[\mathbb{E}_\sigma[u_R(\omega, \tilde{\alpha}(m))|\rho(\omega)]]$ for any $\tilde{\rho}, \tilde{\alpha}$.

(iii) Honesty: For all $\omega, \omega' \in \Omega$, $\mathbb{E}_{\pi^\omega}[u_S(a)] \geq \mathbb{E}_{\pi^{\omega'}}[u_S(a)]$, where $\pi^\omega \in \Delta(A)$ is the conditional probability given $\omega \in \Omega$.

We say that $\pi \in \Delta(\Omega \times A)$ is a communication equilibrium (CE) outcome if it satisfies (i), (ii), and (iii).

3 Belief-based Approach to Mediated Communication

Instead of focusing on CE outcomes, we consider distributions over the receiver's posteriors $\tau \in \Delta(\Delta(\Omega))$ and the sender's indirect utility $V : \Delta(\Omega) \to \mathbb{R}$ in terms of the receiver's posterior. Define the indirect value correspondence $\mathbf{V} : \Delta(\Omega) \rightrightarrows \mathbb{R}$ by

$$\mathbf{V}(\mu) := \operatorname{co}\left(u_S\left(\operatorname*{argmax}_{a \in A} \mathbb{E}_\mu[u_R(\omega, a)]\right)\right).$$

For every posterior $\mu \in \Delta(\Omega)$, the set $\mathbf{V}(\mu)$ collects all the possible (expected) sender's payoffs that can be attained by some (potentially mixed) receiver's best response at posterior μ. By Berge's Theorem, $\mathbf{V}$ is upper hemi-continuous, compact, convex, and non-empty valued, that is, it is a *Kakutani correspondence*.[13] Define the functions $\overline{V}(\mu) = \max \mathbf{V}(\mu)$ and $\underline{V}(\mu) = \min \mathbf{V}(\mu)$, which are respectively upper and lower semi-continuous.[14]

Any CE outcome π induces a distribution over posterior beliefs $\tau^\pi \in \Delta(\Delta(\Omega))$ as follows: $\tau^\pi(D) = \int \mathbb{I}[\pi^a \in D] \, \mathrm{d}\pi$ for all Borel $D \subseteq \Delta(\Omega)$. It also induces an indirect utility for the sender $V^\pi : \Delta(\Omega) \to \mathbb{R}$ defined for τ^π-almost all posterior beliefs by

$$V^\pi(\mu) := \int u_S(a) \, \mathrm{d}\pi(a \mid \pi^a = \mu),$$

where $\pi(\cdot \mid \pi^a = \mu)$ is the conditional probability over $\Omega \times A$ given that $\pi^a = \mu$.

Definition 1. A distribution of posteriors $\tau \in \Delta(\Delta(\Omega))$ and a measurable function $V : \Delta(\Omega) \to \mathbb{R}$ are *induced by* some CE outcome $\pi \in \Delta(\Omega \times A)$ if $\tau = \tau^\pi$ and $V(\mu) = V^\pi(\mu)$ for τ–almost all μ.

For our main analysis, we focus on pairs (τ, V) that are induced by some CE outcome. For any $\tau \in \Delta(\Delta(\Omega))$, we say τ attains value $s \in \mathbb{R}$ if there exists $V \in \mathbf{V}$ such that $\int V \, \mathrm{d}\tau = s$.

[13] Observe that we can alternatively start our analysis from a primitive Kakutani correspondence $\mathbf{V}$ capturing the set of the sender's continuation values given each posterior. This setting is strictly more general than the one presented in the main text as, for example, it would allow us to capture non-EU preferences for the receiver (see Example 3).

[14] See Lemma 17.30 in Aliprantis and Border (2006).

Our first result characterizes the set of implementable distributions over posteriors and indirect utility functions using three conditions parallel to Consistency, Obedience, and Honesty. In particular, as the sender's preference is state-independent, their expected payoff should be the same conditional on every state report.

Theorem 1. *If a distribution of receiver's beliefs $\tau \in \Delta(\Delta(\Omega))$ and a measurable sender's indirect utility function $V : \Delta(\Omega) \to \mathbb{R}$ are induced by some CE outcome, then they satisfy*

(i) Consistency:*

$$\int \mu \, \mathrm{d}\tau(\mu) = p; \tag{BP}$$

(ii) Obedience: For τ-almost all $\mu \in \Delta(\Omega)$, $V(\mu) \in \mathbf{V}(\mu)$;*

(iii) Honesty:*

$$\mathrm{Cov}_\tau[V(\mu), \mu] = \mathbf{0}. \tag{zeroCov}$$

Conversely, if (τ, V) satisfy (i),(ii), and (iii), then there exists a CE outcome $\pi \in \Delta(\Omega \times A)$ such that $\mathbb{E}_\tau[V] = \mathbb{E}_\pi[u_S]$.[15]

The set of implementable distributions over posteriors under mediation is

$$\mathcal{T}_{MD}(p) := \{\tau \in \Delta(\Delta(\Omega)) : \exists\, V \in \mathbf{V} \text{ such that (BP) and (zeroCov) hold}\}.$$

We now sketch the derivation of equation zeroCov. For simplicity, consider the singleton-valued case: $\mathbf{V}(\mu) = V(\mu)$. Under transparent motives, the Honesty constraint implies that $\mathbb{E}_{\tau^\omega}[V(\mu)] = \mathbb{E}_\tau[V(\mu)]$ for all $\omega \in \Omega$, where τ^ω is the conditional distribution of the receiver's beliefs given ω. Furthermore, Consistency* implies that for all $\omega \in \Omega$, τ^ω is absolutely continuous with respect to τ with Radon-Nikodym derivative $\frac{\mathrm{d}\tau^\omega}{\mathrm{d}\tau}(\mu) = \frac{\mu(\omega)}{p(\omega)}$. We then obtain:

$$\int V(\mu) \frac{\mu(\omega)}{p(\omega)} \, \mathrm{d}\tau(\mu) = \int V(\mu) \, \mathrm{d}\tau(\mu) \iff \mathrm{Cov}_\tau[V(\mu), \mu] = \mathbf{0}.$$

Therefore, whenever the indirect value correspondence has a single selection, we fully characterize the set of implementable distributions over posteriors under mediation.

Corollary 1. *If the indirect value correspondence is singleton-valued $\mathbf{V} = V$, then τ is implementable under mediation if and only if (τ, V) satisfy Consistency* and Honesty*.*

[15] Here, $\mathrm{Cov}_\tau[V(\mu), \mu]$ is a $(n-1)$–dimensional vector of one-dimensional covariances $\mathrm{Cov}_\tau[V(\mu), \mu(\omega)]$ between the sender's indirect utility and the receiver's posterior at each of $n-1$ states ω. One state is clearly redundant, hence the dimensionality is $n-1$.

An important case where the correspondence $\mathbf{V}$ is singleton-valued is when the receiver has a single best response $a^*(\mu) \in A$ to every possible posterior, for example when this is the conditional expectation of ω given the message received from the mediator.[16]

The zero covariance condition states that there cannot be any correlation between the payoff of the sender and the belief of the receiver. For example, consider the binary-state case $\Omega = \{\underline{\omega}, \overline{\omega}\}$ with a singleton-valued $\mathbf{V} = V$. The realized posterior belief is represented by the probability $\mu \in [0, 1]$ that the state is $\overline{\omega}$. Suppose that a candidate information structure induces a finitely supported distribution with possible pairs of sender's payoff and receiver's belief given by $\{(\mu_i, V(\mu_i))\}_{i=1}^k \subseteq \mathbb{R}^2$. In statistical terms, the zeroCov condition says that if we draw the regression line for the variable $V(\mu)$ with respect to the variable μ, then this line must be flat. Importantly, the property of having a flat regression line does not imply that there is no stochastic dependence between $V(\mu)$ and μ.

3.1 The Optimal Value of Mediation

Applying our Theorem 1, we can rewrite the mediator's problem in the belief space. The mediator chooses a distribution over receiver's posterior $\tau \in \Delta(\Delta(\Omega))$ and a measurable selection $V \in \mathbf{V}$ to maximize the sender's expected payoff:

$$\sup_{V \in \mathbf{V}, \tau \in \Delta(\Delta(\Omega))} \int V(\mu)\, \mathrm{d}\tau(\mu)$$

$$\text{subject to:} \int \mu\, \mathrm{d}\tau(\mu) = p \tag{BP}$$

$$\int V(\mu)(\mu - p)\, \mathrm{d}\tau(\mu) = \mathbf{0}, \tag{TT}$$

where (TT) is just a rewriting of (zeroCov). Let $g \in \mathbb{R}^n$ denote an arbitrary Lagrange multipliers for (TT) and, for any selection $V \in \mathbf{V}$, define the corresponding *virtual* indirect value function of the sender as $V^g(\mu) := (1 + \langle g, \mu - p \rangle)V(\mu)$. Each $V^g(\mu)$ is the belief-based version of the *virtual utility* in Myerson (1997) and Salamanca (2021) and, like those, takes into account a fixed shadow price g of the constraint TT.[17]

We next use these objects to characterize the optimal value of mediation. For any measurable function $U : \Delta(\Omega) \to \mathbb{R}$, let $\mathrm{cav}(U)(p)$ denote the concavification of U evaluated at p, that is, the pointwise infimum over all concave functions that majorize U.

[16] Kolotilin et al. (2024) give simple sufficient conditions on u_R such that the receiver has a single, yet possibly nonlinear, best response to every belief.

[17] Recall that the virtual utilities in both Myerson (1997) and Salamanca (2021) are defined on outcomes as opposed to beliefs.

Proposition 1. *The mediation problem admits solution* (V^*, τ^*) *and this solution can be implemented using a communication mechanism with no more than* $2n-1$ *messages. Moreover, the sender's optimal value under mediation is given by*

$$\mathcal{V}_{MD}(p) = \max_{V \in \mathbf{V}} \inf_{g \in \mathbb{R}^n} \text{cav}(V^g)(p).$$

We show the existence of a solution by constructing an auxiliary program in the space of joint distributions of the sender's expected values and receiver's posteriors that has also been analyzed in Lipnowski et al. (2022). Since $\mathbf{V}$ is upper hemi-continuous and closed-valued, its graph is closed, so the auxiliary program admits a solution. This implies our existence result. Note that (BP) and (TT) are in the form of moment conditions à la Winkler (1988), which implies that optimal mediation can be achieved with finitely many messages. Because the truth-telling constraint can be incorporated into the objective function via Lagrange multipliers, by Sion's minimax theorem, the sender's optimal value under mediation is the lower envelope of a family of concavified virtual utilities.

3.2 Bayesian Persuasion and Cheap Talk

We now recall how to analyze Bayesian persuasion and cheap talk using the belief-based approach. The classical interpretation of Bayesian persuasion is that the sender can commit to an information structure for the receiver before the state is realized. An alternative, yet mathematically equivalent interpretation, is that there is a mediator with commitment power that is completely aligned with the sender but, unlike in standard mediation, does not need to elicit the state from the sender. In this case, the mediator's problem drops (TT) and directly maximizes the expectation of the upper envelope $\overline{V}$ over all distributions over posteriors τ that satisfy (BP). Let $\mathcal{T}_{BP}(p)$ and $\mathcal{V}_{BP}(p)$ respectively denote the set of implementable distributions over posteriors and the optimal value under persuasion.

Under cheap talk, we completely bypass the mediator: after having observed the state, the sender sends a cheap talk message to the receiver. As the sender does not have commitment power, in equilibrium they must be indifferent among all the messages they send. Thus, the sender's problem under cheap talk replaces (TT) with the following stronger incentive compatibility constraint: the selected indirect value function $V(\mu)$ is constant over $\text{supp}(\tau)$. Therefore, the set of implementable distributions under cheap talk is $\mathcal{T}_{CT}(p) :=$ $\{\tau \in \mathcal{T}_{BP}(p) : \exists\, V \in \mathbf{V} \text{ such that } V \text{ is constant on } \text{supp}(\tau)\}$. An alternative way to represent the constraint under cheap talk is a zero variance constraint $\text{Var}_\tau[V] = 0$. Compared with the zero covariance condition (zeroCov), this illustrates the statistical difference between mediation and cheap talk: Under mediation, there cannot be any correlation between

μ and $V(\mu)$, whereas under cheap talk, these two must be stochastically independent.

Let $\mathcal{V}_{CT}(p)$ denote sender's value in their preferred cheap talk equilibrium. Because the sets of implementable distributions are nested, we have $\mathcal{V}_{BP}(p) \geq \mathcal{V}_{MD}(p) \geq \mathcal{V}_{CT}(p)$. Let $\bar{V}_{CT} : \Delta(\Omega) \to \mathbb{R}$ and $\underline{V}_{CT} : \Delta(\Omega) \to \mathbb{R}$ denote the quasiconcave envelope and the quasiconvex envelope of $\mathbf{V}$, respectively. That is, $\bar{V}_{CT}$ ($\underline{V}_{CT}$) is the pointwise infimum (supremum) over all quasiconcave (quasiconvex) functions that majorize $\bar{V}$ (are majorized by $\underline{V}$). Theorem 2 in Lipnowski and Ravid (2020) shows that the value of the sender's preferred cheap talk equilibrium coincides with the quasiconcave envelope of $\mathbf{V}$, that is $\bar{V}_{CT} = \mathcal{V}_{CT}$. Moreover, define the cheap-talk correspondence by $\mathbf{V}_{CT}(\mu) = [\underline{V}_{CT}(\mu), \bar{V}_{CT}(\mu)]$ for every $\mu \in \Delta(\Omega)$, the sender's value under a cheap talk equilibrium lies in $\mathbf{V}_{CT}$.[18]

Say that a distribution over posteriors τ is *deterministic* if $|\operatorname{supp} \tau^{\omega}| = 1$ for all $\omega \in \Omega$. When this is not the case and τ is implementable under mediation, then it must be induced by a random (direct) communication mechanism, that is $\sigma : \Omega \to \Delta(A)$ such that σ_{ω} is non-degenerate for some $\omega \in \Omega$.

Corollary 2. *A deterministic distribution over posteriors τ is implementable under mediation if and only if it is implementable under cheap talk.*

The full disclosure distribution $\tau_{FD} \coloneqq \sum_{\omega \in \Omega} p(\omega)\delta_{\omega}$ is deterministic, so it is implementable under mediation if and only if there exists $V \in \mathbf{V}$ such that $V(\delta_{\omega})$ is constant. Therefore, when full disclosure, or any other deterministic distribution τ, is sender optimal under mediation at p, we have $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$. Conversely, whenever $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$, Corollary 2 implies that *every* optimal distribution of beliefs under mediation must be induced by a random communication mechanism.

4 Persuasion vs. Mediation

In this section, we compare the sender's optimal value under persuasion and mediation.

Theorem 2. *Elicitation has no value if and only if commitment has no value, that is,*

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

Theorem 2 implies there are only three possible relationships among the values: $\mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$, $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$, or $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$. Com-

[18] See Lipnowski and Ravid (2020) Appendix C.2.1, which defines the quasiconcave and quasiconvex envelopes with an extra semi-continuity assumption. Our definition is the same since our state space Ω is finite. Alternatively, as shown in Aumann and Hart (2003) and Lipnowski and Ravid (2020), $\mathbf{V}_{CT}$ is the correspondence whose graph coincides with the di-convexification of the graph of $\mathbf{V}$.

bined with the geometric characterizations of the optimal persuasion value (Kamenica and Gentzkow (2011)) and the optimal cheap talk value (Lipnowski and Ravid (2020)), Theorem 2 also provides a geometric comparison between the sender's optimal value under commitment and their optimal value under any truthful communication mechanism: these are the same if and only if the concave and quasiconcave envelopes of the sender's value function coincide at the prior. Therefore, if the sender cannot achieve the optimal persuasion value using single-round cheap talk, then they cannot attain this via any communication mechanism without sender commitment (e.g. multiple-round cheap talk, noisy cheap talk).

The if direction of Theorem 2 is obvious. To see the only if direction, recall that the optimal persuasion value is attained from above by the minimal affine functional (i.e., a hyperplane) that dominates $\overline{V}(\mu)$ pointwise. Let $L_p(\mu) = \langle f_p, \mu \rangle$ denote this affine functional, where $f_p \in \mathbb{R}^n$ is its representing vector, and fix a finitely supported distribution τ that is optimal under persuasion and that is implementable under mediation.[19] The duality result in Dworczak and Kolotilin (2024) implies that $\overline{V}(\mu) = \langle f_p, \mu \rangle$ for all μ in the support of τ. In other words, f_p represents the regression hyperplane that passes through all the points $\{(\mu, \overline{V}(\mu))\}_{\mu \in \text{supp}(\tau)}$. The zeroCov condition of Theorem 1 implies that there exists an intercept $\alpha \in \mathbb{R}$ such that $\overline{V}(\mu) = \langle f_p, \mu \rangle = \alpha$ for all $\mu \in \text{supp}(\tau)$. Therefore, τ must be implementable under cheap talk because it induces a constant optimal value for the sender, hence $\mathcal{V}_{BP}(p) = \mathcal{V}_{CT}(p)$.

When Ω is non-binary, comparing the concave envelope and the quasiconcave envelope is not easy in general. Thus, we take a constructive approach and provide a sufficient condition for persuasion to strictly outperform mediation. To state the formal condition, we need the following definition.

Definition 2. The cheap talk hull is defined as

$$H^* := \{\mu \in \Delta(\Omega) : \exists \alpha > 1 \text{ such that } \overline{V}_{CT}(p) \in \mathbf{V}_{CT}(\alpha p + (1 - \alpha)\mu)\} \tag{1}$$

This is the set of beliefs such that the cheap talk value at p can be still attained when the prior is slightly perturbed.[20] Moreover, H^* is non-empty as $p \in H^*$ and convex.

Theorem 2 leads to the following sufficient condition for persuasion to strictly outperform mediation – it suffices to check whether there exists $\mu \in H^*$ where the sender's most preferred cheap talk equilibrium with prior μ is strictly better than the optimal cheap talk equilibrium with prior p.

[19] Given that we restrict to finitely many states, the finite-support assumption is innocuous.

[20] Observe that, by quasiconcavity and quasiconvexity respectively of $\overline{V}_{CT}$ and $\underline{V}_{CT}$ if $\mu \in H^*$ with respect to $\alpha > 1$, then every $\alpha' \in (1, \alpha]$ would work as well.

Proposition 2. *If there exists $\mu \in H^*$ such that $\overline{V}_{CT}(\mu) > \overline{V}_{CT}(p)$, then $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p)$.*

The proof is constructive. For any posterior $\mu \in H^*$ with $\overline{V}_{CT}(\mu) > \overline{V}_{CT}(p)$, there exists $\tau \in \mathcal{T}_{CT}(\mu)$ that attains $\overline{V}_{CT}(\mu)$ and $\tau' \in \mathcal{T}_{CT}((1-\alpha)\mu + \alpha p)$ that attains $\overline{V}_{CT}(p)$ for some $\alpha > 1$. The mixture $\frac{1}{\alpha}\tau' + \frac{\alpha-1}{\alpha}\tau$ is a distribution of beliefs centered at prior p that attains a value strictly higher than $\overline{V}_{CT}(p)$.[21]

We next introduce the notion of *full dimensionality* which allows us to make tighter the comparison between persuasion and mediation in this section and the one between cheap talk and mediation in the next section.

Definition 3. The full-dimensionality condition holds at p if $H^* = \Delta(\Omega)$.

In other words, the full-dimensionality condition requires that $\overline{V}_{CT}(p)$ can still be attained under cheap talk even when the prior is slightly perturbed toward *any* arbitrary direction.

Corollary 3. *Assume that the full-dimensionality condition holds at p. Then, $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p)$ if and only if there exists $\mu \in \Delta(\Omega)$ such that $\overline{V}_{CT}(\mu) > \overline{V}_{CT}(p)$.*

When does the full-dimensionality condition hold? In the binary-state case, it holds if the maximum cheap talk value is strictly higher than the maximum value achievable under no disclosure. In general, the next lemma provides a characterization of full dimensionality as well as an easy-to-verify sufficient condition.

Lemma 1. *The full-dimensionality condition holds at p if and only if there exists $\alpha > 1$ such that $\overline{V}_{CT}(p) \in \bigcap_{\omega\in\Omega} \mathbf{V}_{CT}(\alpha p + (1-\alpha)\delta_\omega)$. In particular, it holds provided that $\overline{V}_{CT}$ is locally constant around p.*

The first condition states that it is enough to check that all the Dirac beliefs are in the cheap talk hull at p. The second condition is particularly useful when the action set A is finite, because then $\overline{V}_{CT}$ is locally constant around p for almost every prior p as shown in Corollary 2 of Lipnowski and Ravid (2020). Combining this observation with our Corollary 3 yields that, when the action set is finite, for almost all priors, either cheap talk achieves the global maximum value or elicitation is strictly valuable.

Example 1. In the context of the illustrative example of Section 1.1, if the mediator has the expertise to assess the quality of the goods without relying on the seller's reports, they design (and commit to) a test/information structure about the quality of the goods that is

[21]A similar construction idea is applied in Corollary 2 of Lipnowski and Ravid (2020), which focuses on the optimal cheap talk value and implements this construction when $H^* = \Delta(\Omega)$. See the discussion about this full-dimensionality case below. In the proof of Proposition 2 we also provide an explicit lower bound on the difference $\mathcal{V}_{BP}(p) - \mathcal{V}_{MD}(p)$.

revealed to the buyer. The seller has a strict incentive to take this option because it relaxes the truth-telling constraint and allows the seller to induce any Bayesian persuasion outcome. For instance, the mediator 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 is greater than the payoff of the mediation plan we illustrated. Indeed, since the value of commitment is strictly positive, our Theorem 2 implies that the value of elicitation is strictly positive as well. Figure 1 (a) plots the CT value (red), the MD value (yellow), and the BP value (blue) over all the priors. Both elicitation and mediation are strictly valuable at every $p \in (0, 0.55)$. △

4.1 The Think Tank Revisited

We now illustrate the ideas introduced in this section with a three-state example. Think tanks often act as *research mediators* between an interest group and lawmakers. Here, we revisit the think-tank example in Lipnowski and Ravid (2020) by assuming that the sender is an interest group, say a lobbyist with private knowledge of the state, the receiver is a lawmaker with the option to maintain the status quo or to choose a new policy, and the mediator is a think tank which is completely aligned to the interest group.[22]

There are three possible states of the world $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and the lawmaker can take one of four actions $A = \{a_0, a_1, a_2, a_3\}$. Each action a_i for $i \in \{1, 2, 3\}$ represents a costly and risky policy that pays if and only if the state is ω_i. Differently, action a_0 is safe and represents the status quo. Formally, the lawmaker's payoff $u_R(\omega_i, a_j)$ is 1 if $i = j \neq 0$, 0 if $j = 0$, and $-c$ otherwise for some $c > 1$. The lobbyist is informed about the state of the world, but their preferences are misaligned with respect to the lawmaker. In particular, the lobbyist's payoff is $u_S(a) = \sum_{i=0}^{3} v_i \mathbb{I}[a = a_i]$ with $v_3 > v_2 > v_1 > v_0 = 0$, that is, the lobbyist prefers higher indexed policies and maintaining the status quo yields zero payoff.

Given belief $\mu \in \Delta(\Omega)$, the lawmaker's best response is to take action a_i if and only if $\mu(\omega_i) > \frac{c}{1+c}$, and they are indifferent between a_i and a_0 when $\mu(\omega_i) = \frac{c}{1+c}$. This is illustrated in the left panel of Figure 2. The colored regions at the vertexes of the simplex represent the beliefs such that the lobbyist's payoff is equal to v_i for some $i \in \{1, 2, 3\}$. The central hexagon is the region of the lawmaker's beliefs where their optimal response is to maintain the status quo, yielding a zero payoff for the lobbyist. Observe that the boundary segments between each colored region and the zero-payoff region represent the beliefs such that the

[22] In Lipnowski and Ravid (2020), the think tank does not have commitment power but does not need to elicit information from an interest group. Therefore, in their cheap-talk example, the think tank is the sender and tries to influence the lawmaker, i.e., the receiver.

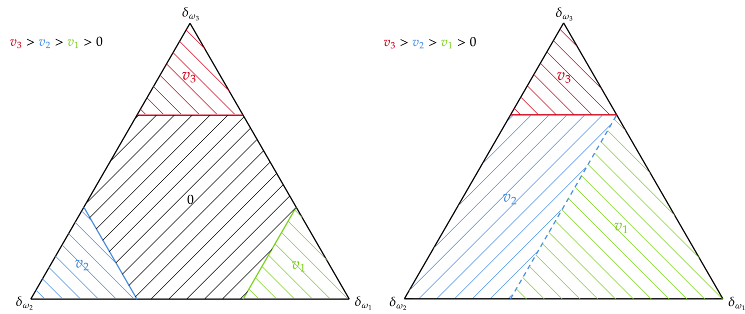

Figure 2: Lobbyist's value function and its quasiconcave envelope

Left panel: lobbyist's value correspondence over the lawmaker's belief space. Right panel: lobbyist's optimal cheap talk value (i.e., quasiconcave envelope) over the lawmaker's belief space. This illustrates the case where $c = 2$.

lawmaker is indifferent between the status quo and one of the new policies.

Suppose first that the lobbyist communicates with the lawmaker without the think tank mediation. This corresponds to the cheap-talk case and the lobbyist's optimal value as a function of the prior belief p is the quasiconcave envelope $\overline{V}_{CT}(p)$:

$$\mathcal{V}_{CT}(p) = \begin{cases} v_3 & \text{if } p(\omega_3) \geq \frac{c}{1+c} \\ v_1 & \text{if } p(\omega_1) \geq \frac{1}{1+c} \\ v_2 & \text{otherwise.} \end{cases}$$

The right panel of Figure 2 shows the level sets of the quasiconcave envelope over the simplex. When the prior is in one of the three colored regions in the left panel, then the babbling equilibrium is optimal for the lobbyist. Instead, the status-quo region can be split into two subregions. For priors that lie between the v_2 and v_3 regions, there exists an equilibrium distribution of the lawmaker's beliefs supported on posteriors where a_2 is uniquely optimal and posteriors where the lawmaker is indifferent between the status quo and a_3. Differently, for priors to the right of the blue dashed line, (BP) implies that any optimal equilibrium must induce a posterior where a_1 is optimal, implying the highest value attainable is v_1.

Given that the action set is finite, the full-dimensionality condition holds at almost all priors p in the simplex. For example, suppose that the prior p lies between the v_2 and v_3 region as in Figure 3. Around this prior, the quasiconcave envelope $\overline{V}_{CT}$ is constant and equal to v_2. This value is attained by the lobbyist-optimal distribution of the lawmaker's beliefs supported over $\{\mu_1, \mu_2, \mu_3, \mu_4\}$ as shown in Figure 3. At posteriors μ_2 and μ_3 the lawmaker takes action a_2, whereas on μ_1 and μ_4 the lawmaker mixes between the status quo and action a_3 so to induce exactly a payoff equal to v_2 for the lobbyist.[23]

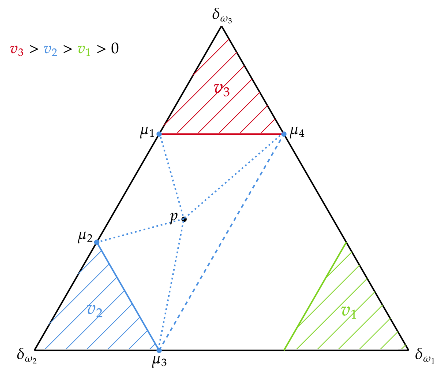

Figure 3: Construction of a cheap-talk equilibrium with a full-dimensional cheap talk hull.

Assume now that the lobbyist and the lawmaker communicate through the mediation of the think tank. We can apply Corollary 3 to establish when the think tank mediation secures to the lobbyist the Bayesian persuasion value. This happens if and only if the prior lies in the red triangle in the left panel of Figure 2. In this case, no disclosure is optimal for all three of the communication protocols considered. As soon as the prior p is outside this region, that is when $p(\omega_3) < \frac{c}{1+c}$, we have $\overline{V}_{CT}(\mu) > \overline{V}_{CT}(p)$ for all μ in the v_3 region, yielding that $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p)$. Thus, for a large set of prior beliefs, a lobbyist with commitment power would be strictly better off than the case where they communicate through an uninformed think tank with commitment, that is, the value of elicitation is strictly positive.

5 Mediation vs. Cheap Talk

This section provides separate sufficient and necessary conditions for the mediator to strictly outperform direct communication. These conditions collapse under the full-dimensionality condition introduced in the previous section, yielding a tight geometric characterization of when mediation is strictly valuable. As we have seen in Corollary 2, the mediator must randomize to strictly improve on cheap talk. Here, we show that they must randomize over posteriors with a value strictly above and below the optimal cheap talk value.[24]

Definition 4. We say that cheap talk is (locally) improvable at p if there exist $\mu \in \Delta(\Omega)$ ($\mu \in H^*$) and $\lambda \in (0,1)$ such that

$$\overline{V}_{CT}(\lambda\mu + (1-\lambda)p) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\mu).$$

[23] In our belief-based approach, this amounts to take a v_2 as a selection from $\mathbf{V}(\mu_1) = \mathbf{V}(\mu_4) = [0, v_3]$.
[24] Recall that the cheap talk hull H^* is defined in (1).

In words, cheap talk is locally improvable at p if there are alternative priors $\mu \in H^*$ and $\mu' = \lambda\mu + (1-\lambda)p$ such that there exists a cheap talk equilibrium at μ and one at μ' that respectively yield a strictly lower and a strictly higher expected payoff to the sender. Importantly, the prior μ' corresponding to the high-value equilibrium has to be "closer" to the original prior p, in the sense that μ' lies in the open segment (p, μ).

We can now state the main result of this section.

Theorem 3. *The following hold:*

1. *If cheap talk is locally improvable at* p*, then* $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$ *and every optimal distribution of beliefs under mediation is induced by a random communication mechanism.*

2. *Conversely, if cheap talk is not improvable at* p*, then* $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$*.*

Moreover, if the full-dimensionality condition holds at p*, then* $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$ *if and only if cheap talk is improvable at* p*.*

We sketch the proof of Theorem 3 here. If cheap talk is locally improvable at p, then there exist two beliefs $\mu_- \in H^*, \mu_+ \in (p, \mu_-)$, and two cheap talk equilibria $\tau_- \in \mathcal{T}_{CT}(\mu_-), \tau_+ \in \mathcal{T}_{CT}(\mu_+)$ that respectively attain a value strictly lower and strictly higher than $\bar{V}_{CT}(p)$. Because $\mu_- \in H^*$, there exists μ_0 that lies on the half line with endpoint μ_- through p, such that $\bar{V}_{CT}(p)$ can be attained by a cheap talk equilibrium τ_0 centered at μ_0. Next, consider a new distribution of beliefs τ obtained by randomizing over the three cheap talk equilibria τ_+, τ_-, and τ_0. Because the barycenters of these three distributions are colinear, finding the weights on each of them such that (BP) and (TT) are satisfied reduces to a 1-dimensional problem. Moreover, since μ_+ is "closer" to the prior p compared to μ_-, (TT) requires the mediator to assign a relatively higher weight to τ_+ compared to τ_-, so the sender's expected utility is strictly higher than $\bar{V}_{CT}(p)$ with this randomization. See Figure 4 in subsection 5.1 for a graphical illustration of this construction. Note that this procedure also provides a lower bound on the value of mediation, which depends on the barycenters and cheap talk equilibria in the construction.[25]

The proof of the converse statement is more technical. If cheap talk is not improvable at p, then there exists a hyperplane H that properly separates all posteriors with values strictly higher than $\bar{V}_{CT}(p)$ from those with values strictly lower than $\bar{V}_{CT}(p)$. Moreover, the prior p lies in the same closed half-space as the posteriors with a value strictly below $\bar{V}_{CT}(p)$. A normal vector $g \in \mathbb{R}^n$ of H is a Lagrange multiplier for the (TT) constraint such that $(V(\mu) - \bar{V}_{CT}(p))\langle g, \mu\rangle \leq 0$ for every $\mu \in \Delta(\Omega)$. Hence, for any (τ, V) implementable

[25]See equation 5 in Appendix A.4 for an explicit expression of this lower bound.

under mediation, we have

$$0 \geq \int (V(\mu) - \overline{V}_{CT}(p))\langle g, \mu\rangle \, \mathrm{d}\tau(\mu) = \left(\int V(\mu) \, \mathrm{d}\tau(\mu) - \overline{V}_{CT}(p)\right) \langle g, p\rangle,$$

by (zeroCov) and (BP). When p does not lie on H, we conclude that $\int V \, \mathrm{d}\tau \leq \overline{V}_{CT}(p)$. Otherwise, for τ-almost-all posteriors μ, either $V(\mu) = \overline{V}_{CT}(p)$ or $\mu \in H \cap \Delta(\Omega)$, which is a strictly lower-dimensional set. We can find another separating hyperplane H' while restricting attention to $H \cap \Delta(\Omega)$ and then repeat the same argument until p is not in the separating hyperplane or until the intersection of all separating hyperplanes $H \cap H' \cap \Delta(\Omega)$ is a singleton p. Either case leads to the desired conclusion that $\mathcal{V}_{MD}(p) \leq \overline{V}_{CT}(p)$.

For the last part of the theorem, full dimensionality implies that cheap talk is locally improvable at p if and only if it is improvable at p. In general, full dimensionality holds when the quasiconcave envelope $\overline{V}_{CT}$ is locally flat at p (see Lemma 1), which is the case for almost every prior p when the action set A is finite.

When the sender's payoff correspondence is singleton-valued and no disclosure is not a sender's optimal cheap talk equilibrium, it is possible to simplify the characterization of Theorem 3 as follows.

Corollary 4. *Assume that $\mathbf{V} = V$ is singleton-valued, that the full-dimensionality condition holds at p, and that no disclosure is suboptimal for cheap talk at p (i.e., $\overline{V}_{CT}(p) > V(p)$). Then $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$ if and only if there exists $\mu \in \Delta(\Omega)$ such that*

$$\overline{V}_{CT}(\mu) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\mu).$$

In this case, it is sufficient to find a single alternative prior μ that admits two cheap talk equilibria respectively inducing a strictly higher and a strictly lower sender's payoff than the sender's optimal cheap talk value at p.

Finally, it is natural to ask whether, under the sufficient condition of Theorem 3, mediation also strictly improves the expected utility of the receiver, $\int V_R(\mu) d\tau(\mu)$, where $V_R(\mu) := \max_{a \in A} \mathbb{E}_\mu[u_R(\omega, a)]$ is the receiver's utility given posterior μ. This is indeed the case provided that $\mathbf{V} = V$ is singleton-valued and that $V_R(\mu) = \phi(V(\mu))$ for some strictly increasing and convex function $\phi : \mathbb{R} \to \mathbb{R}$. In Section 6.1, we analyze a large class of games where this condition is satisfied. Beyond this condition, it is not always easy to adapt our approach to conclude whether there exists a mediation plan that improves both the sender's and receiver's expected payoff compared to their payoffs under some sender-preferred cheap talk equilibrium. However, this is the case in the illustrative example in the introduction as well as in the illustration in Section 5.1.[26]

5.1 Valuable Mediation in the Think-Tank Example

Consider again the setting of Section 4.1 with a lobbyist (sender) trying to influence a lawmaker (receiver) through a think tank (mediator). Here, we use the results of this section to show when the mediation of the think tank is strictly valuable. Recall that in this case, the full dimensionality condition holds at almost every prior.

Suppose first that the prior p lies between the v_2 and v_3 region as in Figure 3. Observe that the lawmaker's beliefs μ' such that $\overline{V}_{CT}(\mu') > \overline{V}_{CT}(p) = v_2$ are those in the v_3 region (the red triangle). Therefore, it is not possible to find a belief μ and a point μ' in the segment (p, μ) as described in Definition 4. To see this, note that if $\overline{V}_{CT}(\mu') > v_2$ for some $\mu' \in (p, \mu)$, then μ must be in the v_3 region except the boundary red line where the lobbyist is indifferent between a_3 and a_0, yielding that $\underline{V}_{CT}(\mu) = \overline{V}_{CT}(\mu) = v_3$. This logic holds for all priors in the central hexagon and at the left of the dashed blue line in Figure 3. That is, for any p with $p(\omega_1) < \frac{1}{1+c}$, cheap talk is not improvable at p, so the think tank is worthless.

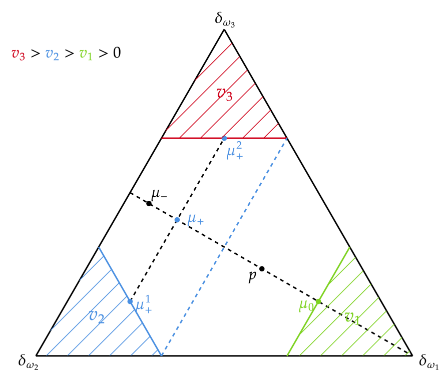

Figure 4: Construction of an improving distribution of beliefs under mediation

Differently, consider a prior p to the right of the same dashed blue line as in Figure 4, that is such that $p(\omega_1) > \frac{1}{1+c}$. At all these priors, cheap talk is improvable, so by Theorem 3 mediation strictly improves on direct communication. Intuitively, mediation helps strictly when the lawmaker has a pessimistic prior belief. Figure 4 graphically constructs an improving distribution of beliefs that is feasible under mediation following the logic of Theorem 3. First, recall from Figure 2 that $\overline{V}_{CT}(p) = v_1 > 0$. Next, fix μ_- and $\mu_+ \in (p, \mu_-)$ lying in the same segment as in Figure 4. Both these two beliefs are to the left of the blue dashed line, implying that $\overline{V}_{CT}(\mu_+) = \overline{V}_{CT}(\mu_-) = v_2 > v_1$. Moreover, $\underline{V}_{CT}(\mu_-) = \underline{V}_{CT}(\mu_+) = 0$, the payoff of the babbling equilibrium. This shows that cheap talk is improvable at p. Next, consider a

[26]See also the discussion at the end of Section 7.

distribution τ_+ of the lawmaker's beliefs that is supported on $\{\mu_+^1, \mu_+^2\}$ and with barycenter μ_+. This is a feasible distribution of beliefs under cheap talk at prior μ_+ since we can select a mixed best response at μ_+^2 that induces expected payoff v_2 for the lobbyist. Importantly, this distribution of beliefs and selection gives an expected payoff $\overline{V}_{CT}(\mu_+) = v_2 > v_1$ to the lobbyist. Consider also two degenerate distributions of beliefs $\tau_- = \delta_{\mu_-}$ and $\tau_0 = \delta_{\mu_0}$, where μ_0 lies at the intersection of the previous segment and the boundary between the status-quo region and the v_1 region.[27] Because the barycenters of $\tau_+, \tau_-,$ and τ_0 are colinear, we can mix these distributions as in the illustrative example in the introduction such that (BP) and (TT) are satisfied while strictly improving the expected payoff of the lobbyist.

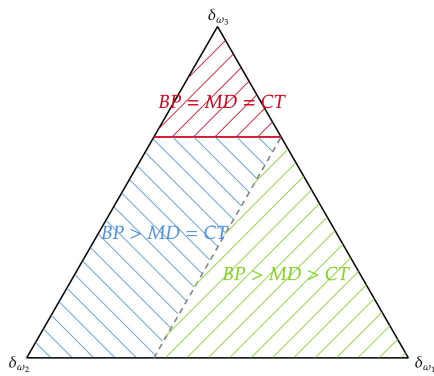

Figure 5: Relationships among communication protocols

For p with $p(\omega_1) \geq \frac{c}{1+c}$, no disclosure is optimal under cheap talk and suboptimal under mediation. Hence, the optimal mediation solution is strictly more informative than an optimal cheap talk equilibrium under these priors. Moreover, as the cost c increases, the region where the cheap talk is improvable expands, and it converges to the entire simplex as $c \to \infty$. Therefore, mediation by a think tank is more likely to be valuable for high-stakes decisions. In general, the dotted blue line in Figure 4 separates the status-quo hexagon into two regions: to its left elicitation is strictly valuable but mediation is not, to its right both elicitation and mediation are strictly valuable. The relations among the three protocols are summarized in Figure 5. All the three possible scenarios that we mentioned after Theorem 2 are present in the current example: For priors p in the red region we have $\mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$, for p in the blue region $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$, and for p in the green region $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$.[28]

Finally, we show that for every p such that $p(\omega_1) \in \left(\frac{1}{1+c}, \frac{c}{1+c}\right)$, there is a distribution

[27] In principle, there are multiple ways to construct μ_0 and τ_0, and μ_0 is not required to lie in the v_1 region. By full dimensionality, any μ_0 in a neighborhood of p attains v_1 under cheap talk. Hence, for any selection of μ_-, we can choose a μ_0 in the extended segment $(\mu_-, p]$ through p where v_1 is attained under cheap talk with some distribution τ_0. We choose the simplest one for illustration here.

[28] The dotted grey line in Figure 5 is a zero-measure region where full dimensionality does not hold.

of beliefs $\tau \in \mathcal{T}_{MD}(p)$ under which both the lobbyist's and lawmaker's expected payoff are strictly higher than their payoff under a lobbyist-preferred cheap talk equilibrium. Consider a lobbyist-preferred cheap talk equilibrium $\tau' \in \mathcal{T}_{CT}(p)$ that is supported on μ_3, μ_4 and some posteriors on the boundary of the v_1 region as in Figure 3. At every posterior in the support of τ', the lawmaker is indifferent between a_0 and some other action, so the lawmaker's expected payoff is 0 under τ'. We've illustrated that mixing among three cheap talk equilibria τ_+, τ_- and τ_0 with different but colinear barycenters yields a $\tau \in \mathcal{T}_{MD}(p)$ that strictly improves the lobbyist's payoff. Different from the illustration, we now take a τ_0 that supports on μ_3, μ_4 and δ_{ω_1}. The lawmaker takes action a_1 with certainty at posterior δ_{ω_1}, so their expected payoff at δ_{ω_1} is 1. Hence, the lawmaker's expected utility under τ is strictly positive.

5.2 Binary-state case and quasi single-crossing

Here, we provide a geometric comparison between mediation and cheap talk when Ω is binary. This is captured by a weaker version of the single-crossing condition.

Definition 5. A compact-valued correspondence $\mathbf{U} : [0,1] \rightrightarrows \mathbb{R}$ is *quasi single-crossing at* $x_0 \in [0,1]$ *from below (above)* if $\bar{U}(x_0) = 0$ and

- for all $x \le x_0$ $(x \ge x_0)$, $\bar{U}(x) \le 0$;
- for all $x' > x > x_0$ $(x' < x < x_0)$, $\bar{U}(x) > 0$ implies $\underline{U}(x') \ge 0$,

where $\bar{U}(x) = \max \mathbf{U}(x)$, $\underline{U}(x) = \min \mathbf{U}(x)$. $\mathbf{U}$ is *quasi single-crossing at* x_0 if it is quasi single-crossing at x_0 either from below or from above.[29]

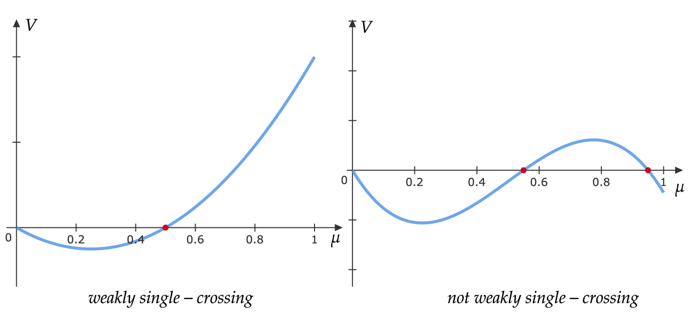

Figure 6: Comparison of quasi single-crossing and not quasi single-crossing functions

[29]When $\mathbf{U}$ is singleton-valued, we obtain the corresponding definition for functions. In Appendix C.1, we relate this notion of quasi single-crossing function with the standard definition of single-crossing function as well as with the notion of weak single-crossing function in Shannon (1995).

Corollary 5. *Assume that $|\Omega| = 2$ and that the full-dimensionality condition holds at p. Then $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ if and only if $\mathbf{V}_{CT} - \overline{V}_{CT}(p)$ is quasi single-crossing at p.*[30]

Corollary 5 follows from Theorem 3' that we introduce below. In particular, recall that when the state is binary, the full-dimensionality condition holds at p provided that the babbling equilibrium is not sender optimal. In Appendix C.2, we document other corollaries for the binary-state case, such as a geometric sufficient condition for mediation to be valueless.

The failure of the weak single-crossing property captures the idea of countervailing incentives, that is, the sender would like to induce more optimistic beliefs for some realized posterior and more pessimistic beliefs for some others. Here, countervailing incentives allow us to construct a strictly improving mediation plan as in the proof of Theorem 3.[31]

We can recast our geometric characterization in the general case with an arbitrary number of states in terms of the weak single-crossing property. Let $\mathcal{E}$ be the set of line segments with two endpoints on the boundary of $\Delta(\Omega)$. Any $\ell \in \mathcal{E}$ can be parameterized as $\ell = \{t\mu + (1-t)\mu' : t \in [0,1]\}$ for some boundary $\mu, \mu' \in \Delta(\Omega)$. For any $\ell \in \mathcal{E}$, fix a parametrization and let $\ell(\mu) \in [0,1]$ denote the parametrized representation of $\mu \in \ell$. Finally, $\mathbf{V}_{CT|\ell} : [0,1] \rightrightarrows \mathbb{R}$ denotes the restriction of the cheap talk correspondence over $\ell \in \mathcal{E}$.

We can now restate our main theorem in terms of the weak single-crossing property.

Theorem 3'. Suppose the full-dimensionality condition holds at p. Then $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ if and only if, for every $\ell \in \mathcal{E}$ with $p \in \ell$, $\mathbf{V}_{CT|\ell} - \overline{V}_{CT}(p)$ is quasi single-crossing at $\ell(p)$.

Under full dimensionality, it is enough to find a line segment ℓ containing p such that $\mathbf{V}_{CT|\ell} - \overline{V}_{CT}(p)$ fails the weak single-crossing condition to imply that we can construct a strictly improving mediation plan as in the proof of Theorem 3.

6 Moment Mediation: Quasiconvex Utility

In this section, we apply the results from Section 5 to *moment-measurable mediation*. For $1 \leq k \leq n-1$, a k-dimensional moment is a linear function $T : \Delta(\Omega) \to \mathbb{R}^k$ such that the set of relevant moments $X := T(\Delta(\Omega))$ has dimension k. We assume that $\mathbf{V} = V$ is singleton-valued and specifically that $V(\mu) = v(T(\mu))$ for some continuous $v : \mathbb{R}^k \to \mathbb{R}$ and

[30] Recall that $\mathbf{V}_{CT}(\mu) - \overline{V}_{CT}(p) := \{v - \overline{V}_{CT}(p) \in \mathbb{R} : v \in \mathbf{V}_{CT}(\mu)\}$ for all $\mu \in \Delta(\Omega)$. When the state is binary, we abuse the notation so that μ denotes the first entry of the receiver's posterior.

[31] Corollary 5 does not apply when the full-dimensionality fails at p. In this case, no disclosure is optimal under cheap talk, so $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ if and only if no disclosure is optimal under mediation. When $\mathbf{V} = V$ is singleton-valued, applying the results in Dworczak and Kolotilin (2024), no disclosure is optimal under mediation if there exists $g \in \mathbb{R}$ such that $V^g(\mu)$ in Proposition 1 is superdifferentiable at p. This becomes an if and only if when strong duality holds for the mediation program. Differently from Bayesian persuasion, this is not true in general for mediation as we show via example in Appendix E.

k-dimensional moment $T(\mu)$. Here, we focus on the multidimensional case ($k > 1$) under the assumption that $v(x)$ is strictly quasiconvex. This is the main case considered in past works on multidimensional cheap talk under transparent motives (see Chakraborty and Harbaugh (2010) and Lipnowski and Ravid (2020)).[32] The analysis of the one-dimensional case ($k = 1$) for general $v(x)$ is similar to that for the binary-state case in Appendix C.2 and is relegated to Appendix F.

When $v(x)$ is strictly quasiconvex and the full-dimensionality condition holds at p, only two extreme cases can happen: either all the communication protocols attain the global max of V or the optimal sender's value across communication protocols, including no disclosure, are all strictly separated. Hence, elicitation, mediation, and communication are all strictly valuable in the latter case.

Theorem 4. *Assume that $V(\mu) = v(T(\mu))$ for some k-dimensional moment T ($k \geq 2$) and continuous and strictly quasiconvex $v(x)$. If the full-dimensionality condition holds at p, then exactly one of these cases holds:*

(1) $\max V = \mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p) > V(p)$;

(2) $\max V > \mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p) > V(p)$.

Corollary 6 in Lipnowski and Ravid (2020) shows that under strict quasiconvexity no disclosure is suboptimal under cheap talk. In addition, we show that strict quasiconvexity and full-dimensionality imply that cheap talk is improvable at p if and only if its optimal value is strictly below the global max of V. Finally, the strict separation between Bayesian persuasion and mediation in (2) comes from Theorem 2.

While Theorem 4 dramatically simplifies the comparison among communication protocols in the present setting, it still relies on the full-dimensionality condition. We now provide an easy-to-check condition that implies the existence of a non-trivial set of priors that satisfy full dimensionality when v is strictly quasiconvex. Define $X_T := \{T(\delta_\omega) \in \mathbb{R}^k : \omega \in \Omega\}$.

Definition 6. We say that $v(x)$ is *minimally edge non-monotone given* T if there exists $\underline{x} \in \operatorname{argmin}_{\tilde{x} \in X_T} v(\tilde{x})$ such that for all $x \in X_T \setminus \{\underline{x}\}$, the one-dimensional function $\hat{v}_x(\lambda) := v(\lambda x + (1-\lambda)\underline{x})$ is not weakly increasing in $\lambda \in [0, 1]$.

The utility function $v(x)$ is minimally edge non-monotone given T whenever the one-dimensional restrictions of $v \circ T$ over the segments between the worst possible degenerate belief and any alternative degenerate belief are all non-monotone. This property captures

[32] Quasiconvex sender's utilities play an important role also in the informed information design model of Koessler and Skreta (2023).

the idea of countervailing incentives that we mentioned in the introduction. When $v(x)$ is both strictly quasiconvex and minimally edge non-monotone given T, it follows that the one-dimensional function $\hat{v}_x(\lambda)$ defined above is strictly single-dipped with a unique minimum at some $\lambda_x \in (0,1)$.

Proposition 3. *Assume that $V(\mu) = v(T(\mu))$ for some k-dimensional moment T ($k \geq 2$) and that $v(x)$ is continuous, strictly quasiconvex, and minimally edge non-monotone given T. Then there exists an $(n-1)$–simplex $\tilde{\Delta} \subseteq \Delta(\Omega)$ such that the full-dimensionality condition holds for all $p \in \operatorname{int}\tilde{\Delta}$. For every such p, point (2) of Theorem 4 holds if and only if $\min_{x\in X_T} v(x) < \max_{x\in X} v(x)$.*

In the proof, we derive an explicit expression for the simplex $\tilde{\Delta}$, that is,

$$\tilde{\Delta} \coloneqq \operatorname{co}\{\delta_{\underline{\omega}}, \{\mu_\omega \in \Delta(\Omega) : \omega \in \Omega \setminus \{\underline{\omega}\}\}\},$$

where $\underline{\omega}$ is an element in $\operatorname{argmin}_{\omega\in\Omega} v(T(\delta_\omega))$ and, for every $\omega \in \Omega \setminus \{\underline{\omega}\}$, μ_w is the unique element of the one-dimensional segment $(\delta_{\underline{\omega}}, \delta_\omega]$ such that $v(T(\delta_{\underline{\omega}})) = v(T(\mu_\omega))$.[33] Full dimensionality holds at every $p \in \operatorname{int}\tilde{\Delta}$ as strict quasiconvexity implies that at every such prior, there exists an optimal cheap talk equilibrium supported on *all* extreme points of $\tilde{\Delta}$.

6.1 Acceptance Games and Pareto Improving Mediation

In this section, we apply our results for moment mediation to a class of *acceptance games* and provide sufficient conditions such that mediation is strictly (ex-ante) Pareto improving for both the sender and the receiver with respect to unmediated communication.

The receiver has a binary choice: whether to accept or reject a certain prospect. Given a posterior belief μ, the prospect's value $R(x)$ for the receiver depends on k-dimensional moments $x = T(\mu) \in \mathbb{R}^k$, where $R : X \to \mathbb{R}$ is a continuous value function. For example, when $\Omega \in \mathbb{R}$, the prospect's value may depend on both the mean and the variance of ω.

The receiver compares the ex-post value $R(x)$ to an outside option with value $\varepsilon \in \mathbb{R}$ which is their private information and is drawn from a strictly increasing and continuously differentiable CDF G. We assume that $R(X) \subseteq \operatorname{supp}\varepsilon$ so the outside option is competitive.

The sender acts to maximize the ex-ante probability that the receiver accepts the prospect given moments x, that is, $v(x) \coloneqq G(R(x))$. Observe that the receiver's ex-ante payoff given x is $v_R(x) \coloneqq H(R(x))$ where $H(r) \coloneqq \int \max\{\varepsilon, r\}\, \mathrm{d}G(\varepsilon)$ is a strictly increasing function. In

[33] For every $\omega \in \Omega \setminus \{\underline{\omega}\}$, μ_ω is well-defined because of strict quasiconvexity and minimal edge non-monotonicity.

other words, the payoffs of both the sender and the receiver are strictly increasing transformation of the prospect's value $R(x)$.[34]

To avoid the trivial case where unmediated communication attains the global maximum payoffs for both the sender and the receiver, we assume that $\min_{x \in X_T} R(x) < \max_{x \in X} R(x)$ for the rest of the section (cf. Proposition 3). We say that mediation is *strictly (ex-ante) Pareto improving* at p if there exists a distribution $\tau^* \in \mathcal{T}_{MD}(p)$ feasible under mediation such that

$$\int v(T(\mu))\,\mathrm{d}\tau^*(\mu) > \int v(T(\mu))\,\mathrm{d}\tau(\mu) \quad \text{and} \quad \int v_R(T(\mu))\,\mathrm{d}\tau^*(\mu) > \int v_R(T(\mu))\,\mathrm{d}\tau(\mu),$$

for all distributions $\tau \in \mathcal{T}_{CT}(p)$ feasible under cheap talk. Recall that G is log-concave if $\log(G)$ is a concave function. Log-concavity is a standard assumption that is satisfied by several parametric families of distributions on the real line.

Proposition 4. *If R is strictly quasiconvex and minimally edge non-monotone given T, and G is log-concave, then there exists an $(n-1)$–simplex $\tilde{\Delta} \subseteq \Delta(\Omega)$ such that mediation is strictly (ex-ante) Pareto improving at p for all $p \in \operatorname{int} \tilde{\Delta}$.*

The assumptions on R and the fact that G is strictly increasing allow us to invoke Proposition 3 to conclude that the sender is strictly better off under mediation than under chap talk for all $p \in \operatorname{int} \tilde{\Delta}$. Next, observe that $v_R(x) = \phi(v(x))$ where $\phi(z) := H(G^{-1}(z))$, and that log-concavity of G implies that ϕ is a convex function. With this, Jensen inequality and the fact that the sender's payoff must be constant over all the possible receiver's belief imply that the receiver is strictly better off under mediation than under chap talk for all $p \in \operatorname{int} \tilde{\Delta}$.

We stress that the assumptions on R are not needed for the last step, that is, in these acceptance games the receiver is strictly better off under mediation than under cheap talk whenever the sender is strictly better off and G is log-concave.

We now illustrate Proposition 4 through a simple example of an acceptance game.

Example 2 (Uncertain projects and costly investments)**.** An innovator (sender) is proposing a bundle of two projects with potential payoffs $\omega \in \Omega = \{(1,0), (0,1), (0,2)\}$ for a firm (receiver), where the i-th coordinate of ω represents project i's payoff with $i \in \{1,2\}$. In each state only one project pays out, but project 2 is potentially more valuable in one state. Given a proposal, the firm can accept or decline it. If the firm declines, it obtains

[34] Observe that this is not sufficient by itself to conclude that a strict improvement in the ex-ante payoff of the sender must imply a strict improvement of the receiver's ex-ante payoff due to the potential nonlinearity of G and H.

an outside option ε with a value uniformly distributed in $[0,1]$. If the firm accepts, it exerts efforts $e_1, e_2 \in [0,1]$ with a cost $e_1^2 + e_2^2$ to implement the projects, and project i pays off with probability e_i. Hence, given posterior belief $\mu \in \Delta(\Omega)$, the firm's expected utility from accepting is a function of the posterior mean $x = \mathbb{E}_\mu(\omega) \in \mathbb{R}^2$, which is $R(x) = \max_{e_1,e_2 \in [0,1]} \{e_1 x_1 - e_1^2 + e_2 x_2 - e_2^2\} = \frac{1}{4}(x_1^2 + x_2^2)$. R is strictly convex and hence strictly quasiconvex. The one-dimensional function $f_1(\lambda) = R(\lambda(0,1)+(1-\lambda)(1,0)) = \frac{1}{4}(2\lambda^2 - 2\lambda + 1)$ and $f_2(\lambda) = R(\lambda(0,2) + (1-\lambda)(1,0)) = \frac{1}{4}(5\lambda^2 - 2\lambda + 1)$ are non-monotone in $\lambda \in [0,1]$. Therefore, R is also minimally edge non-monotone. As the distribution of the outside option is log-concave, Proposition 4 implies that there exists an $(n-1)$–simplex $\tilde{\Delta} \subseteq \Delta(\Omega)$ such that mediation is strictly Pareto improving for $p \in \operatorname{int} \tilde{\Delta}$. △

We close this section by showing that the *strict* quasiconvexity assumption in Proposition 4 can be relaxed in some cases.[35]

Example 3 (Financial Intermediation under Mean-Variance Preferences)**.** A financial issuer tries to convince an investor to invest in an asset with unknown return $\omega \in \Omega \subseteq \mathbb{R}$. The investor is risk-averse and cares about both the expected payoff and the variance. That is, the investor's payoff from investing is $\mathbb{E}_\mu(\omega) - \gamma \operatorname{Var}_\mu(\omega)$ for some $\gamma > 0$. Defining the two moments $x_1 = \mathbb{E}_\mu(\omega)$, $x_2 = \mathbb{E}_\mu(\omega^2)$, we may rewrite the investor's payoff given μ as $R(x) = \gamma x_1^2 + x_1 - \gamma x_2$. These preferences capture that investors must satisfy some risk requirements for their investment. In particular, γ can be interpreted as the shadow price on the constraint on the maximum variance in a portfolio selection problem. Importantly, these preferences are not necessarily monotone with respect to first-order stochastic dominance.

Suppose there are n states $0 = \omega_0 < \omega_1 < \ldots < \omega_{n-1} = 1$ with $n \geq 3$. Assume that the investor is risk averse enough: $\gamma > 1/\omega_i$ for all $\omega_i > 0$; and that the investor's outside option follows a uniform distribution on $[0,1]$. Let $\alpha_i = 1 - \frac{1}{\gamma \omega_i}$ and $\mu_i = \alpha_i \delta_{\omega_i} + (1-\alpha_i)\delta_0$. We next show that for all p in the interior of $\tilde{\Delta} = \operatorname{co}\{\delta_0, \{\mu_i : i = 1, \ldots, n-1\}\}$, the full-dimensionality condition holds and that mediation is strictly better than cheap talk.

Note that the issuer's payoff function $v(x) = R(x)$ is convex but not strictly quasiconvex in x, so we cannot directly apply Theorem 4 and Proposition 3. However, the same idea as in the proof could also help us to verify the claim. Fix any $\omega_i \neq 0$, we show the seller's payoff $V(\mu)$ is non-monotone on the edge of $\Delta(\Omega)$ that connects δ_0 and δ_{ω_i}. For every $\alpha \in [0,1]$, we have $V(\alpha \delta_{\omega_i} + (1-\alpha)\delta_0) = \alpha \omega_i - \gamma \alpha (1-\alpha) \omega_i^2$. This is a quadratic function that is non-monotone on $[0,1]$ and intersects 0 at $\alpha = 0$ or $1 - \frac{1}{\gamma \omega_i}$.

By construction, for all $p \in \operatorname{int} \tilde{\Delta}$, there exists $\tau \in \mathcal{T}_{CT}(p)$ that attains value 0. Note that V is convex by the convexity of v and linearity of T, so the set of posteriors that attains

[35] In Appendix D we consider an expanded version of the following example as well as an additional example without strict quasiconvexity.

value higher than 0 is contained in $\Delta(\Omega) \setminus \tilde{\Delta}$. Lemma 6 then implies 0 is the optimal cheap talk value for priors in $\tilde{\Delta}$. Finally, note that $v(x) \leq 0$ gives $x_2 \geq x_1^2 + x_1/\gamma$, so the lower contour set $\{v \leq 0\}$ is strictly convex. In Appendix D.1, we use an analogous argument to that of Theorem 4 to show that $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p) > V(p)$ for every $p \in \operatorname{int} \tilde{\Delta}$.

The issuer strictly benefits from mediation when the investor's prior is sufficiently pessimistic. Moreover, when the investor becomes more risk-averse (γ increases), then α_i also increases for all $i = 1, \ldots, n-1$. So the region where the issuer strictly benefits from mediation expands as the investor becomes more risk-averse. Finally, because the distribution of the outside option is log-concave, the investor also strictly benefits from mediation when the prior is in $\operatorname{int} \tilde{\Delta}$, that is mediation is strictly Pareto improving there. △

7 Discussion and Extensions

This section discusses some of the points left out from the main analysis and future research.

Correlated equilibria in long cheap talk and repeated games Under transparent motives, our results shed light on Nash and correlated equilibria payoffs in static and repeated games with asymmetric information.[36] We restrict to the finite-action case, an assumption that is consistent with most of the literature on this topic.

The sender-receiver games we studied in this paper are called *basic decision problems* in Forges (2020). First, consider the cheap-talk extended version of this game with (potentially infinite) rounds of pre-play communication, which is known as the long cheap talk (Aumann and Hart, 2003). Lipnowski and Ravid (2020) show the highest sender's expected payoff that is induced by a Nash equilibrium of this long cheap talk game coincides with the one-shot highest cheap talk value $\mathcal{V}_{CT}(p)$. For correlated equilibria, Forges (1985) shows that the highest sender's expected payoff coincides with the payoff induced by the sender's preferred communication equilibrium, that is $\mathcal{V}_{MD}(p)$.[37] Our results then imply that, for almost all priors p, correlated equilibria strictly increase the expected payoff of the sender if and only if cheap talk is improvable at p (Theorem 3).

Next, we consider the infinitely repeated version of this sender-receiver game where both players take actions but only the receiver's action is payoff relevant. This is the transparent-motive case of the repeated games of *pure information transmission* as defined in Forges (2020). The overall payoff of the players is given by the undiscounted time average of the one-period payoffs. Forges (1985) shows that the set of correlated equilibrium payoffs of this

[36] See the recent survey by Forges (2020).

[37] In this case, a single round of pre-play communication is sufficient.

game corresponds to the one induced by the communication equilibria of the stage game. Moreover, the results in Hart (1985) and Habu et al. (2021) imply that every sender's Nash-equilibrium payoff of this game corresponds to a sender's payoff of a one-stage cheap talk equilibrium. Theorem 3 implies that if cheap talk is improvable at p, then the sender's largest correlated-equilibrium payoff in the repeated game is strictly higher than their best Nash-equilibrium payoff. See Appendix H for more details.

Sender's interim efficiency Theorem 2 established that under transparent motives and with a single receiver (or multiple receivers and public information), mediation attains the ex-ante efficient value (i.e., Bayesian persuasion) if and only if the same value can be attained under cheap talk. This result can be generalized by replacing this notion of ex-ante efficiency with a notion of interim efficiency inspired by the analysis in Doval and Smolin (2024).

We say that $\tau \in \mathcal{T}_{BP}(p)$ is *fully interim efficient* if there exist $V \in \mathbf{V}$ and $\lambda \in \Delta(\Omega)$ with $\lambda(\omega) > 0$ for all $\omega \in \Omega$, such that

$$(\tau, V) \in \underset{\tilde{\tau} \in \mathcal{T}_{BP}(p), \tilde{V} \in \mathbf{V}}{\operatorname{argmax}} \sum_{\omega \in \Omega} \left(\int_{\Delta(\Omega)} \tilde{V}(\mu) \, \mathrm{d}\tilde{\tau}^{\omega}(\mu) \right) \lambda(\omega), \tag{2}$$

and we say τ is fully interim efficient with selection V if (τ, V) satisfies (2). When $\mathbf{V} = V$ is singleton-valued, fully interim efficient distributions τ induce interim sender's values $w = (\mathbb{E}_{\tau^{\omega}}[V]))_{\omega \in \Omega} \in \mathbb{R}^{\Omega}$ that are on the Pareto frontier of the Bayes welfare set introduced in Doval and Smolin (2024).[38] This set represents all the sender's interim expected payoffs that can be induced by some Blackwell experiments without requiring that the truth-telling constraint holds. Therefore, the points on its Pareto frontier represent vectors of interim sender's payoffs that cannot be Pareto improved by an alternative experiment. Here, we restrict to the fully efficient outcome where every state has a strictly positive Pareto weight, that is $\lambda(\omega) > 0$ for all $\omega \in \Omega$.

In Lemma 4 in the appendix we show that if $\tau \in \mathcal{T}_{MD}(p)$ is fully interim efficient, then $\tau \in \mathcal{T}_{CT}(p)$. This allows us to extend Theorem 2: A mediator can induce an efficient vector of the sender's interim payoffs if and only if the same vector can be induced via unmediated communication. Observe that Theorem 2 immediately follows from this more general result by just setting $\lambda = p$.

This result can also be interpreted as a *mediation's trilemma*. Consider the three following properties: (1) Information is public; (2) The sender's payoff is state-independent; (3) Mediation is fully interim efficient and strictly better than cheap talk. The previous result implies that these three properties are incompatible. Moreover, this is a proper trilemma in

[38] This immediately follows from their Theorem 2.

the sense that if we relax any one of (1) or (2), then mediation can be interim efficient and strictly better than cheap talk. We show this with two examples in Appendix G.1.

Receiver's utility and informativeness In some cases, our results can be used to show that communication mechanisms improving the sender's expected payoff also improve the receiver's expected payoff, that is mediation yields a strict ex-ante Pareto improvement (see Section 6.1). In general, our techniques can be extended beyond these cases. However, focusing on the receiver's expected utility would present a key new challenge, namely that the objective function in the mediation problem would be different from the utility function in the truth-telling constraint. A related point is the comparison of informativeness across the sender's optimal communication and cheap talk equilibria respectively. In general, this comparison seems ambiguous as suggested by our examples. In the illustration in the introduction, when the prior p is in a neighborhood of 0.6, the sender's optimal cheap-talk equilibrium would be no disclosure while the sender's optimal communication equilibrium would involve some nontrivial form of disclosure (see Figure 1). Conversely, in Appendix G.2, we modify this example and show that in this case there exists a neighborhood of priors p such that full disclosure is sender optimal under cheap talk but not under mediation. We leave both these interesting questions for future research.

A Proofs

A.1 Preliminaries

The proofs of the next two ancillary lemmas are standard and relegated to Appendix B.

Lemma 2 (Lemma 3 of Lipnowski and Ravid (2020), and a symmetric version)**.**

(1) *If* $F : [0,1] \rightrightarrows \mathbb{R}$ *is a Kakutani correspondence with* $\min F(0) \leq 0 \leq \max F(1)$*, and* $\bar{x} = \inf\{x \in [0,1] : \max F(x) \geq 0\}$*, then* $0 \in F(\bar{x})$*.*

(2) *If* $F : [0,1] \rightrightarrows \mathbb{R}$ *is a Kakutani correspondence with* $\max F(0) \geq 0 \geq \min F(1)$*, and* $\bar{x} = \inf\{x \in [0,1] : \min F(x) \leq 0\}$*, then* $0 \in F(\bar{x})$*.*

Proof of (1) is in Lipnowski and Ravid (2020) and we give a proof of (2) in Appendix B.

Lemma 3. *An outcome distribution* $\pi \in \Delta(\Omega \times A)$ *satisfies Obedience if and only if for every measurable* $\tilde{a} : A \to \Delta(A)$*,* $\int u_R(\omega, a)\, \mathrm{d}\pi(\omega, a) \geq \int u_R(\omega, \tilde{a})\, \mathrm{d}\pi(\omega, a)$*, where* $u_R(\omega, \tilde{a}) = \int_A u_R(\omega, a)\, \mathrm{d}\tilde{a}(a)$*.*

A.2 The Mediation Problem

Given two measurable spaces $(X,\Sigma),(X',\Sigma')$, a measure τ on (X,Σ), and a measurable function $f: X \to X'$, we let $(f)_{\#}\tau$ denote the pushforward measure of τ under f.

Proof of Theorem 1. We first show the only if direction. Suppose that $\tau \in \Delta(\Delta(\Omega))$ and $V: \Delta(\Omega) \to \mathbb{R}$ are induced by some communication equilibrium outcome $\pi \in \Delta(\Omega \times A)$. Note that τ is the pushforward measure of $\operatorname{marg}_A \pi \in \Delta(A)$ under the measurable function $\phi: A \to \Delta(\Omega)$ with $\phi(a) = \pi^a$. For every $\omega \in \Omega$,

$$\begin{aligned}\int_{\Delta(\Omega)} \mu(\omega)\,\mathrm{d}\tau(\mu) &= \int_A \phi(a)(\omega)\,\mathrm{d}\operatorname{marg}_A \pi(a) = \int_A \pi^a(\omega)\,\mathrm{d}\operatorname{marg}_A \pi(a)\\ &= \int_{\Omega\times A} \mathbb{I}[\tilde{\omega} = \omega]\,\mathrm{d}\pi(\tilde{\omega}, a) = p(\omega),\end{aligned}$$

where $\mathbb{I}$ denotes the indicator function and where the first equality follows because $\tau = (\phi)_{\#}\operatorname{marg}_A \pi$, the second equality follows by definition, the third equality follows by the law of iterated expectations, and the last equality follows by Consistency of π. Hence, τ satisfies Consistency*.

Since V is induced by π, $V(\mu)$ is the conditional expectation of u_S with respect to $\operatorname{marg}_A \pi$, conditional on $\phi(a) = \mu$. Note that by Obedience, π is supported on $a \in A^*(\mu)$ only, where $A^*(\mu) = \operatorname{argmax}_{a\in A} \mathbb{E}_\mu[u_R(\omega, a)]$ is nonempty-compact-valued and weakly measurable by the measurable maximum theorem (Aliprantis and Border, 2006, Theorem 18.19). Therefore, $V(\mu) \in [\min_{a\in A^*(\mu)} u_S(a), \max_{a\in A^*(\mu)} u_S(a)]$ and V is measurable, so Obedience* is satisfied. By Honesty of π and the fact that u_S does not depend on ω, we have $\mathbb{E}_{\pi^\omega}[u_S] = \mathbb{E}_{\pi^{\omega'}}[u_S]$ for any $\omega, \omega' \in \Omega$. By Consistency, we have $\frac{\mathrm{d}\pi^\omega}{\mathrm{d}\operatorname{marg}_A \pi}(a) = \frac{\pi^a(\omega)}{p(\omega)}$ for all $\omega \in \Omega$. Therefore,

$$\begin{aligned}&\int_A u_S(a)\,\mathrm{d}\pi^\omega(a) = \int_A u_S(a)\tfrac{\pi^a(\omega)}{p(\omega)}\,\mathrm{d}\operatorname{marg}_A \pi(a)\\ &= \int_A \mathbb{E}\left[u_S(a)\tfrac{\pi^a(\omega)}{p(\omega)} \mid \phi(a) = \mu\right]\mathrm{d}\operatorname{marg}_A \pi(a) = \int_A \mathbb{E}[u_S \tfrac{\mu(\omega)}{p(\omega)} | \phi^{-1}(\mu)]\,\mathrm{d}\operatorname{marg}_A \pi(a)\\ &= \int_A V(\phi(a))\tfrac{\phi(a)(\omega)}{p(\omega)}\,\mathrm{d}\operatorname{marg}_A \pi(a) = \int_{\Delta(\Omega)} V(\mu)\tfrac{\mu(\omega)}{p(\omega)}\,\mathrm{d}\tau(\mu),\end{aligned}$$

where the second equality is by iterated expectation, the third one is simply rewriting, the fourth one is by $V = V^\pi$, and the last equality is by the fact that $\tau = (\phi)_{\#}\operatorname{marg}_A \pi$. Therefore, there exists a constant $c \in \mathbb{R}$ such that $\int_{\Delta(\Omega)} V(\mu)\frac{\mu(\omega)}{p(\omega)}\,\mathrm{d}\tau(\mu) = \int_{\Delta(\Omega)} V(\mu)\frac{\mu(\omega')}{p(\omega')}\,\mathrm{d}\tau(\mu) = c$

for every $\omega, \omega' \in \Omega$. It follows that for all $\omega \in \Omega$, $\int_{\Delta(\Omega)} V(\mu)\mu(\omega)\, \mathrm{d}\tau(\mu) = c \cdot p(\omega)$, so

$$\int_{\Delta(\Omega)} V(\mu)\, \mathrm{d}\tau(\mu) = \sum_{\omega' \in \Omega} \int_{\Delta(\Omega)} V(\mu)\mu(\omega')\, \mathrm{d}\tau(\mu) = c \cdot \sum_{\omega' \in \Omega} p(\omega') = c.$$

As we have shown that τ satisfies (BP), it follows that for any $\omega \in \Omega$,

$$\begin{aligned}\int_{\Delta(\Omega)} V(\mu)\mu(\omega)\, \mathrm{d}\tau(\mu) &= \left(\int_{\Delta(\Omega)} V(\mu)\, \mathrm{d}\tau(\mu)\right) p(\omega) \\ &= \left(\int_{\Delta(\Omega)} V(\mu)\, \mathrm{d}\tau(\mu)\right) \left(\int_{\Delta(\Omega)} \mu(\omega)\, \mathrm{d}\tau(\mu)\right),\end{aligned}$$

which implies that $\mathrm{Cov}_\tau(V(\mu), \mu(\omega)) = 0$ for every $\omega \in \Omega$, so Honesty* holds.

Next, we show by construction that for any $\tau \in \Delta(\Delta(\Omega))$ and $V \in \mathbf{V}$ that satisfy Consistency* and Honesty*, there exists a communication equilibrium outcome π with $\mathbb{E}_\tau[V] = \mathbb{E}_\pi[u_S]$. Since $V \in \mathbf{V}$, by Lemma 2 of Lipnowski and Ravid (2020), there exists a measurable $\lambda : \Delta(\Omega) \to \Delta(A)$ such that for all $\mu \in \Delta(\Omega)$, $\lambda(\mu) \in \mathrm{argmax}_{\alpha \in \Delta(A)} \mathbb{E}_\mu[u_R(\alpha, \omega)]$ is a mixed best response for the receiver with posterior μ, and $V(\mu) = \int_A u_S(a)\, \mathrm{d}\lambda(\mu)(a)$.

Define $\pi \in \Delta(\Omega \times A)$ by $\pi(\{\omega\} \times D) = \int_{\Delta(\Omega)} \mu(\omega)\lambda(\mu)(D)\, \mathrm{d}\tau(\mu)$ for any $\omega \in \Omega$ and any Borel $D \subseteq A$. We show that π is a desired communication equilibrium outcome. First, note that for any $\omega \in \Omega$, $\pi(\omega, A) = \int_{\Delta(\Omega)} \mu(\omega)\lambda(\mu)(A)\, \mathrm{d}\tau(\mu) = \int_{\Delta(\Omega)} \mu(\omega)\, \mathrm{d}\tau(\mu) = p(\omega)$ by Consistency*, so π satisfies Consistency.

Note that by construction we have $\pi^\omega(D) = \int_{\Delta(\Omega)} \frac{\mu(\omega)}{p(\omega)} \lambda(\mu)(D)\, \mathrm{d}\tau(\mu)$ for any Borel $D \subseteq A$. That is, π^ω is an average of $\lambda(\mu) \in \Delta(A)$. So for any function $u : \Omega \times A \to \mathbb{R}$, we have

$$\begin{aligned}\mathbb{E}_\pi[u] = \sum_{\omega \in \Omega} p(\omega)\mathbb{E}_{\pi^\omega}[u] &= \sum_{\omega \in \Omega} p(\omega) \int_{\Delta(\Omega)} \frac{\mu(\omega)}{p(\omega)} \left(\int_A u(\omega, a)\, \mathrm{d}\lambda(\mu)(a)\right) \mathrm{d}\tau(\mu) \\ &= \int_{\Delta(\Omega)} \mathbb{E}_\mu \left(\int_A u(\omega, a)\, \mathrm{d}\lambda(\mu)(a)\right) \mathrm{d}\tau(\mu),\end{aligned}$$

where the first equality follows from iterated expectation, the second one is by linearity, and the third one is simply rewriting.

To see Obedience, take any measurable $\tilde{a} : A \to \Delta(A)$, by definition of λ, we have $\mathbb{E}_\mu \left(\int_A u_R(\omega, a)\, \mathrm{d}\lambda(\mu)(a)\right) \geq \mathbb{E}_\mu \left(\int_A u_R(\omega, \tilde{a})\, \mathrm{d}\lambda(\mu)(a)\right)$ for any $\mu \in \Delta(\Omega)$. Taking expectation with respect to τ, we have $\int u_R(\omega, a)\, \mathrm{d}\pi(\omega, a) \geq \int u_R(\omega, \tilde{a})\, \mathrm{d}\pi(\omega, a)$, and π satisfies Obedience by Lemma 3.

Finally, by definition of λ, we have

$$\mathbb{E}_{\pi^\omega}[u_S] = \int_{\Delta(\Omega)} \tfrac{\mu(\omega)}{p(\omega)} \left(\int_A u_S(a)\,\mathrm{d}\lambda(\mu)(a) \right) \mathrm{d}\tau(\mu) = \int_{\Delta(\Omega)} \tfrac{\mu(\omega)}{p(\omega)} V(\mu)\,\mathrm{d}\tau(\mu) = \mathbb{E}_\pi[u_S].$$

Hence, π satisfies Honesty as τ satisfies Honesty*. ∎

Next, define $I := [\min_{\mu\in\Delta(\Omega)} \underline{V}(\mu), \max_{\mu\in\Delta(\Omega)} \overline{V}(\mu)]$, where the minimum and maximum are attained because of the semi-continuity of $\underline{V}, \overline{V}$. We now introduce an auxiliary program

$$\sup_{\eta\in\Delta(\Delta(\Omega)\times I)} \int_{\Delta(\Omega)\times I} s\,\mathrm{d}\eta(\mu, s) \tag{η-MD}$$

$$\text{subject to:} \int_{\Delta(\Omega)\times I} \mu\,\mathrm{d}\eta(\mu, s) = p \tag{η-BP}$$

$$\eta(\mathrm{Gr}(\mathbf{V})) = 1 \tag{η-OB}$$

$$\int_{\Delta(\Omega)\times I} s(\mu - p)\,\mathrm{d}\eta(\mu, s) = \mathbf{0}, \tag{η-TT}$$

where $\mathrm{Gr}(\mathbf{V}) \subseteq \Delta(\Omega) \times I$ denotes the graph of $\mathbf{V}$. The three constraints (η-BP), (η-OB), and (η-TT) correspond to Consistency*, Obedience* and Honesty*, respectively. Note that for any η feasible in this program, $\tau = \mathrm{marg}_{\Delta(\Omega)}\,\eta$ and $V(\mu) = \mathbb{E}_\eta[s|\mu]$ are feasible under mediation. Moreover, for any (τ, V) feasible under mediation, $\eta(\mu, s) = \tau(\mu)\mathbb{I}[s = V(\mu)]$ is also feasible under the auxiliary program. So mediation has the same value as this auxiliary program, and the existence of a solution for one program implies the existence of a solution for the other one and vice versa.

Proof of Proposition 1. We first show the auxiliary program has an optimal solution η^*. Note that the integrand of the first and third constraints are continuous. Hence, for any sequence of feasible η_n that converges weakly to η, we have $p = \int \mu\,\mathrm{d}\eta_n \to \int \mu\,\mathrm{d}\eta$ and $\mathbf{0} = \int s(\mu - p)\,\mathrm{d}\eta_n \to \int s(\mu - p)\,\mathrm{d}\eta$. Note that $\mathrm{Gr}(\mathbf{V})$ is closed since $\mathbf{V}$ is upper hemi-continuous and closed-valued, so $1 = \limsup_n \eta_n(\mathrm{Gr}(\mathbf{V})) \leq \eta(\mathrm{Gr}(\mathbf{V}))$ by the Portmanteau Theorem. Hence, $\eta(\mathrm{Gr}(\mathbf{V})) = 1$, and η is also feasible under the auxiliary program. Therefore, the feasibility set of the auxiliary program is compact in the weak topology. As the objective function is continuous, there exists $\eta^* \in \Delta(\Delta(\Omega) \times I)$ that solves the auxiliary program. Then, $\tau^* = \mathrm{marg}_{\Delta(\Omega)}\,\eta^*$ and $V^*(\mu) = \mathbb{E}_{\eta^*}[s|\mu]$ are the desired solution that solves the mediation problem.

Fix the optimal V^* we constructed, and consider the mediation problem with a fixed selection V^*. We endow $\Delta(\Delta(\Omega))$ with the weak* topology induced by bounded and measurable functions over $\Delta(\Omega)$. Then, the objective $\int V^*\,\mathrm{d}\tau$ is affine and continuous in τ since

V^* is bounded and measurable. Since the maps $\mu \mapsto \mu(\omega)$ and $\mu \mapsto V^*(\mu)(\mu(\omega) - p(\omega))$ are measurable for all $\omega \in \Omega$, the set $\mathcal{T}_{MD}(p \mid V^*) \coloneqq \{\tau \in \mathcal{T}_{BP}(p) : (V^*, \tau) \text{ satisfies (TT)}\}$ is closed. Theorem 1 of Maccheroni and Marinacci (2001) then implies that $\mathcal{T}_{MD}(p \mid V^*)$ is compact. This set is also convex, Bauer's maximum principle implies that there exists a solution τ' which is an extreme point of $\mathcal{T}_{MD}(p \mid V^*)$. Theorem 2.1 of Winkler (1988) then implies the size of the support of τ' is bounded by the number of linearly independent moment constraints plus one, that is, $|\operatorname{supp}(\tau')| \leq 2(n-1)+1 = 2n-1$.

Finally, fix any measurable selection $V \in \mathbf{V}$ and consider the mediation problem with fixed selection V. We can rewrite the value of the problem using a Lagrange multiplier $g \in \mathbb{R}^n$ on the truth-telling constraint

$$\sup_{\tau \in \mathcal{T}_{BP}(p)} \inf_{g \in \mathbb{R}^n} \int_{\Delta(\Omega)} V(\mu)(1 + \langle g, \mu - p \rangle) \, \mathrm{d}\tau(\mu). \tag{3}$$

Next, define the function $M(\tau, g) \coloneqq \int_{\Delta(\Omega)} (1 + \langle g, \mu - p \rangle) V(\mu) \, \mathrm{d}\tau(\mu)$. Again, we endow $\Delta(\Delta(\Omega))$ with the weak* topology induced by bounded and measurable functions over $\Delta(\Omega)$. The function $M(\tau, g)$ is continuous by definition because $V(\mu)$ is measurable and bounded. In the same topology, the set $\mathcal{T}_{BP}(p)$ is closed because the map $\mu \mapsto \mu(\omega)$ is measurable for all $\omega \in \Omega$. With this, Theorem 1 in Maccheroni and Marinacci (2001) implies that $\mathcal{T}_{BP}(p)$ is compact. Finally, given that $M(\tau, g)$ is affine and continuous, and that both $\mathcal{T}_{BP}(p)$ and $\mathbb{R}^n$ are convex, we can apply Sion's minimax theorem to exchange the sup and inf in (3). Therefore, the value can be rewritten as $\inf_{g \in \mathbb{R}^n} \sup_{\tau \in \mathcal{T}_{BP}(p)} \int V(\mu)(1 + \langle g, \mu - p \rangle) \, \mathrm{d}\tau(\mu) = \inf_{g \in \mathbb{R}^n} \operatorname{cav}(V^g)(p)$, where $V^g(\mu) = V(\mu)(1 + \langle g, \mu - p \rangle)$, and the last equality follows from Kamenica and Gentzkow (2011). Maximizing over all measurable selections, we have the desired representation of the optimal mediation value. ■

A.3 Persuasion vs. Mediation

The following lemma leads to a general version of Theorem 2: mediation is fully interim efficient[39] if and only if cheap talk is fully interim efficient.

Lemma 4. *If $\tau \in \mathcal{T}_{MD}(p)$ is fully interim efficient with selection $V \in \mathbf{V}$ such that $\int V(\mu)(\mu - p) \, \mathrm{d}\tau = \mathbf{0}$, then $\tau \in \mathcal{T}_{CT}(p)$.*

Proof. For every $\omega \in \Omega$, the conditional distribution $\tau^\omega \in \Delta(\Delta(\Omega))$ satisfies the Radon-Nikodym derivative $\frac{\mathrm{d}\tau^\omega}{\mathrm{d}\tau}(\mu) = \frac{\mu(\omega)}{p(\omega)}$, so

[39] See the definition in Section 7, equation 2.

$$\sum_{\omega\in\Omega}\left(\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau^{\omega}(\mu)\right)\lambda(\omega)=\sum_{\omega\in\Omega}\left(\int_{\Delta(\Omega)} V(\mu)\tfrac{\mu(\omega)}{p(\omega)}\,\mathrm{d}\tau(\mu)\right)\lambda(\omega)=\int_{\Delta(\Omega)} V(\mu)\langle\tfrac{\lambda}{p},\mu\rangle\,\mathrm{d}\tau(\mu)$$

Since τ, V solves the optimization problem as in (2), $V=\overline{V}$ almost surely with respect to τ. Otherwise, suppose there exists a measurable set $D\subseteq\Delta(\Omega)$ such that $\tau(D)>0$ and $\overline{V}(\mu)>V(\mu)$ for all $\mu\in D$. Since λ is strictly positive, $\int\overline{V}(\mu)\langle\frac{\lambda}{p},\mu\rangle\,\mathrm{d}\tau(\mu)>\int V(\mu)\langle\frac{\lambda}{p},\mu\rangle\,\mathrm{d}\tau(\mu)$, yielding a contradiction.

By Corollary 1 of Dworczak and Kolotilin (2024), there exists $f\in\mathbb{R}^n$ such that $\overline{V}(\mu)\langle\frac{\lambda}{p},\mu\rangle\leq\langle f,\mu\rangle$ for all $\mu\in\Delta(\Omega)$ and $\overline{V}(\mu)\langle\frac{\lambda}{p},\mu\rangle=\langle f,\mu\rangle$ for all $\mu\in\operatorname{supp}(\tau)$. Since τ satisfies truth-telling with selection $V=\overline{V}$, (iii) of Theorem 1 implies $\operatorname{Cov}_\tau(\overline{V},\langle f,\cdot\rangle)=0$. Let $Z(\mu):=\langle\frac{\lambda}{p},\mu\rangle$ and define $\tilde{\tau}\in\Delta(\Delta(\Omega))$ by the Radon-Nikodym derivative $\frac{\mathrm{d}\tilde{\tau}}{\mathrm{d}\tau}(\mu)=Z(\mu)$. Then, $\operatorname{Cov}_\tau(\overline{V},\langle f,\cdot\rangle)=\operatorname{Cov}_\tau(\overline{V},\overline{V}Z)=\mathbb{E}_\tau[\overline{V}^2Z]-\mathbb{E}_\tau[\overline{V}]\mathbb{E}_\tau[\overline{V}Z]=\mathbb{E}_\tau[\overline{V}^2Z]-\mathbb{E}_\tau[\overline{V}Z]^2=\operatorname{Var}_{\tilde{\tau}}[\overline{V}]$, where the second last equality is by (TT) and the last equality is by the definition of $\tilde{\tau}$. Therefore, $\overline{V}$ is constant over $\operatorname{supp}(\tilde{\tau})=\operatorname{supp}(\tau)$ since $Z(\mu)>0$ for all $\mu\in\Delta(\Omega)$. ■

Proof of Theorem 2. The if direction is immediate. The only if direction, follows from Lemma 4 by observing that if $\tau\in\mathcal{T}_{MD}(p)$ attains the optimal Bayesian persuasion value, then τ is fully interim efficient for $\lambda=p$. ■

Before proving the rest of the results in this section, we need an additional definition.

Definition 7. For every $s\geq\overline{V}(p)$ attainable under cheap talk, the cheap talk hull of s is

$$H(s):=\{\mu\in\Delta(\Omega):\exists\alpha>1\text{ such that }s\in\mathbf{V}_{CT}(\alpha p+(1-\alpha)\mu)\}\tag{4}$$

Observe that $H^*:=H(\mathcal{V}_{CT}(p))$ by Definition 2.

Proof of Proposition 2. We show the following lemma, which implies the desired result.

Lemma 5. *For every $s\geq\overline{V}(p)$ attainable under cheap talk, if there exists $\mu\in H(s)$ such that $\overline{V}_{CT}(\mu)>s$, then there exists $\tau\in\mathcal{T}_{BP}(p)$ such that $\int\overline{V}(\mu)\,\mathrm{d}\tau(\mu)>s$.*

To see this, take any $s\geq\overline{V}(p)$ attainable under cheap talk such that there exists $\hat{\mu}\in H(s)$ with $\overline{V}_{CT}(\hat{\mu})>s$. Hence, there exists $\hat{\tau}\in\mathcal{T}_{CT}(\hat{\mu})$ that attains a higher value than s. As $\hat{\mu}\in H(s)$, there exists $\alpha>1$ and $\tau'\in\mathcal{T}_{CT}(\alpha p+(1-\alpha)\hat{\mu}$ such that τ' attains s. Hence, $\tilde{\tau}=\frac{1}{\alpha}\tau'+\frac{\alpha-1}{\alpha}\hat{\tau}$ is feasible under Bayesian persuasion, as $\frac{1}{\alpha}(\alpha p+(1-\alpha)\hat{\mu})+\frac{\alpha-1}{\alpha}\hat{\mu}=p$. Note that $\int\overline{V}\,\mathrm{d}\tilde{\tau}>s$ since $\int\overline{V}\,\mathrm{d}\hat{\tau}>s$ and $\frac{\alpha-1}{\alpha}>0$.

By Lemma 5, if there exists $\mu\in H^*$ such that $\overline{V}_{CT}(\mu)>\overline{V}_{CT}(p)$, then $\mathcal{V}_{BP}(p)>\mathcal{V}_{CT}(p)$. By Theorem 2, this implies $\mathcal{V}_{BP}(p)>\mathcal{V}_{MD}(p)$. ■

Proof of Corollary 3. The if direction holds by Proposition 2, because there exists $\mu \in \Delta(\Omega) = H^*$ such that $\overline{V}_{CT}(\mu) \geq \overline{V}(\mu) > \overline{V}_{CT}(p)$. For the only if direction, suppose $\overline{V}(\mu) \leq \mathcal{V}_{CT}(p)$ for any $\mu \in \Delta(\Omega)$. As cheap talk attains the global maximum, we have $\mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$. ■

A.4 Mediation and Cheap Talk

Recall that $\mathbf{V}_{CT} : \Delta(\Omega) \rightrightarrows \mathbb{R}$ be the correspondence of sender's payoff under some cheap talk equilibrium with prior $\mu \in \Delta(\Omega)$By Corollary 3 and Section C.2.1 of Lipnowski and Ravid (2020), $\mathbf{V}_{CT}$ is non-empty, convex, and compact-valued.

We say a distribution $\tau \in \mathcal{T}_{CT}(p)$ attains value s (under cheap talk) if $s \in \cap_{\mu \in \text{supp}(\tau)} \mathbf{V}(\mu)$, and a value $s \in \mathbb{R}$ is attainable under cheap talk if there exists $\tau \in \mathcal{T}_{CT}(p)$ that attains it. By Theorem 1 in Lipnowski and Ravid (2020), $s \geq \overline{V}(p)$ is attainable under cheap talk if and only if $p \in \text{co}\left\{\overline{V} \geq s\right\}$.[40]

We start with a useful lemma that extends Theorem 1 in Lipnowski and Ravid (2020).[41]

Lemma 6. *For every $s \in \mathbb{R}$, $\overline{V}_{CT}(p) > s$ if and only if $p \in \text{co}\{\overline{V} > s\}$, and $\underline{V}_{CT}(p) < s$ if and only if $p \in \text{co}\{\underline{V} < s\}$.*

This lemma implies that there exists a cheap talk equilibrium that attains a *strictly* higher (lower) value than s if and only if the prior lies in the convex hull of posteriors with highest (lowest) value strictly above (below) s.

Proof. For any $s \geq \overline{V}(p)$, the first equivalence follows from Theorem 1 of Lipnowski and Ravid (2020). For the only if direction, suppose $\overline{V}_{CT}(p) > s$, then there exists $\tau \in \mathcal{T}_{CT}(p)$ that attains a value $s' > s$. Theorem 1 of Lipnowski and Ravid (2020) implies that $p \in \text{co}\{\overline{V} \geq s'\} \subseteq \text{co}\{\overline{V} > s\}$. For the if direction, suppose $p \in \text{co}\{\overline{V} > s\}$, then there exists finitely many points $\{\mu_i\}_{i=1}^k \subseteq \{\overline{V} > s\}$ such that $p = \sum \alpha_i \mu_i$ for some $\{\alpha_i\}_{i=1}^k \subseteq [0, 1]$, $\sum_{i=1}^k \alpha_i = 1$. Let $\hat{s} := \min_i \overline{V}(\mu_i)$, we have $p \in \text{co}\{\overline{V} \geq \hat{s}\}$. Theorem 1 of Lipnowski and Ravid (2020) then implies that $\overline{V}_{CT}(p) \geq \hat{s} > s$.

For any $s < \overline{V}(p)$, the first equivalence is true as both $\overline{V}_{CT}(p) \geq \overline{V}(p) > s$ and $p \in \text{co}\{\overline{V} > s\}$ are true. The second equivalence follows from a symmetric argument.[42] ■

Proof of Theorem 3. Let $s = \overline{V}_{CT}(p)$ in this proof.

[40] Here, we use the notation $\left\{\overline{V} \geq s\right\} = \left\{\mu \in \Delta(\Omega) : \overline{V}(\mu) \geq s\right\}$.

[41] Theorem 1 of Lipnowski and Ravid (2020) establishes the weak inequality versions of the first equivalence in Lemma 6. We extend this result to strict inequalities.

[42] See footnote 15 of Lipnowski and Ravid (2020).

First Statement: This statement can be shown through an explicit construction. To show this, we consider the auxiliary program (η-MD) as in the proof of Proposition 1. The variable in (η-MD) is a probability measure $\eta \in \Delta(\Delta(\Omega) \times I)$, and we use (μ, r) to denote arbitrary elements in $\Delta(\Omega) \times I$.

Suppose cheap talk is locally improvable at p. By assumption, there exists $\tilde{\mu} \in H^*$ and $\lambda \in (0,1)$ such that $\overline{V}_{CT}(\lambda\tilde{\mu} + (1-\lambda)p) > s > \underline{V}_{CT}(\tilde{\mu})$. Let $\hat{\mu} := \lambda\tilde{\mu} + (1-\lambda)p$. By Lemma 6, there exist $\tau^+ \in \mathcal{T}_{CT}(\hat{\mu})$ that attains value $s + V^+$ for some $V^+ > 0$ and $\tau^- \in \mathcal{T}_{CT}(\tilde{\mu})$ that attains value $s - V^-$ for some $V^- > 0$. For any $r \in I$, define $\phi^r : \Delta(\Omega) \to \Delta(\Omega) \times I$ by $\varphi^r(\mu) = (\mu, r)$. Let $\eta^+ := (\varphi^{s+V^+})_{\#}\tau^+$ and $\eta^- := (\varphi^{s-V^-})_{\#}\tau^-$.

Let $\xi := \frac{\frac{1}{\lambda}V^-}{V^+ + \frac{1}{\lambda}V^-}$. Then,

$$\begin{aligned}\mathbb{E}_{(\xi\eta^+ + (1-\xi)\eta^-)}\left[(r-s)(\mu-p)\right] =& \xi V^+(\hat{\mu}-p) - (1-\xi)V^-(\tilde{\mu}-p) \\ =& \left(\lambda\xi V^+ - (1-\xi)V^-\right)(\tilde{\mu}-p) = \mathbf{0}.\end{aligned}$$

Let $\mu^* := \xi\hat{\mu} + (1-\xi)\tilde{\mu}$. Since H^* is convex, $\mu^* \in H^*$. So there exists $\alpha > 1$ and $\tau' \in \mathcal{T}_{CT}(\alpha p + (1-\alpha)\mu^*)$ such that τ' attains s. Let $\eta' = (\varphi^s)_{\#}\tau'$.

Next, consider $\tilde{\eta} := \frac{1}{\alpha}\eta' + \frac{\alpha-1}{\alpha}\xi\eta^+ + \frac{\alpha-1}{\alpha}(1-\xi)\eta^-$. By construction, $\tilde{\eta}$ satisfies (η-BP) and (η-OB). It also satisfies (η-TT) since

$$\mathbb{E}_{\tilde{\eta}}\left[r(\mu-p)\right] = s\mathbb{E}_{\tilde{\eta}}\left[\mu-p\right] + \tfrac{\alpha-1}{\alpha}\mathbb{E}_{(\xi\eta^+ + (1-\xi)\eta^-)}\left[(r-s)(\mu-p)\right] = \mathbf{0},$$

where the last equality is by (η-BP) and our construction of ξ. With this, we have

$$\mathbb{E}_{\tilde{\eta}}[r] = s + \tfrac{\alpha-1}{\alpha}\xi V^+ - \tfrac{\alpha-1}{\alpha}(1-\xi)V^- = s + (\tfrac{1}{\lambda} - 1)\tfrac{\alpha-1}{\alpha}\tfrac{V^+V^-}{V^+ + \frac{1}{\lambda}V^-} > s. \tag{5}$$

Finally, take $\tilde{\tau} = \text{marg}_{\Delta(\Omega)}\,\tilde{\eta}$ and $\tilde{V}(\mu) = \mathbb{E}_{\tilde{\eta}}[r|\mu]$. $(\tilde{V}, \tilde{\tau})$ is implementable under mediation and attains exactly the same value as $\tilde{\eta}$, which is higher than s.

Second Statement: By definition, cheap talk is improvable at p if and only if there exists $\mu \in \{\underline{V}_{CT} < s\}$ such that

$$\{\overline{V}_{CT} > s\} \cap [p, \mu) \neq \emptyset,$$

where $[p, \mu)$ denote the line segment connecting p and μ, including the end point p while excluding μ.[43] Let $D_+ := \{\overline{V} > s\}$ and $D_- := \{\underline{V} < s\}$. By Lemma 6, $\{\overline{V}_{CT} > s\} = \text{co}\, D_+$ and $\{\underline{V}_{CT} < s\} = \text{co}\, D_-$. Suppose s is not improvable at p, then for any $\mu \in \text{co}\, D_-$,

[43] Recall that $p \notin \{\overline{V}_{CT} > s\}$ since $s = \overline{V}_{CT}(p)$.

$\text{co}\, D_+ \cap [p, \mu) = \emptyset$. Therefore,

$$\text{co}\, D_+ \bigcap \left(\bigcup_{\mu \in \text{co}\, D_-} [p, \mu) \right) = \emptyset.^{44} \tag{6}$$

For any affine set $M \subseteq \mathbb{R}^n$, we say that M is *orthogonal* to s if for every $(\tau, V) \in \mathcal{T}_{MD}(p) \times \mathbf{V}$ satisfying (TT), $\tau(\{V = s\} \cup M) = 1$. The second statement of Theorem 3 then follows from the following lemma.

Lemma 7. *Suppose* (6) *holds and that there exists an affine set* $M \subseteq \mathbb{R}^n$ *such that* $p \in M$ *and such that* M *is orthogonal to* s. *Then either* $\mathcal{V}_{MD}(p) \leq s$ *or there is an affine set* $M' \subseteq \mathbb{R}^n$ *such that* $\dim M' = \dim M - 1$, $p \in M'$, *and such that* M' *is orthogonal to* s.

With this lemma, we may start from an initial affine set M_0 to be the affine hull of $\Delta(\Omega)$. Note that $p \in M_0$ and M_0 is orthogonal to s. The claim implies either $\mathcal{V}_{MD}(p) \leq s$, which is the desired property, or that there exists an $n-2$-dimensional affine set M_1 such that $p \in M_1$ and such that M_1 is orthogonal to s. Repeat this algorithm, and it terminates either when the desired property $\mathcal{V}_{MD}(p) \leq s$ holds, or when we reach a 0-dimensional affine set $M_{n-1} = \{p\}$. In the latter case, by orthogonality, $V(\mu) \leq s$ τ-almost surely for any $\tau \in \mathcal{T}_{MD}(p)$ and $V \in \mathbf{V}$ that (τ, V) satisfies (TT), so $\int V \, d\tau \leq s$. By assumption, s is attainable under cheap talk, so we have $\mathcal{V}_{MD}(p) = s$.

Now we prove the lemma. Suppose $D_+ = \emptyset$, then the claim is trivially true since $\mathcal{V}_{MD}(p) \leq s$ holds. Suppose $D_+ \neq \emptyset$. We next show that (6) implies that there exists a $g \in \mathbb{R}^n$ such that $\langle g, \mu \rangle \leq 0$ for all $\mu \in S_+ \coloneqq \text{co}(D_+ \cap M)$ and $\langle g, \mu \rangle \geq 0$ for all $\mu \in S_- \coloneqq \bigcup_{\mu \in \text{co}(D_- \cap M)} [p, \mu)$.

To see this, first observe that S_- is convex. If $\text{co}(D_- \cap M) = \emptyset$, then $S_- = \{p\}$. If $\text{co}(D_- \cap M) \neq \emptyset$, take any $\mu_1 = \alpha_1 \hat{\mu}_1 + (1 - \alpha_1)p$, $\mu_2 = \alpha_2 \hat{\mu}_2 + (1 - \alpha_2)p$ for some $\hat{\mu}_1, \hat{\mu}_2 \in \text{co}(D_- \cap M)$ and $\alpha_1, \alpha_2 \in (0, 1)$. For any $\lambda \in (0, 1)$, $\lambda \mu_1 + (1 - \lambda)\mu_2 = (\lambda \alpha_1 + (1 - \lambda)\alpha_2) \left(\frac{\lambda \alpha_1}{\lambda \alpha_1 + (1-\lambda)\alpha_2} \hat{\mu}_1 + \frac{(1-\lambda)\alpha_2}{\lambda \alpha_1 + (1-\lambda)\alpha_2} \hat{\mu}_2 \right) + (\lambda(1 - \alpha_1) + (1 - \lambda)(1 - \alpha_2))p$, where $\frac{\lambda \alpha_1}{\lambda \alpha_1 + (1-\lambda)\alpha_2} \hat{\mu}_1 + \frac{(1-\lambda)\alpha_2}{\lambda \alpha_1 + (1-\lambda)\alpha_2} \hat{\mu}_2 \in \text{co}(D_- \cap M)$.[45]

Since S_+ and S_- are disjoint nonempty convex sets, Theorem 11.3 of Rockafellar (1970) then implies there exists a hyperplane in $\mathbb{R}^{n-1}$ separating S_+ and S_- properly. That is, there exists $\hat{g} \in \mathbb{R}^n$ such that $\langle \hat{g}, \mu \rangle \geq c \geq \langle \hat{g}, \mu' \rangle$ for all $\mu \in S_-$, $\mu' \in S_+$ for some $c \in \mathbb{R}$, and hyperplane $\{\mu \in \mathbb{R}^n : \langle \mu, \hat{g} \rangle = c\}$ does not contain both sets. Take $g = \hat{g} - \mathbf{c} \in \mathbb{R}^n$,[46] we have the desired hyperplane $H \coloneqq \{\mu \in \mathbb{R}^n : \langle \mu, g \rangle = 0\}$ that separates S_+ and S_- properly.

[44] We use the convention that $\cup_{\mu \in S} [p, \mu) = \{p\}$ if $S = \emptyset$.

[45] When $\frac{\lambda \alpha_1}{\lambda \alpha_1 + (1-\lambda)\alpha_2} \hat{\mu}_1 + \frac{(1-\lambda)\alpha_2}{\lambda \alpha_1 + (1-\lambda)\alpha_2} \hat{\mu}_2 = p$, it follows that $\lambda \mu_1 + (1 - \lambda)\mu_2 = p \in S_-$.

[46] Here, $\mathbf{c} = (c, \ldots, c) \in \mathbb{R}^n$.

Note that $\mathrm{co}(D_- \cap M) \subseteq \mathrm{cl}\, S_-$, so $D_- \cap M$ is contained in the same closed half-space determined by H as S_-. This implies that $(V(\mu) - s)\langle g, \mu\rangle \leq 0$ for all $\mu \in \Delta(\Omega) \cap M$ and $V \in \mathbf{V}$. For any $\tau \in \mathcal{T}_{MD}(p)$ and $V \in \mathbf{V}$ such that (τ, V) satisfies (TT), since M is orthogonal to s, $(V(\mu) - s)\langle g, \mu\rangle \leq 0$ τ-almost surely, and thereby

$$0 \geq \int_{\Delta(\Omega)} (V(\mu) - s)\langle g, \mu\rangle \, \mathrm{d}\tau(\mu) = \left(\int_{\Delta(\Omega)} V(\mu) \, \mathrm{d}\tau(\mu) - s \right) \langle g, p\rangle, \tag{7}$$

where the last equality is by (zeroCov) and (BP).

By construction $p \in S_-$, so $\langle g, p\rangle \geq 0$. If $\langle g, p\rangle > 0$, (7) implies that $\int V \, \mathrm{d}\tau \leq s$ for any $\tau \in \mathcal{T}_{MD}(p)$ and $V \in \mathbf{V}$ such that (τ, V) satisfies (TT), so $\mathcal{V}_{MD}(p) \leq s$. If $\langle g, p\rangle = 0$, we show that $H \cap M$ is an affine set of dimension $\dim M - 1$ which is orthogonal to s. Note that H does not contain M as it separates S_+ and S_- properly, and $H \cap M$ is non-empty because it contains p. Therefore, $H \cap M$ is an affine set of dimension $\dim M - 1$. Since $\langle g, p\rangle = 0$, (7) implies that for every $\tau \in \mathcal{T}_{MD}(p)$ and $V \in \mathbf{V}$ such that (τ, V) satisfies (TT), $(V(\mu) - s)\langle g, \mu\rangle = 0$ τ-almost surely. This means that for every implementable (V, τ) under mediation, with probability one, either $V(\mu) = s$ or $\langle g, \mu\rangle = 0$; that is, $\tau(\{V = s\} \cup H) = 1$. Since the intersection of two probability-one events also has probability one, $\tau(\{V = s\} \cup (H \cap M)) = 1$, so $H \cap M$ is orthogonal to s, which establishes the lemma. ∎

Proof of Corollary 4. Since the full-dimensionality condition holds at p, Theorem 3 implies that $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$ if and only if cheap talk is improvable at p.

Suppose cheap talk is improvable at p, then there exists $\mu \in \Delta(\Omega)$ such that $\overline{V}_{CT}(\lambda\mu + (1-\lambda)p) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\mu)$ for some $\lambda \in (0, 1)$. By assumption, $V(p) < \overline{V}_{CT}(p)$, so $\lambda\mu + (1-\lambda)p \in \mathrm{co}\{V < \overline{V}_{CT}(p)\} = \{\underline{V}_{CT} < \overline{V}_{CT}(p)\}$ by Lemma 6. Therefore, $\overline{V}_{CT}(\lambda\mu + (1-\lambda)p) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\lambda\mu + (1-\lambda)p)$.

Suppose there exists $\mu \in \Delta(\Omega)$ such that $\overline{V}_{CT}(\mu) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\mu)$, then $\mu \in \{\overline{V}_{CT} > \overline{V}_{CT}(p)\} = \mathrm{co}\{V > \overline{V}_{CT}(p)\}$ by Lemma 6. Since V is continuous, $\{V > \overline{V}_{CT}(p)\}$ is open and so is its convex hull. Moreover, we have $\mu \neq p$ because $\overline{V}_{CT}(\mu) > \overline{V}_{CT}(p)$. Thus, there is $\lambda \in (0, 1)$ such that $\overline{V}_{CT}(\lambda\mu + (1-\lambda)p) > \overline{V}_{CT}(p)$, so cheap talk is improvable at p. ∎

Proof of Theorem 3'. We show cheap talk is improvable at p if and only if there exists $\ell \in \mathcal{E}$ with $p \in \ell$ such that $\mathbf{V}_{CT|\ell} - \overline{V}_{CT}(p)$ is not quasi single-crossing at $\ell(p)$, the result then follows from Theorem 3.

Suppose cheap talk is improvable at p, by definition, there exists $\hat{\mu} \in \Delta(\Omega)$ and $\tilde{\mu} \in (p, \hat{\mu})$ such that $\overline{V}_{CT}(\tilde{\mu}) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\hat{\mu})$. Now consider the line segment $\ell \in \mathcal{E}$ crossing $p, \hat{\mu}$, we show $\mathbf{V}_{CT|\ell} - \overline{V}_{CT}(p)$ is not quasi single-crossing at $\ell(p)$. Without loss, we may assume $\ell(\hat{\mu}) > \ell(\tilde{\mu}) > \ell(p)$. By definition, $\overline{V}_{CT}(\tilde{\mu}) > \overline{V}_{CT}(p)$ implies that $\mathbf{V}_{CT|\ell} - \overline{V}_{CT}(p)$ is not quasi single-crossing at $\ell(p)$ from above, and $\overline{V}_{CT}(\tilde{\mu}) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\hat{\mu})$ implies that it is also not quasi single-crossing at $\ell(p)$ from below.

Conversely, suppose there exists $\ell \in \mathcal{E}$ with $p \in \ell$ such that $\mathbf{V}_{CT|\ell} - \overline{V}_{CT}(p)$ is not quasi single-crossing at $\ell(p)$. Then there must exist $\tilde{\mu} \in \ell$ such that $\overline{V}_{CT}(\tilde{\mu}) > \overline{V}_{CT}(p)$. Without loss, we consider $\ell(\tilde{\mu}) > \ell(p)$. Then by definition, $\mathbf{V}_{CT|\ell} - \overline{V}_{CT}(p)$ is not quasi single-crossing at $\ell(p)$ from above. By construction of $\mathbf{V}_{CT}$, $\overline{V}_{CT}(\mu) \leq \overline{V}_{CT}(p)$ for all $\mu \in \ell$ with $\ell(\mu) \leq \ell(p)$. Since $\mathbf{V}_{CT|\ell} - \overline{V}_{CT}(p)$ is also not quasi single-crossing at $\ell(p)$ from below, by definition, there must be $\hat{\mu}, \hat{\mu}' \in \ell$ with $\ell(\hat{\mu}') > \ell(\hat{\mu}) > \ell(p)$ such that $\overline{V}_{CT}(\hat{\mu}) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\hat{\mu}')$, which implies cheap talk is improvable at p. ∎

A.5 Moment Mediation: Quasiconvex Utility

Proof of Theorem 4. By Corollary 6 of Lipnowski and Ravid (2020), when T is multi-dimensional and v strictly quasiconvex, no disclosure is suboptimal under cheap talk. Suppose the full-dimensionality condition holds at p, by Corollary 3, $\mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p)$ if and only if $\{V > \mathcal{V}_{CT}(p)\} = \emptyset$, which means cheap talk attains the global maximum value. This leads to the dichotomy in the theorem statement: If $\max V = \mathcal{V}_{CT}(p)$, then (1) holds trivially. It suffices to show $\max V > \mathcal{V}_{CT}(p)$ implies (2).

Note that if $\mathcal{V}_{BP}(p) = \max V$, it must be the case that $V(\mu) = \max V$ for all μ in the support of any optimal $\tau \in \mathcal{T}_{BP}(p)$, which implies $\mathcal{V}_{BP}(p) = \mathcal{V}_{CT}(p)$, contradiction. Hence, what remains to show is that $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$. By Corollary 4 and Lemma 6, $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ if and only if $\operatorname{co} D_+ \cap \operatorname{co} D_- = \emptyset$, where $D_+ = \{V > \mathcal{V}_{CT}(p)\}$ and $D_- = \{V < \mathcal{V}_{CT}(p)\}$. We next show that under strict quasiconvexity and $\max V > \mathcal{V}_{CT}(p)$, the intersection is always non-empty, hence $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$.

Let $\bar{D}_+ = \{x \in X : v(x) > \mathcal{V}_{CT}(p)\}$ and $\bar{D}_- = \{x \in X : v(x) < \mathcal{V}_{CT}(p)\}$, both are open by continuity of v. We first show that $\operatorname{co} \bar{D}_+ \cap \operatorname{co} \bar{D}_- \neq \emptyset$. Since $\max V > \mathcal{V}_{CT}(p)$, we have $\bar{D}_+ \neq \emptyset$. Take any open ball in $\bar{D}_+$, there exist two points x_1, x_2 in this open ball such that x_1, x_2, and $T(p)$ are not colinear. Note that by strict quasiconvexity, we have $T(p) \in \bar{D}_-$. Moreover, there exists a unique $\lambda_i \in (0,1)$ such that $v(\lambda_i x_i + (1-\lambda_i)T(p)) = \mathcal{V}_{CT}(p)$ for $i = 1, 2$ since v is continuous and strictly quasiconvex. Here, existence follows by the intermediate value theorem, whereas strict quasiconvexity implies uniqueness. By strict quasiconvexity, $\frac{1}{2}(\lambda_1 x_1 + \lambda_2 x_2) + (1 - \frac{1}{2}(\lambda_1 + \lambda_2))T(p) \in \bar{D}_-$. Since $\bar{D}_-$ is open, there exists $\varepsilon > 0$ such that $\frac{1}{2}((\lambda_1 + \varepsilon)x_1 + (\lambda_2 + \varepsilon)x_2) + (1 - \frac{1}{2}(\lambda_1 + \lambda_2 + 2\varepsilon))T(p) \in \bar{D}_-$. Note that $x_i' = (\lambda_i + \varepsilon)x_i + (1 - \lambda_i - \varepsilon)T(p) \in \bar{D}_+$, and we have $\frac{1}{2}x_1' + \frac{1}{2}x_2' \in \bar{D}_-$, so $\operatorname{co} \bar{D}_+ \cap \operatorname{co} \bar{D}_- \neq \emptyset$.

Finally, take any $\mu_i \in \Delta(\Omega)$ such that $T(\mu_i) = x_i'$ for $i = 1, 2$, we have $\mu_i \in D_+$. Since $T(\frac{1}{2}\mu_1 + \frac{1}{2}\mu_2) = \frac{1}{2}x_1' + \frac{1}{2}x_2'$, $\frac{1}{2}\mu_1 + \frac{1}{2}\mu_2 \in D_-$, the claim holds. ∎

Proof of Proposition 3. Since v is minimally edge non-monotone, there exists a state $\underline{w} \in \operatorname{argmin}_{\omega\in\Omega} V(\delta_\omega)$ such that for any $\omega \in \Omega \setminus \{\underline{\omega}\}$, $f_\omega(\lambda) := V(\lambda\delta_\omega + (1-\lambda)\delta_{\underline{\omega}})$ is not weakly increasing in $\lambda \in [0,1]$.

We show that f_ω is strictly quasiconvex on $[0,1]$. Note that for any $\lambda \neq \lambda' \in [0,1]$

$$\begin{aligned} f_\omega(\alpha\lambda + (1-\alpha)\lambda') &= v(\alpha T(\mu) + (1-\alpha)T(\mu')) \\ &\leq \max\{v(T(\mu)), v(T(\mu'))\} = \max\{f_\omega(\lambda), f_\omega(\lambda')\}, \end{aligned}$$

where $\mu = \lambda\delta_\omega + (1-\lambda)\delta_{\underline{\omega}}$, $\mu' = \lambda'\delta_\omega + (1-\lambda')\delta_{\underline{\omega}}$. The first equality is by definition and linearity of T, the inequality is by (strict) quasiconvexity of v, and the last equality is by definition. Moreover, the inequality is strict if and only if $T(\mu) \neq T(\mu')$. Suppose $T(\mu) = T(\mu')$, then by linearity of T, $T(\delta_\omega) = T(\delta_{\underline{\omega}})$, which means f_ω is a constant on $[0,1]$. This contradicts with the assumption that f_ω is non-monotone, hence $T(\mu) \neq T(\mu')$ and f_ω is strictly quasiconvex.

As f_ω is strictly quasiconvex and non-monotone, there must be a unique $\lambda_\omega \in (0,1]$ such that $f_\omega(\lambda_\omega) = f_\omega(0)$. Suppose $f_\omega(\lambda) > f_\omega(0)$ for all $\lambda > 0$, then there exists $\lambda_2 > \lambda_1 > 0$ such that $f_\omega(\lambda_1) > f_\omega(\lambda_2) > f_\omega(0)$ (otherwise f_ω is weakly increasing). But $\lambda_1 \in (0, \lambda_2)$, so $f_\omega(\lambda_1) > f_\omega(\lambda_2) > f_\omega(0)$ violates the strict quasiconvexity, contradiction. So there must be a $\hat{\lambda}_\omega \in (0,1]$ such that $f_\omega(\hat{\lambda}_\omega) \leq f_\omega(0)$. By continuity of v, there exists $\lambda_\omega \in [\hat{\lambda}_\omega, 1]$ such that $f_\omega(\lambda_\omega) = f_\omega(0)$. The uniqueness is by strict quasiconvexity.

The argument above holds for any $\omega \in \Omega \setminus \{\underline{\omega}\}$. Let $\mu_\omega := \lambda_\omega\delta_\omega + (1-\lambda_\omega)\delta_{\underline{\omega}}$, we have $V(\mu_\omega) = V(\delta_{\underline{\omega}})$ for any $\omega \in \Omega \setminus \{\underline{\omega}\}$. Set $\tilde{\Delta} := \operatorname{co}\{\delta_{\underline{\omega}}, \{\mu_\omega : \omega \in \Omega \setminus \{\underline{\omega}\}\}\}$. This is an $(n-1)$–simplex as $\{\delta_{\underline{\omega}}, \{\mu_\omega : \omega \in \Omega \setminus \{\underline{\omega}\}\}\}$ is affinely independent with n points. Moreover, for any $p \in \operatorname{int} \tilde{\Delta}$, there is $\tau \in \mathcal{T}_{CT}(p)$ that supports on $\{\delta_{\underline{\omega}}, \{\mu_\omega : \omega \in \Omega \setminus \{\underline{\omega}\}\}\}$ that attains $V(\delta_{\underline{\omega}})$. Since $v(\cdot)$ is strictly quasiconvex, the composition $V = v \circ T$ is quasiconvex, hence $V(\mu) \leq V(\delta_{\underline{\omega}})$ for any $\mu \in \tilde{\Delta}$. This shows that $\{V > V(\delta_{\underline{\omega}})\}$ is contained in the convex set $\Delta(\Omega) \setminus \tilde{\Delta}$, by Lemma 6, $\mathcal{V}_{CT}(p) \leq V(\delta_{\underline{\omega}})$ for any $p \in \tilde{\Delta}$. Therefore, the full-dimensionality condition holds for all priors $p \in \operatorname{int} \tilde{\Delta}$. Moreover, if $V(\delta_{\underline{\omega}}) < \max_{\mu\in\Delta(\Omega)} V(\mu)$, then for any $p \in \operatorname{int} \tilde{\Delta}$, $\mathcal{V}_{CT}(p) < \max V$. As the full-dimensionality condition holds, Theorem 4 shows that $\max V > \mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p) > V(p)$. ∎

Proof of Proposition 4. Fix an acceptance game such that R is strictly quasiconvex, minimally edge non-monotone given T, and satisfies $\min_{x\in X_T} R(x) < \max_{x\in X} R(x)$. In addition, assume that G is strictly increasing and log-concave. It follows that $v(x) = G(R(x))$ is strictly quasiconvex and minimally edge non-monotone given T. By Proposition 3, there exists an $(n-1)$–simplex $\tilde{\Delta} \subseteq \Delta(\Omega)$ such that, for all $p \in \operatorname{int} \tilde{\Delta}$ there exists $\tau^* \in \mathcal{T}_{MD}(p)$ with $\int v(T(\mu))\, d\tau^*(\mu) > \int v(T(\mu))\, d\tau(\mu)$ for all $\tau \in \mathcal{T}_{CT}(p)$. Fix p and τ^* as above. Next, we show that $\phi(z) \coloneqq H(G^{-1}(z))$ is convex. First, observe that

$$\phi(z) = H(G^{-1}(z)) = \int_0^1 \max\left\{G^{-1}(t), G^{-1}(z)\right\} \mathrm{d}t = G^{-1}(z)z + \int_z^1 G^{-1}(t)dt,$$

hence that

$$\phi'(z) = z(G^{-1})'(z) = \frac{z}{g(G^{-1}(z))} = \frac{G(G^{-1}(z))}{g(G^{-1}(z))} = \frac{G}{g} \circ G^{-1}(z),$$

where g is the density of G. By log-concavity of G, $\frac{G}{g}$ is increasing. It follows that ϕ' is the composition of two increasing functions, hence it is an increasing function. This shows that ϕ is convex. Finally, for all $\tau \in \mathcal{T}_{CT}(p)$, we have

$$\begin{aligned}\int v_R(T(\mu))\,\mathrm{d}\tau^*(\mu) &= \int \phi(v(T(\mu)))\,\mathrm{d}\tau^*(\mu) \geq \phi\left(\int v(T(\mu))\,\mathrm{d}\tau^*(\mu)\right)\\ &> \phi\left(\int v(T(\mu))\,\mathrm{d}\tau(\mu)\right) = \int \phi(v(T(\mu)))\,\mathrm{d}\tau(\mu) = \int v_R(T(\mu))\,\mathrm{d}\tau(\mu),\end{aligned}$$

where the first inequality follows by Jansen inequality, the second inequality follows by the fact that $\phi(z)$ is convex and strictly increasing, and the second equality from the fact that $V(\mu) = v(T(\mu))$ must be constant over the support of τ since $\tau \in \mathcal{T}_{CT}(p)$. ■

References

ALIPRANTIS, C. D. AND K. C. BORDER (2006): *Infinite Dimensional Analysis: a Hitchhiker's Guide*, Berlin; London: Springer.

ARIELI, I., Y. BABICHENKO, R. SMORODINSKY, AND T. YAMASHITA (2023): "Optimal persuasion via bi-pooling," *Theoretical Economics*, 18, 15–36.

ATHEY, S. (1998): "Comparative statics under uncertainty: single crossing properties and log supermodularity," Working Paper.

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.

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.

CORRAO, R. (2023): "Mediation Markets: The Case of Soft Information," Working Paper.

CRAWFORD, V. P. AND J. SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.

Doval, L. and V. Skreta (2024): "Constrained information design," *Mathematics of Operations Research*, 49, 78–106.

Doval, L. and A. Smolin (2024): "Persuasion and Welfare," *Journal of Political Economy*, 132, 2451–2487.

Dworczak, P. and A. Kolotilin (2024): "The Persuasion Duality," *Theoretical Economics (forthcoming)*, working Paper.

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

Forges, F. (1985): "Correlated equilibria in a class of repeated games with incomplete information," *International Journal of Game Theory*, 14, 129–149.

——— (1986): "An approach to communication equilibria," *Econometrica: Journal of the Econometric Society*, 1375–1385.

——— (2020): "Games with incomplete information: from repetition to cheap talk and persuasion," *Annals of Economics and Statistics*, 3–30.

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.

Habu, T., E. Lipnowski, and D. Ravid (2021): "Knowing the Informed Player's Payoffs and Simple Play in Repeated Games," *Available at SSRN 3837519*.

Hart, S. (1985): "Nonzero-sum two-person repeated games with incomplete information," *Mathematics of Operations Research*, 10, 117–153.

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

Koessler, F. and V. Skreta (2023): "Informed information design," *Journal of Political Economy*, 131, 3186–3232.

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

Kolotilin, A., R. Corrao, and A. Wolitzky (2024): "Persuasion and Matching: Optimal Productive Transport," *Journal of Political Economy (forthcoming)*.

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

Lin, X. and C. Liu (2024): "Credible persuasion," *Journal of Political Economy*, 132, 2228–2273.

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.

Maccheroni, F. and M. Marinacci (2001): "A Heine-Borel Theorem for ba (Sigma)," *Rivista internazionale di scienze economiche e commerciali*, 48, 353–362.

Mathevet, L., J. Perego, and I. Taneva (2020): "On information design in games," *Journal of Political Economy*, 128, 1370–1404.

Milgrom, P. and C. Shannon (1994): "Monotone comparative statics," *Econometrica: Journal of the Econometric Society*, 157–180.

Myerson, R. B. (1982): "Optimal coordination mechanisms in generalized principal–agent problems," *Journal of Mathematical Economics*, 10, 67–81.

——— (1997): *Game theory: analysis of conflict*, Harvard university press.

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.

Soltan, V. (2019): *Lectures on convex sets*, World Scientific.

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.

Supplement to "The Bounds of Mediated Communication"

B Omitted Proofs of Technical Lemmas

Proof of Lemma 2. (1) is shown in Lipnowski and Ravid (2020), and (2) can be shown using a similar argument. Since $\bar{x}$ is the infimum, there exists a weakly decreasing sequence $\{x_n^-\}_{n=1}^\infty \subseteq [\bar{x}, 1]$ that converges to $\bar{x}$ and $\min F(x_n^-) \leq 0$ for all $n = 1, 2, \ldots$. Take a strictly increasing sequence $\{x_n^+\}_{n=1}^\infty \subseteq [0, \bar{x}]$ that converges to $\bar{x}$ (and constant 0 sequence if $\bar{x} = 0$). By definition of $\bar{x}$, we have $\max F(x_n^+) \geq 0 \geq \min F(x_n^-)$ for all $n = 1, 2, \ldots$.

Taking subsequence if necessary, $\{\min F(x_n^-)\}_{n=1}^\infty$ converges to $y \leq 0$. By upper hemi-continuity of F, $y \in F(\bar{x})$, hence $\min F(\bar{x}) \leq 0$. A similar argument shows that $0 \leq \max F(\bar{x})$. As F is convex-valued, $0 \in F(\bar{x})$. ■

Proof of Lemma 3. The only if direction follows from the law of iterated expectations and the definition of Obedience. For every measureable $\tilde{a} : A \to \Delta(A)$, we have

$$\begin{aligned}\int u_R(\omega, a)\,\mathrm{d}\pi(\omega, a) &= \int_A \mathbb{E}[u_R(\omega, a) \mid \pi^a]\,\mathrm{d}\,\mathrm{marg}_A\,\pi(a)\\ &\geq \int_A \mathbb{E}[u_R(\omega, \tilde{a}) \mid \pi^a]\,\mathrm{d}\,\mathrm{marg}_A\,\pi(a) = \int u_R(\omega, \tilde{a})\,\mathrm{d}\pi(\omega, a).\end{aligned}$$

For the if direction, suppose Obedience is not satisfied, then there exists a measurable $S \subseteq A$ with $\mathrm{marg}_A\,\pi(S) > 0$ such that $\mathbb{E}[u_R(\omega, a) \mid \pi^a] < \max_{a' \in A} \mathbb{E}[u_R(\omega, a') \mid \pi^a]$ for all $a \in S$. By the measurable maximum theorem (Aliprantis and Border, 2006, Theorem 18.19), there exists a measurable $\hat{a} : A \to A$ with $\hat{a}(a) \in \mathrm{argmax}_{a' \in A}\,\mathbb{E}[u_R(\omega, a') \mid \pi^a]$. We then have $\int u_R(\omega, a)\,\mathrm{d}\pi(\omega, a) < \int u_R(\omega, \hat{a})\,\mathrm{d}\pi(\omega, a)$, contradiction. ■

C Single-crossing and the Binary-state Case

C.1 Quasi Single-crossing Functions

In this appendix, we define the notion of quasi single-crossing function and relate it to the standard notion of single-crossing function as well as Shannon (1995)'s notion of weakly single-crossing function.

In Section 5.2, we defined the notion of quasi single-crossing correspondence. When a correspondence is singleton valued, we obtain the corresponding definition for functions. A function $U : [0,1] \to \mathbb{R}$ is *quasi single-crossing at x_0 from below (above)* if for all $x \leq (\geq) x_0$, $U(x) \leq U(x_0) = 0$ and for all $x' > x > x_0$ $(x' < x < x_0)$, $U(x) > 0$ implies $U(x') \geq 0$, and U is *quasi single-crossing at x_0* if it is quasi single-crossing at x_0 either from below or from above.

Next, following Athey (1998), we provide a local version of the standard notion of single-crossing single-crossing functions in Milgrom and Shannon (1994). We say that U is *single-crossing* at $x_0 \in [0,1]$ if U is single-crossing and $U(x_0) = 0$. Finally, we recall the notion of weakly single-crossing function in Shannon (1995). We say that U is *weakly single crossing from below (above)* if for any $x' > x$ $(x' < x)$, $U(x) > 0$ implies $U(x') \geq 0$, and U is *weakly single-crossing at x_0* if it is weakly single-crossing at x_0 either from below or from above.

Lemma 8. *Fix a function $U : [0,1] \to \mathbb{R}$ and $x_0 \in [0,1]$. The following facts hold:*

1. *If U is single-crossing at x_0, then U is quasi single-crossing at x_0;*

2. *If U is quasi single-crossing at x_0, then U is weakly single-crossing.*

Proof. Suppose U is single-crossing at x_0 from below, by definition, $U(x) \geq (>) 0$ implies $U(x') \geq (>) 0$ for any $x' > x$. Since $U(x_0) = 0$, $U(x) \leq 0$ for any $x \leq x_0$. Therefore, U is also quasi single-crossing at x_0 from below. This establishes statement 1.

Suppose U is quasi single-crossing at x_0 from below, then for any $x' > x$, if $x > x_0$, then $U(x) > 0$ implies $U(x') \geq 0$ by definition; if $x \leq x_0$, then $U(x) \leq 0$, which means $U(x) > 0$ implies $U(x') \geq 0$ is true automatically. Therefore, U is also weakly single-crossing from below. This establishes statement 2. ∎

C.2 Binary-state Case

When the state space Ω is binary, our main results can be dramatically simplified. As Ω is binary and $\Delta(\Omega)$ is 1-dimensional, with a slight abuse of notation, we use μ to denote the first entry of the receiver's posterior belief.

Proposition 5. *The following are equivalent:*

(i) $\mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p)$*;*

(ii) $\mathcal{V}_{BP}(p) = \mathcal{V}_{CT}(p)$*;*

(iii) $p \in \mathrm{co}(\mathrm{argmax}\, \overline{V})$ *or* $\overline{V}$ *is superdifferentiable at* p*.*

Proof of Proposition 5. The equivalence between (i) and (ii) is immediate from Theorem 2 (see the proof in Appendix A.3). We now show the equivalence between (ii) and (iii). The if direction is immediate. For the only if direction, suppose that $\mathcal{V}_{BP}(p) = \mathcal{V}_{CT}(p)$. Take any optimal $\tau^* \in \mathcal{T}_{CT}(p)$ with finite support and a selection $V \in \mathbf{V}$ such that $V(\mu) = \mathcal{V}_{BP}(p)$ τ^*–almost surely. Moreover, $V = \overline{V}$ τ^*–almost surely, otherwise persuasion would attain a strictly higher value. By Corollary 1 of Dworczak and Kolotilin (2024), there exists $f \in \mathbb{R}^2$ such that $\overline{V}(\mu) \leq \langle f, \mu \rangle$ for all $\mu \in \Delta(\Omega)$ and $\overline{V}(\mu) = \langle f, \mu \rangle$ for all $\mu \in \text{supp}(\tau^*)$. When τ^* is non-degenerate, $f = (\mathcal{V}_{BP}(p), \mathcal{V}_{BP}(p))$, hence $\mathcal{V}_{BP}(p) \geq \overline{V}(\mu)$ on $\Delta(\Omega)$. This means that $\mathcal{V}_{BP}(p) = \mathcal{V}_{CT}(p)$ is the maximum value of $\overline{V}$. Then $p \in \text{co}(\text{supp}(\tau^*)) \subseteq \text{co}(\text{argmax}\, \overline{V})$. If τ^* is degenerate, then $\overline{V}(p) = \langle f, p \rangle$ and $\overline{V}(\mu) \leq \langle f, \mu \rangle$ for all $\mu \in \Delta(\Omega)$, which means $\overline{V}$ is superdifferentiable at p. ∎

Corollary 5 is an immediate consequence of Theorem 3'. When the sender's payoff is uniquely defined given the receiver's posterior and we strengthen the weak single-crossing condition of Corollary 5 to the standard single-crossing condition, the equivalence between mediation and chap talk is much stronger as we show next.

Corollary 6. *Assume that $\mathbf{V} = V$ is singleton-valued. If $V(\mu) - \overline{V}_{CT}(p)$ is single-crossing at $\mu = p$, then $\mathcal{T}_{MD}(p) = \mathcal{T}_{CT}(p)$ and all cheap talk equilibria attain the same value for the sender.*[47] *In this case, no disclosure is optimal for mediation.*

Proof of Corollary 6. Since $V(\mu) - \mathcal{V}_{CT}(p)$ is single-crossing at p, $\mathcal{V}_{CT}(p) = V(p)$ and $[V(\mu) - \overline{V}_{CT}(p)](\mu - p)$ is non-negative/non-positive for any $\mu \in \Delta(\Omega)$. Therefore, the shifted truth-telling constraint for the mediation problem $\int_{\Delta(\Omega)} [V(\mu) - \mathcal{V}_{CT}(p)](\mu - p)\, \mathrm{d}\tau(\mu) = 0$ implies that $V(\mu) = \mathcal{V}_{CT}(p)$ for any $\mu \in \text{supp}(\tau)$, hence $\mathcal{T}_{MD}(p) = \mathcal{T}_{CT}(p)$. As no disclosure is optimal under cheap talk, no disclosure is also optimal under mediation. ∎

The assumptions of Corollary 6 hold whenever $\mathbf{V} = V$ is monotone. Therefore, countervailing incentives (i.e., V non-monotone) are necessary for mediation to strictly outperform cheap talk with binary states.

Corollaries 5 and 6 imply that cheap talk and mediation attain the same sender-optimal value for several canonical shapes of the sender's payoff.

Corollary 7. *Assume that $\mathbf{V} = V$ is singleton-valued. If V is concave or quasiconvex, then $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ for all $p \in (0,1)$. There exists a non-monotone quasiconcave V and $p \in (0,1)$ such that $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$.*

[47] A function $U : \mathbb{R} \to \mathbb{R}$ is single-crossing at $\hat{x}$ if U is single-crossing and $U(\hat{x}) = 0$.

Proof of Corollary 7. The claim is straightforward when V is concave. If V is quasiconvex, then either 0 or 1 attains its maximum value. Without loss of generality, assume $V(0) \leq V(1)$, and let $\tilde{p} := \sup\{\mu \in [0,1] : V(\mu) = V(0)\}$. By continuity of V, $V(\tilde{p}) = V(0)$. For every $\mu \in [0, \tilde{p}]$, we have $V(\mu) \leq V(0)$ by quasiconvexity, while $V(\mu) > V(0)$ for every $\mu \in (\tilde{p}, 1]$ by the definition of $\tilde{p}$.

For every prior $p \in (0, \tilde{p}]$, we have $\mathcal{V}_{CT}(p) = V(0)$. The argument above shows that $\{\mu \in [0,1] : V(\mu) < V(0)\} \subseteq [0, \tilde{p}]$ and $\{\mu \in [0,1] : V(\mu) > V(0)\} \subseteq (\tilde{p}, 1]$, so cheap talk is not improvable at p. By Theorem 3, $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ for every $p \in (0, \tilde{p})$.

For every prior $p \in (\tilde{p}, 1)$, we have $V(p) > V(0)$. Quasiconvexity of V implies that $V(\mu) \geq V(p)$ for every $\mu > p$. Otherwise, if there exists $\hat{\mu} > p$ with $V(\hat{\mu}) < V(p)$, then $V(p) > \max\{V(0), V(\hat{\mu})\}$, contradicting quasiconvexity. A similar argument shows that $V(\mu) \leq V(p)$ for every $\mu < p$, hence, $\mathcal{V}_{CT}(p) = V(p)$. As $\{V > V(p)\} \subseteq (p, 1]$ and $\{V < V(p)\} \subseteq [0, p)$, cheap talk is not improvable at p, so $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ by Theorem 3.

Finally, consider $V(\mu) = 0$ for $\mu \in [0, 1/2)$ and $V(\mu) = -(\mu - 1/2)(\mu - 3/4)$ for $\mu \in [1/2, 1]$. This V is non-monotone and quasiconcave. At any $p \in (0, 1/2)$, cheap talk is improvable and the full-dimensionality condition holds at p. By Theorem 3, $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$. ■

In Sections 4 and 5, we provided generalizations of these results to settings with an arbitrary number of states. When the belief space is multidimensional, we restrict consideration to one-dimensional line segments to use one-dimensional notions such as the weak single-crossing or single-crossing properties to characterize when mediation is strictly valuable. These properties are also relevant when the sender's payoff depends on a one-dimensional statistic of the receiver's posterior (see Appendix F.2 below).

D Moment-measurable Illustrations

D.1 Financial Intermediation under Mean-Variance Preferences

In this example, the issuer's payoff function is $v(x) = R(x) = \gamma x_1^2 + x_1 - \gamma x_2$ for some $\gamma > 0$. This is convex but not strictly quasiconvex in x, so we cannot conclude as in Section 6 that no disclosure is always suboptimal under cheap talk. However, we can show this explicitly for every $p \in \tilde{\Delta}$ as constructed in Example 3. Let $t := \sum_{j=1}^{n-1} \frac{p(\omega_j)}{\alpha_j}$ and $\hat{\mu}_i := \alpha_i t \delta_{\omega_i} + (1 - \alpha_i t)\delta_0$ for all $i = 1, \ldots, n-1$. Observe that $p = \sum_{i=1}^{n-1} \frac{p(\omega_i)}{\alpha_i t} \hat{\mu}_i$, and since $p \in \tilde{\Delta}$, $t \leq 1$, as otherwise none of $\hat{\mu}_i$ lies in the line segment $[\delta_0, \mu_i]$, which implies that $p \in \Delta(\Omega) \setminus \tilde{\Delta}$, a contradiction. Hence, $\hat{\mu}_i \in [\delta_0, \mu_i]$ for every i. By the convexity of v, $V = v \circ T$ is also convex, so $V(p) \leq \sum \frac{p(\omega_i)}{\alpha_i t} V(\hat{\mu}_i)$. We have shown in the main text that for every $i = 1, \ldots, n-1$, V is strictly convex along the edge of the simplex connecting

δ_0 and δ_{ω_i}. Recall that $V(\mu_i) = V(\delta_0) = 0$ for every i, which implies that $V(\hat{\mu}_i) < 0$ by strict convexity of V along the segment $[\delta_0, \mu_i]$, so $V(p) < 0$. This shows that there exists a distribution of posteriors feasible under cheap talk that secures a payoff to the sender that is strictly higher than that under no disclosure.

We now show that for any $p \in \operatorname{int} \tilde{\Delta}$, $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p) > V(p)$. Since $\mathcal{V}_{CT}(p) = 0 < V(\delta_{\omega_{n-1}})$, cheap talk does not attain the global maximum value, which implies $\mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p)$ by Proposition 2. By Corollary 4 and Lemma 6, $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ if and only if $\operatorname{co} D_+ \cap \operatorname{co} D_- = \emptyset$, where $D_+ = \{\mu \in \Delta(\Omega) : V(\mu) > 0\}$ and $D_- = \{\mu \in \Delta(\Omega) : V(\mu) < 0\}$.

As in the proof of Theorem 4, we consider the upper and lower contour sets of v at value $\mathcal{V}_{CT}(p) = 0$, that is, $\bar{D}_+ = \{x \in X : x_1^2 + x_1/\gamma > x_2\}$ and $\bar{D}_- = \{x \in X : x_1^2 + x_1/\gamma < x_2\}$, both are open by continuity of v. Since $\max V > \mathcal{V}_{CT}(p)$, we have $\bar{D}_+ \neq \emptyset$. Take any open ball in $\bar{D}_+$, there exist two points x, x' in this open ball such that x, x' and $T(p)$ are not colinear. Since $V(p) < 0$, we have $T(p) \in \bar{D}_-$. Moreover, there exists a unique $\lambda \in (0,1)$ such that $v(\lambda x + (1-\lambda)T(p)) = 0$. Here, existence follows from the continuity of v, and uniqueness comes from the fact that any line segment with two endpoints in $\bar{D}_+$ and $\bar{D}_-$ can cross the set $\{x \in X : x_1^2 + x_1/\gamma = x_2\}$ at most once. Similarly, there exists a unique $\lambda' \in (0,1)$ such that $v(\lambda' x' + (1-\lambda')T(p)) = 0$.

Note that $\{x \in X : x_1^2 + x_1/\gamma \leq x_2\}$ is strictly convex, so $\frac{1}{2}(\lambda x + \lambda' x') + (1 - \frac{1}{2}(\lambda + \lambda'))T(p) \in \bar{D}_-$. Since $\bar{D}_-$ is open, there exists $\varepsilon > 0$ such that $\frac{1}{2}((\lambda+\varepsilon)x + (\lambda'+\varepsilon)x') + (1 - \frac{1}{2}(\lambda + \lambda' + 2\varepsilon))T(p) \in \bar{D}_-$. Note that $\hat{x} = (\lambda+\varepsilon)x + (1-\lambda-\varepsilon)T(p) \in \bar{D}_+$ and $\hat{x}' = (\lambda'+\varepsilon)x' + (1-\lambda'-\varepsilon)T(p) \in \bar{D}_+$, and we have $\frac{1}{2}\hat{x} + \frac{1}{2}\hat{x}' \in \bar{D}_-$, so $\operatorname{co} \bar{D}_+ \cap \operatorname{co} \bar{D}_- \neq \emptyset$.

Finally, take any $\mu, \mu' \in \Delta(\Omega)$ such that $T(\mu) = \hat{x}$ and $T(\mu') = \hat{x}'$, we have $\mu, \mu' \in D_+$ and $\frac{1}{2}\mu + \frac{1}{2}\mu' \in D_-$, so $\operatorname{co} D_+ \cap \operatorname{co} D_- \neq \emptyset$ and $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$.

D.2 Salesman with Reputation Concerns

We extend our illustration in the introduction to multidimensional states and revisit the salesman example in Chakraborty and Harbaugh (2010) and Lipnowski and Ravid (2020). A seller is trying to convince a buyer to purchase a good with multiple features $\omega \in \Omega \subseteq \mathbb{R}_+^k$ and assume that $\mathbf{0} \in \Omega$. The buyer is uncertain about ω, and their payoff from purchasing this good only depends on the posterior mean on the quality of these features $T(\mu) = \mathbb{E}_\mu(\omega) \in \mathbb{R}^k$. In particular, we assume that Ω is a finite set such that T is full-rank. In the main text, this assumption is implied by the fact that $\Omega = \{0,1\}^k$ with $k > 1$. In general, recall that $X = T(\Delta(\Omega))$ and that in this case $X_T = \{T(\delta_\omega) \in X : \omega \in \Omega\} = \Omega$.

The buyer's payoff with posterior mean x is $R(x)$ for some function $R : \mathbb{R}^k \to \mathbb{R}$ that is continuously differentiable, strictly convex, and strictly increasing with $R(\mathbf{0}) = 0$.[48] The buyer has an outside option with value $\varepsilon \in \mathbb{R}$ with distribution G that has a continuous density, is strictly increasing and strictly convex, and such that $R(X) \subseteq \operatorname{supp} \varepsilon$. Therefore, the buyer purchases the good if and only if $R(x) \geq \varepsilon$.

The expected revenue for the seller when x is the realized vector of conditional expectations is $G(R(x))$. The seller has also reputation concerns, that is, the overall seller's expected payoff with posterior mean x is $v(x) = G(R(x)) - \langle \rho, x \rangle$, where $\rho \in \mathbb{R}^k_{++}$ measures the seller's reputation concern. Our key assumption on the seller's payoff is

$$G(R(x)) > \langle \rho, x \rangle > \langle G'(0)\nabla R(\mathbf{0}), x \rangle \qquad \forall x \in \Omega \setminus \{\mathbf{0}\}, \tag{8}$$

where $\nabla R(\mathbf{0})$ is the gradient of R at $\mathbf{0}$. This implies that the seller's payoff when the buyer is sure that the state is $\mathbf{0}$ is strictly lower than any other degenerate buyer's belief, that is, $G(R(x)) - \langle \rho, x \rangle > 0$ for all $x \in \Omega \setminus \{\mathbf{0}\}$. In general, this assumption captures the fact that the reputation concerns of the seller are mild.

By assumption, the composition $G \circ R$ is strictly convex, hence the seller's payoff v is strictly convex. We show that the seller's payoff $v(x)$ is minimally edge non-monotone given T. Fix any $x \in \Omega \setminus \{\mathbf{0}\}$. It suffices to check that $\phi(\alpha) \coloneqq v(\alpha x)$ is non-monotone in $\alpha \in [0, 1]$. The derivative of ϕ is $\phi'(\alpha) = G'(R(\alpha x))\langle \nabla R(\alpha x), x \rangle - \langle \rho, x \rangle$. By assumption 8, we have

$$\phi'(0) = \langle G'(0)\nabla R(\mathbf{0}), x \rangle - \langle \rho, x \rangle < 0$$

and

$$\phi(1) = G(R(x)) - \langle \rho, x \rangle > 0 = G(R(\mathbf{0})) - \langle \rho, \mathbf{0} \rangle = \phi(0).$$

Because ϕ' is continuous, it follows that ϕ is non-monotone.

By Proposition 3, there exists an $(n-1)$–simplex $\tilde{\Delta} \subseteq \Delta(\Omega)$ where the full-dimensionality condition holds in its interior. This simplex can be explicitly constructed. For all $x \in \Omega \setminus \{\mathbf{0}\}$, let $\alpha_x \in (0, 1)$ denote the unique solution of $v(\alpha x) = 0$ and define $\mu_x = \alpha_x \delta_x$. With this,

$$\tilde{\Delta} \coloneqq \operatorname{co}\{\delta_{\mathbf{0}}, \{\mu_x : x \in \Omega \setminus \{\mathbf{0}\}\}\}$$

is the desired simplex. Proposition 3 also implies that the seller strictly benefits from hiring a mediator when the prior is in $\operatorname{int} \tilde{\Delta}$. Moreover, since the seller's payoff at state $\mathbf{0}$ is strictly lower than other states, the dichotomy in Theorem 4 implies that the seller attains an even higher payoff under Bayesian persuasion than mediation at priors in $\operatorname{int} \tilde{\Delta}$.

[48] Strictly increasing in the sense that $R(x) < R(x')$ for all $x < x'$ componentwise.

E Non-existence of Dual Solution

In this section, we present a binary-state example where the dual problem of optimal mediation does not have a solution.

Assume that $\mathbf{V} = V$ is singleton-valued. The dual problem of mediation is to find two Lagrange multipliers $f, g \in \mathbb{R}^n$ that solve the following minimization problem:

$$\begin{aligned} & \inf_{f,g\in\mathbb{R}^n} \langle f, p\rangle \\ \text{subject to:} & \\ & \langle f, \mu\rangle \geq (1 + \langle g, \mu - p\rangle)V(\mu) \qquad \forall \mu \in \Delta(\Omega). \end{aligned} \tag{D}$$

We now exhibit a binary-state example where the minimum in (D) is not attained. Suppose the sender has preference $V(\mu) = 4\mu(\mu - 1/2) + 1/4$. When the common prior $p = 1/2$, the corresponding dual problem of mediation does not have a solution. To see this, note that the dual problem can be written as

$$\begin{aligned} & \inf_{f_0,f_1,g\in\mathbb{R}} \tfrac{1}{2}f_1 + f_0 \\ \text{subject to: } & f_1\mu + f_0 \geq (1 + g(\mu - \tfrac{1}{2}))(\tfrac{1}{4} + 4\mu(\mu - \tfrac{1}{2})). \end{aligned}$$

Let $V^g(\mu) := (1 + g(\mu - \frac{1}{2}))(\frac{1}{4} + 4\mu(\mu - \frac{1}{2}))$. Note that when $g < 0$, the lowest line above V^g is a tangent line of V^g at $\mu^* = \frac{1}{2} - \frac{1}{2g}$ that passes through $(0, V^g(0))$. That is, $f_1 = \frac{g}{4} - \frac{1}{g}$ and $f_0 = \frac{1}{4}(1 - \frac{g}{2}) = V^g(0)$. Then the value $f_1/2 + f_0 = \frac{1}{4} - \frac{1}{2g} \downarrow \frac{1}{4}$ as $g \to -\infty$. Also observe that $g \geq 0$ is never an optimal solution of the dual, since $(V^g(0) + V^g(1))/2 = \frac{5}{4} + \frac{g}{2} > \frac{5}{4}$. Therefore, the infimum value of this dual problem cannot attained by any $f_1, f_0, g \in \mathbb{R}$.

F Mean-measurable Mediation

F.1 Implementation

In this subsection, we consider a special case of the setting of Section 6 where the moment function leads to the receiver's posterior mean. We focus on Euclidean state spaces $\Omega \subseteq \mathbb{R}^k$ for some $k \geq 1$ and moment function $T(\mu) = \mathbb{E}_\mu(\omega)$. Let $X := T(\Delta(\Omega)) \subseteq \mathbb{R}^k$ be the set of all possible posterior means. Assume the sender's payoff only depends on the receiver's posterior mean, i.e., $V(\mu) = v(T(\mu))$ for some continuous $v : \mathbb{R}^k \to \mathbb{R}$.

Differently from Section 6, here we do not focus on distributions over posteriors $\tau \in \Delta(\Delta(\Omega))$, but rather on the induced distributions of posterior means $q \in \Delta(X)$. We say

$q \in \Delta(X)$ is implementable under mediation if there exists $\tau \in \mathcal{T}_{MD}(p)$ that induces q.

In this setting, we can adapt Theorem 1 as follows. For any $q \in \Delta(X)$ and $v : X \to \mathbb{R}$, define the corresponding distorted distribution $q^v \in \Delta(X)$ by

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

Proposition 6. *Let $V(\mu) = v(T(\mu))$. The following are equivalent:*

(i) *$q \in \Delta(X)$ is implementable under mediation;*

(ii) *There exists a dilation*[49] *$\mathbf{D} : X \to \Delta(X)$ such that $\mathbf{D}q = \mathbf{D}q^v = p$;*

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

Note that when there is no truth-telling constraint, by Strassen's Theorem,[50] 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^v are mean-preserving contractions of p.

Proof. We first show that (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, $q = (T)_\# \tau$. We construct a dilation $\mathbf{D} : X \to \Delta(X)$ by $\mathbf{D}_x = \mathbb{E}_\tau[\mu|T(\mu) = x]$. By construction we have $x = \int y\,\mathrm{d}\mathbf{D}_x(y)$ for all x and $\int \mathbf{D}_x\,\mathrm{d}q(x) = \int \mu\,\mathrm{d}\tau(\mu) = p$. Note that $\int \mathbf{D}_x v(x)\,\mathrm{d}q(x) = \int V(\mu)\mu\,\mathrm{d}\tau = p\int V\,\mathrm{d}\tau = p\int v\,\mathrm{d}q$, where the first and third equalities are obtained by iterated expectation and $V(\mu) = v(T(\mu))$, and the second by truth-telling. Hence, the dilation constructed satisfies $\mathbf{D}q = \mathbf{D}q^v = p$.

Conversely, suppose there exists a dilation $\mathbf{D}$ such that $\mathbf{D}q = \mathbf{D}q^v = p$. Then let $\tau \in \Delta(\Delta(\Omega))$ be the pushforward measure of q under dilation $\mathbf{D}$, that is, $\tau(R) = q(\mathbf{D}^{-1}(R))$ for all measurable $R \subseteq \Delta(\Omega)$. By change of variable, we obtain $\int \mu\,\mathrm{d}\tau = \int \mathbf{D}_x\,\mathrm{d}q = p$ and

$$\begin{aligned}\int V(\mu)\mu\,\mathrm{d}\tau &= \int v(x)\mathbf{D}_x\,\mathrm{d}q(x)\\ &= p \cdot \int v(x)\,\mathrm{d}q(x) = p\int V(\mu)\,\mathrm{d}\tau(\mu)\end{aligned}$$

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

[50]Let X be a compact convex metrizable space and p, q are Borel probability measures on X. Strassen's Theorem states that $q \preceq_{cvx} p$ if and only if there exists a dilation $\mathbf{D}$ such that $p = \mathbf{D}q$, 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).

where the first and third equalities follow by a change of variable, and the second one follows by $\mathbf{D}q^v = p$. Overall, this implies that $\tau \in \mathcal{T}_{MD}(p)$.

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

F.2 One-dimensional Mean

In this subsection, we consider another special case of the setting of the previous subsection: the one where the mean function is one-dimensional. Formally, assume that $\Omega \subset \mathbb{R}$ and that $T(\mu) = \mathbb{E}_\mu[\omega]$. That is, the state is one-dimensional, and the sender's value function depends on the receiver's conditional expectation only: $V(\mu) = v(\mathbb{E}_\mu[\omega])$. This is the most studied case in the Bayesian persuasion literature.

Let $\bar{v}$ ($\underline{v}$) denote the quasiconcave (quasiconvex) envelope of v and let $\mathbf{v}(x) = [\underline{v}(x), \bar{v}(x)]$ for every $x \in X$. In general, the quasiconcave envelope of v evaluated at the prior mean can be strictly larger than the actual optimal cheap talk value, that is, we can have $\bar{v}(T(p)) > \mathcal{V}_{CT}(p)$. However, $\mathbf{v}(x)$ is still helpful in studying the value comparison between cheap talk and mediation.

The binary state case is a special case of a one-dimensional mean, and we show that many intuitions from Corollary 5 extend. Unlike the binary case, the full-dimensionality condition may not hold even if no disclosure is suboptimal under cheap talk. In the next proposition, we provide a sufficient condition on the prior p such that a weak single-crossing condition in $\mathbf{v}(x)$ characterizes the comparison between mediation and cheap talk.

Proposition 7. *Suppose $V(\mu) = v(T(\mu))$ for some continuous v on $\mathbb{R}$.*

(1) *If $v(T(p)) = \bar{v}(T(p))$, then no disclosure is optimal under cheap talk. In this case, $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ if and only if no disclosure is optimal for mediation.*

(2) *If $v(T(p)) < \bar{v}(T(p))$ and $p \in \operatorname{int}\operatorname{co}\{V = \bar{v}(T(p))\}$, then $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ if and only if $\mathbf{v}(x) - \bar{v}(T(p))$ is quasi single-crossing at $T(p)$.*

The first statement says that v is equal to its quasiconcave envelope at $T(p)$, then the only way that mediation is not strictly valuable is when no disclosure is optimal. When there is a wedge at $T(p)$ between v and its quasiconcave envelope and the full-dimensionality condition holds, then, similarly to the binary-state case, mediation is worthless if and only if the sender's shifted utility function is quasi single-crossing at $T(p)$. Here, full dimensionality is implied by the condition $p \in \operatorname{int}\operatorname{co}\{V = \bar{v}(T(p))\}$, which also implies that $\mathcal{V}_{CT}(p) = \bar{v}(T(p))$.

Before proving Proposition 7, we introduce the relaxed mediation problem and state and prove a useful lemma. First, by Proposition 6, if $q \in \Delta(X)$ is implementable under mediation then

$$\int_X v(x)(x - T(p)) \,\mathrm{d}q(x) = \int_X v(x)x \,\mathrm{d}q(x) - \left(\int_X v(x) \,\mathrm{d}q(x)\right)\left(\int_X x \,\mathrm{d}q(x)\right)$$
$$= \operatorname{Cov}_{\pi_q}(v, \omega) = 0$$

where $\pi_q \in \Delta(\Omega \times X)$ satisfies the conditions in (iii) of Proposition 6. Second, we use this observation to define the relaxed mediation problem as:

$$\sup_{q \in \Delta(X)} \int_X v(x) \,\mathrm{d}q(x) \tag{9}$$
$$\text{subject to: } \int_X x \,\mathrm{d}q(x) = T(p) \tag{10}$$
$$\int_X v(x)(x - T(p)) \,\mathrm{d}q(x) = 0. \tag{11}$$

The first constraint relaxes (BP) by only requiring consistency with the prior mean as opposed to the entire prior distributions. The second constraint relaxes (zeroCov) as explained above.

Similarly, we can relax the cheap talk problem analyzed in the main text by replacing the zero-variance condition $\operatorname{Var}_\tau(V) = 0$ with a weaker zero-variance condition involving only the distribution of conditional expectations: $\operatorname{Var}_q(v) = 0$. Therefore, the relaxed cheap talk problem is defined as in (9) by replacing the second constraint with the latter zero-variance condition.

Lemma 9. *The following statements are true:*

(1) $\mathcal{V}_{CT}(p) \leq \bar{v}(T(p))$.

(2) *If* $p \in \operatorname{int}\operatorname{co}\{\mu : v(T(\mu)) = \bar{v}(T(p))\}$, *then* $\mathcal{V}_{CT}(p) = \bar{v}(T(p))$ *and the full-dimensionality condition holds at* p.

Proof. (1): Note that $\bar{v}(T(p))$ is the value of the relaxed cheap talk problem. For any $\tau \in \mathcal{T}_{CT}(p)$, $q^\tau = (T)_\# \tau \in \Delta(X)$ is feasible in the relaxed cheap talk problem. As $\int_{\Delta(\Omega)} V(\mu) \,\mathrm{d}\tau(\mu) = \int_X v(x) \,\mathrm{d}q^\tau$, we have $\mathcal{V}_{CT}(p) \leq \bar{v}(T(p))$.

(2): Suppose $p \in \operatorname{int} \operatorname{co}\{\mu : v(T(\mu)) = \bar{v}(T(p))\}$, then there exists an open neighborhood N of p such that $\bar{v}(T(p))$ can be attained under cheap talk under any prior $p' \in N$. By (i), $\mathcal{V}_{CT}(p) = \bar{v}(T(p))$. By Lemma 1, the full-dimensionality condition holds at p. ■

Proof of Proposition 7. (1) is clear by (1) of Lemma 9.

For (2), as $p \in \operatorname{int} \operatorname{co}\{V = \bar{v}(T(p))\}$, (2) of Lemma 9 implies $\mathcal{V}_{CT}(p) = \bar{v}(T(p))$ and the full-dimensionality condition holds at p. So $v(T(p)) < \bar{v}(T(p))$ implies that no disclosure is suboptimal under cheap talk. By Corollary 4 and Lemma 6, $\mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p)$ if and only if $\operatorname{co} D_+ \cap \operatorname{co} D_- = \emptyset$, where $D_+ = \{V > \bar{v}(T(p))\}$ and $D_- = \{V < \bar{v}(T(p))\}$.

Using a similar argument as in the proof of Corollary 4, we can show that $\mathbf{v}(x) - \bar{v}(T(p))$ is quasi single-crossing at $T(p)$ if and only if $\operatorname{co} \bar{D}_+ \cap \operatorname{co} \bar{D}_- = \emptyset$, where $\bar{D}_+ = \{v > \bar{v}(T(p))\}$ and $\bar{D}_- = \{v < \bar{v}(T(p))\}$ are sets in X. It suffices to show $\operatorname{co} \bar{D}_+ \cap \operatorname{co} \bar{D}_- = \emptyset$ if and only if $\operatorname{co} D_+ \cap \operatorname{co} D_- = \emptyset$.

By continuity of v, $\operatorname{co} \bar{D}_+$ and $\operatorname{co} \bar{D}_-$ are open convex subsets of $X \subseteq \mathbb{R}$, which are either empty or open intervals. If any of $\operatorname{co} \bar{D}_+$ and $\operatorname{co} \bar{D}_-$ is empty, then the claim holds trivially, so we focus on the case when both convex hulls are non-empty.

Suppose $\operatorname{co} \bar{D}_+ \cap \operatorname{co} \bar{D}_- = \emptyset$, then there exists $\hat{x} \in X$ that separates $\operatorname{co} \bar{D}_+$ and $\operatorname{co} \bar{D}_-$. Without loss, assume $\sup \operatorname{co} \bar{D}_- \leq \hat{x} \leq \inf \operatorname{co} \bar{D}_+$, and by openness $\bar{D}_- \subseteq \{x < \hat{x}\}$, $\bar{D}_+ \subseteq \{x > \hat{x}\}$. Then for any $\mu \in D_-$, $V(\mu) = v(T(\mu)) < \bar{v}(T(p))$, hence we have $T(\mu) < \hat{x}$. Similarly, any $\mu \in D_+$ is contained in the positive half-space determined by $\{\mu \in \Delta(\Omega) : T(\mu) = \hat{x}\}$. Therefore, $\operatorname{co} D_+$ and $\operatorname{co} D_-$ are strictly separated by the hyperplane $\{\mu \in \Delta(\Omega) : T(\mu) = \hat{x}\}$ and has no intersection.

Suppose $\operatorname{co} \bar{D}_+ \cap \operatorname{co} \bar{D}_- \neq \emptyset$. Then either $\operatorname{co} \bar{D}_+ \cap \bar{D}_- \neq \emptyset$ or $\bar{D}_+ \cap \operatorname{co} \bar{D}_- \neq \emptyset$.[51] Without loss, suppose the former is true. Then there exists $\hat{x} \in \bar{D}_-$ and $\{x_i\}_{i=1}^k \subseteq \bar{D}_+$ such that $\hat{x} = \sum \alpha_i x_i$ for some $\alpha_i \in (0, 1)$, $\sum_i \alpha_i = 1$. Since $X = T(\Delta(\Omega))$, there exists $\mu_i \in \Delta(\Omega)$ such that $T(\mu_i) = x_i$ for all $i = 1, \ldots, k$, hence $\mu_i \in D_+$. Note that $\sum_i \alpha_i \mu_i \in \Delta(\Omega)$ and $T(\sum_i \alpha_i \mu_i) = \hat{x}$, which means $\sum_i \alpha_i \mu_i \in D_-$. Therefore, $\operatorname{co} D_+ \cap \operatorname{co} D_- \neq \emptyset$. ■

Next, we derive a sufficient condition on $\mathbf{v}(x)$ such that there exists a non-trivial set of priors $p \in \Delta(\Omega)$ where the full-dimensionality assumption in Proposition 7 is satisfied.

[51] If there exists $\{x_i\}_{i=1}^k \subseteq \bar{D}_+$ and $\{y_j\}_{j=1}^m \subseteq \bar{D}_-$ with $\sum \alpha_i x_i = \sum \beta_j y_j$ for some $\alpha_i, \beta_j \in (0, 1)$ and $\sum_i \alpha_i = \sum_j \beta_j = 1$. Without loss, assume the points are ordered by indices. Suppose $y_j \notin \operatorname{co}\{x_i\}_{i=1}^k = [x_1, x_k]$ for all $j = 1, \ldots, m$. Then there must be some $y_{j_1} < x_1$ and $y_{j_2} > x_k$, which means $[x_1, x_k]$ is contained in $\operatorname{co}\{y_j\}_{j=1}^m$. It follows that $\operatorname{co} \bar{D}_- \cap \bar{D}_+ \neq \emptyset$.

Proposition 8. *If there exists $\hat{x} \in X$ such that $\bar{v}(\hat{x}) > v(\hat{x})$ and $\mathbf{v}(x) - \bar{v}(\hat{x})$ is not quasi single-crossing at $\hat{x}$, then the set*

$$\Delta(\hat{x}) := \{\mu \in \Delta(\Omega) : T(\mu) = \hat{x}\} \cap \operatorname{int} \operatorname{co}\{\mu \in \Delta(\Omega) : v(T(\mu)) = \bar{v}(\hat{x})\}$$

is nonempty and, for all $p \in \Delta(\hat{x})$, we have $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$.

Proof. We first show $\Delta(\hat{x}) \neq \emptyset$. Note that X is a closed interval in $\mathbb{R}$. Let $\underline{x} = \min X = T(\delta_{\underline{\omega}})$, $\bar{x} = \max X = T(\delta_{\bar{\omega}})$ for some $\underline{\omega}, \bar{\omega} \in \Omega$. Since $\bar{v}(\hat{x}) > v(\hat{x})$, there exists $x_1 < \hat{x} < x_2$ in X such that $v(x_1) = v(x_2) = \bar{v}(\hat{x})$. Moreover, since $\mathbf{v}(x) - \bar{v}(\hat{x})$ is not quasi single-crossing at $\hat{x}$, there exists $x' \neq \hat{x} \in X$ such that $v(x') > \bar{v}(\hat{x})$. By continuity, there exists x in $\operatorname{int} \operatorname{co}\{\hat{x}, x'\}$ with $v(x) = \bar{v}(\hat{x})$. So it is without loss to assume at least one of x_1, x_2 is in the interior of X.

If $x_1 > \underline{x}$, then the hyperplane $H_1 := \{\tilde{\mu} \in \mathbb{R}^n : T(\tilde{\mu}) = x_1\}$ either intersects the interior of $\Delta(\Omega)$ or contains the line segment $\operatorname{co}\{\delta_{\underline{\omega}}, \delta_{\bar{\omega}}\}$. To see this, observe that H_1 contains a point in the relative interior of $\operatorname{co}\{\delta_{\underline{\omega}}, \delta_{\bar{\omega}}\}$ by linearity of T and $x_1 > \underline{x}$. With this, there are two cases. If H_1 contains $\operatorname{co}\{\delta_{\underline{\omega}}, \delta_{\bar{\omega}}\}$ then the claim at the beginning of this paragraph trivially follows. If instead H_1 does not contain the line segment $\operatorname{co}\{\delta_{\underline{\omega}}, \delta_{\bar{\omega}}\}$, Theorem 3.44 of Soltan (2019) implies that H_1 cuts $\operatorname{co}\{\delta_{\underline{\omega}}, \delta_{\bar{\omega}}\}$, that is, the line segment $\operatorname{co}\{\delta_{\underline{\omega}}, \delta_{\bar{\omega}}\}$ intersects both open halfspaces of $\mathbb{R}^n$ determined by H_1, proving the claim also in this case.

Next, observe that it is not possible for H_1 to contain $\operatorname{co}\{\delta_{\underline{\omega}}, \delta_{\bar{\omega}}\}$ as it implies $X = \{x_1\}$ is a singleton, yielding a contradiction. So H_1 intersects the interior of $\Delta(\Omega)$ and $H_1 \cap \Delta(\Omega) = \{\mu \in \Delta(\Omega) : T(\mu) = x_1\}$ has dimension $n-2$ by Corollary 3.45 of Soltan (2019). Hence, there exist $n-2$ affinely independent points $\mu_1, \ldots, \mu_{n-2}$ in $\{\mu \in \Delta(\Omega) : T(\mu) = x_1\}$, paired with any point $\mu_0 \in \{\mu \in \Delta(\Omega) : T(\mu) = x_2\}$, we have an $(n-1)$–simplex that has non-empty intersection with $\{\mu \in \Delta(\Omega) : T(\mu) = \hat{x}\}$. As $x_1 < \hat{x} < x_2$, a similar argument shows that $\{\mu \in \Delta(\Omega) : T(\mu) = \hat{x}\}$ intersects the interior of this $(n-1)$–simplex, hence $\Delta(\hat{x}) \neq \emptyset$. Similarly, if $x_2 < \bar{x}$, we also have $\Delta(\hat{x}) \neq \emptyset$. Proposition 7 then implies that for any prior $p \in \Delta(\hat{x})$, $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$. ■

Similar to Corollary 6, we can derive simple sufficient conditions such that no disclosure is the only implementable outcome under both cheap talk and mediation.

Corollary 8. *If $v(x) - \bar{v}(T(p))$ is single-crossing at $T(p)$, then $\mathcal{T}_{MD}(p) = \mathcal{T}_{CT}(p)$ and all cheap talk equilibria are optimal. Hence, no disclosure is optimal for mediation.*

In particular, for any monotone v, $v(x) - \bar{v}(T(p))$ is single-crossing at $T(p)$. So non-monotonicity on v is necessary for mediation to outperform cheap talk strictly.

Proof. Since $v(x) - \bar{v}(T(p))$ is single-crossing at $T(p)$, $v(T(p)) = \bar{v}(T(p))$ and $[v(x) - \bar{v}(T(p))](x - T(p))$ is non-negative/non-positive for any $x \in X$. Therefore, the shifted truth-telling constraint $\int [v(x) - \bar{v}(T(p))](x - T(p))\,\mathrm{d}q(x) = 0$ for the relaxed mediation problem in (9) implies that $v(x) = \bar{v}(T(p))$ for any $x \in \mathrm{supp}(q)$ and any feasible $q \in \Delta(X)$ under the relaxed mediation problem. Note that for any implementable $\tau \in \mathcal{T}_{MD}(p)$ in the mediation problem, $q^\tau = (T)_\# \tau$ is feasible in the relaxed mediation problem in (9), which means $V(\mu) = \bar{v}(T(p))$ for any $\mu \in \mathrm{supp}(\tau)$, hence $\tau \in \mathcal{T}_{CT}(p)$. As no disclosure is optimal under cheap talk and $\mathcal{T}_{MD}(p) = \mathcal{T}_{CT}(p)$, no disclosure is also optimal under mediation. ■

G Additional Examples

G.1 Mediation's Trilemma

Recall the mediation trilemma that the following three properties cannot hold at the same time: (1) Information is public; (2) The payoff of the sender is state-independent; (3) Mediation is fully interim efficient and strictly better than cheap talk. In this subsection, we provide examples where (3) holds when we relax one of (1) and (2).

An example without transparent motives where (1) and (3) holds: Consider a binary state space $\Omega = \{0, 1\}$ and the prior on $\omega = 1$ is $p = 1/2$. The sender's indirect utility is state-dependent and singleton-valued $V(\mu, \omega) = G(\mu) - \frac{\omega}{\mu}$, where

$$G(\mu) = \begin{cases} 4\mu & \text{if } \mu \in [0, 1/4) \\ -2\mu + 3/2 & \text{if } \mu \in [1/4, 1/2) \\ 2\mu - 1/2 & \text{if } \mu \in [1/2, 3/4) \\ -4\mu + 4 & \text{if } \mu \in [3/4, 1] \end{cases}$$

We show that $\tilde{\tau} = \frac{1}{2}\delta_{1/4} + \frac{1}{2}\delta_{3/4}$ is feasible under mediation and is fully interim efficient for p, and cheap talk is strictly worse than mediation.[52]

By definition (2), $\tilde{\tau}$ is fully interim efficient with respect to p if it solves

$$\max_{\tau \in \mathcal{T}_{BP}(p)} p \int_0^1 V(\mu, 1)\,\mathrm{d}\tau^1(\mu) + (1-p) \int_0^1 V(\mu, 0)\,\mathrm{d}\tau^0(\mu).$$

Bayes-plausibility implies that the objective function becomes $\int_0^1 G(\mu)\,\mathrm{d}\tau - 1$, hence $\tilde{\tau}$ is the

[52]The mediation problem and the definition of fully interim efficiency can be extended to the state-dependent case. Details available upon request.

unique solution of this maximization problem because it is supported on the global maximum of G.

Note that $\int_0^1 \frac{1}{\mu}\,\mathrm{d}\tilde{\tau}^0(\mu) = \frac{1}{2}4\frac{1-1/4}{1-1/2}+\frac{1}{2}\frac{4}{3}\frac{1-3/4}{1-1/2} = 10/3 > 2 = \int_0^1 \frac{1}{\mu}\,\mathrm{d}\tilde{\tau}^1(\mu)$ and $\int_0^1 G(\mu)\,\mathrm{d}\tilde{\tau}^0(\mu) = \int_0^1 G(\mu)\,\mathrm{d}\tilde{\tau}^1(\mu) = 1$. The truth-telling constraints for mediation $\int V(\mu,0)\,\mathrm{d}\tilde{\tau}^0(\mu) \geq \int V(\mu,0)\,\mathrm{d}\tilde{\tau}^1(\mu)$ and $\int V(\mu,1)\,\mathrm{d}\tilde{\tau}^1(\mu) \geq \int V(\mu,1)\,\mathrm{d}\tilde{\tau}^0(\mu)$ are satisfied. So $\tilde{\tau}$ is implementable under mediation. However, cheap talk with state-dependent utility requires $V(\mu,\omega) = V(\mu',\omega)$ for all $\omega \in \{0,1\}$ and $\mu,\mu' \in \operatorname{supp}(\tilde{\tau}^\omega)$. So $\tilde{\tau}$ is not feasible under cheap talk because $V(1/4,1) = -3 \neq -1/3 = V(3/4,1)$. As $\tilde{\tau}$ is the unique solution of the maximization problem (2) and $\tilde{\tau}$ is not feasible under cheap talk, cheap talk attains a strictly lower value than mediation.

An example without public communication where (2) and (3) holds: There is a binary state space $\Omega = \{0,1\}$ and two receivers. The pair of posteriors on $\omega = 1$ is $\mu = (\mu_1,\mu_2) \in [0,1]^2$, and the prior is $p = (1/2,1/2)$. The sender has a state-independent indirect utility $V(\mu) = G(\mu_1) - \rho\mu_2$, where $G : [0,1] \to [0,1]$ is a strictly increasing and strictly convex CDF, and $\rho > 1$ is a constant. A communication mechanism induces a joint distribution of the receivers' posterior beliefs $\tau \in \Delta([0,1]^2)$.

Because V is separable for μ_1 and μ_2, for Bayesian persuasion we can focus on the marginal distributions of posteriors $\tau_i \in \Delta([0,1])$ with $i \in \{1,2\}$. Given that G is strictly convex, the uniquely optimal distribution of posteriors for 1 is the one induced by full disclosure: $\tau_1^* = 1/2\delta_0 + 1/2\delta_1$. Because V is linear in μ_2, any information policy for receiver 2 is optimal because (BP) implies that $\int_0^1 \mu_2\,\mathrm{d}\tau_2(\mu_2) = 1/2$ for all feasible τ_2.

It can be shown using analogous steps to those in the proof of Theorem 1 that the implementation for mediation with additively separable sender's preference can be characterized by the following aggregate truth-telling constraint over marginals[53]

$$\int_0^1 G(\mu_1)(\mu_1 - \tfrac{1}{2})\,\mathrm{d}\tau_1(\mu_1) - \rho\int_0^1 \mu_2(\mu_2 - \tfrac{1}{2})\,\mathrm{d}\tau_2(\mu_2) = 0. \tag{12}$$

We next show that the mediator can attain the optimal persuasion value for the sender while satisfying (12). Consider a candidate pair of marginal distributions of beliefs (τ_1^*, τ_2) where τ_1^* corresponds to full disclosure. Equation 12 then becomes

$$\frac{1}{4} = \rho\int_0^1 \mu_2(\mu_2 - \tfrac{1}{2})\,\mathrm{d}\tau_2(\mu_2).$$

[53]Details of the proof of the characterization of the feasible distributions of receivers' beliefs are available upon request.

Now observe that for all feasible τ_2, we have $\int_0^1 \mu_2(\mu_2 - \frac{1}{2})\,\mathrm{d}\tau_2(\mu_2) \in [0, 1/4]$, where the minimum and maximum elements of the interval are respectively attained by no disclosure and full disclosure for receiver 2. In addition, by convexity of the set of Bayes plausible τ_2, there exists a feasible τ_2 such that $\int_0^1 \mu_2(\mu_2 - \frac{1}{2})\,\mathrm{d}\tau_2(\mu_2) = c$, for every $c \in [0, 1/4]$. Take a Bayes plausible τ_2^* such that $\int_0^1 \mu_2(\mu_2 - \frac{1}{2})\,\mathrm{d}\tau_2^*(\mu_2) = 1/(4\rho)$ and observe that (τ_1^*, τ_2^*) satisfies (12) by construction. In particular, (τ_1^*, τ_2^*) is optimal for Bayesian persuasion, hence the mediator can attain the optimal persuasion value.

A joint distribution τ is implementable under cheap talk (with purely private messages) only if $V(\mu_1, \mu_2) = V(\mu_1', \mu_2')$ for any $\mu, \mu' \in \text{supp}(\tau)$. This implies that full disclosure for receiver 1 is not implementable under cheap talk. To see this, fix two points $(1, \mu_2')$ and $(0, \mu_2)$ in the support of a candidate cheap talk distribution that induces full disclosure for receiver 1, and assume that these posteriors are respectively induced by the pairs of private messages (m_1', m_2') and (m_1, m_2). The sender has a profitable deviation at (m_1, m_2) by privately sending (m_1', m_2) to the receivers. Indeed, $V(1, \mu_2) > V(0, \mu_2)$, that is the deviation yields a strictly higher than the one obtained by sending (m_1, m_2). This shows that no cheap talk equilibrium can sustain full disclosure for receiver 1, hence that the optimal persuasion and mediation value cannot be attained under cheap talk.

G.2 Informativeness of Optimal Mediation

The comparison between the informativeness of the sender's optimal mediation plan and the sender's preferred cheap talk equilibria is ambiguous. In the illustration in the introduction, the sender's optimal cheap talk equilibrium is no disclosure when the prior p is in a neighborhood of 0.6, while the optimal mediation plan discloses some information about the state. We now present an example where there exists an open ball of priors such that full disclosure is optimal under cheap talk but not under mediation.

Consider a binary state space $\Omega = \{0, 1\}$ and let $\mu \in [0, 1]$ denote the posterior belief on $\omega = 1$. The sender's indirect utility function is $V(\mu) = \sin(3\pi\mu - \pi)$. For any prior $p \in (0, 1/3)$, full disclosure is optimal under cheap talk and cheap talk has value 0. Note that no disclosure is suboptimal under cheap talk and V is not quasi single-crossing at p, Proposition 5 implies that full disclosure is suboptimal under mediation.

H Correlated equilibria in long cheap talk and repeated games

In this appendix, we discuss more in detail the implications of our results for the comparison of correlated and Nash equilibria in long cheap talk and repeated games with asymmetric information where the sender's payoff is state independent.

Fix a finite set of states Ω, a finite action set A, and utility functions $u_R(\omega, a)$ and $u_S(a)$ for the receiver and the sender respectively. Following the notation in Forges (2020), let $DP_0(p)$ denote the basic decision problem described by the previous primitive objects.

The long cheap talk game is an extension of the basic decision problem $DP_0(p)$ by allowing the sender and receiver to exchange messages simultaneously for several rounds before the receiver takes an action. Formally, let two finite sets M_S and M_R be the sender and receiver's message spaces, respectively. Following Lipnowski and Ravid (2020)'s notation, we let $H_{<\infty} := \bigsqcup_{t=0}^{\infty}(M_S \times M_R)^t$ and $H_\infty := (M_S \times M_R)^{\mathbb{N}}$. The sender observes the realized state $\omega \in \Omega$ at $t = 0$. Then at each time $t = 1, 2, \ldots$, the sender sends message $m_t \in M_S$ and the receiver sends $\tilde{m}_t \in M_R$ simultaneously. Finally, after seeing the sequence of messages $h_\infty \in H_\infty$, the receiver chooses an action $a \in A$. A strategy for the sender is a measurable function $\sigma : \Omega \times H_{<\infty} \to \Delta M_S$ and a strategy for the receiver is a pair of measurable functions $\tilde{\sigma} : H_{<\infty} \to \Delta M_R$ and $\rho : H_\infty \to \Delta A$. We denote the long cheap talk game as $CT_\infty(p)$.

Under transparent motives, Proposition 4 of Lipnowski and Ravid (2020) shows that every sender payoff attainable in a Nash equilibrium of $CT_\infty(p)$ is also attainable in a perfect Bayesian equilibrium of the one-shot cheap-talk game. Therefore, the highest sender's expected payoff that is induced by a Nash equilibrium of $CT_\infty(p)$ coincides with the one-shot highest cheap talk value $\mathcal{V}_{CT}(p)$. A correlated equilibrium of $CT_\infty(p)$ is a Nash equilibrium of an extension of $CT_\infty(p)$ where the players privately receive correlated signals before the beginning of the game. Forges (1985) shows that the set of correlated equilibrium payoffs of the long cheap talk game $\mathcal{C}(CT_\infty(p))$ is the same as the set of all communication equilibrium payoffs of the basic decision problem $\mathcal{M}(DP_0(p))$. Therefore, the highest sender's expected payoff induced by a correlated equilibrium of $CT_\infty(p)$ coincides with the sender's payoff induced by the sender's preferred communication equilibrium $\mathcal{V}_{MD}(p)$.

A different class of games we consider is a simplified version of the infinitely repeated sender-receiver game introduced in Hart (1985). There are two action sets A_S, A_R for the sender and receiver, respectively. The sender observes the realized state $\omega \in \Omega$ at $t = 0$. Then at each time $t = 1, 2, \ldots$, the sender chooses action $a_t \in A_S$ and the receiver chooses $\tilde{a}_t \in A_R$ simultaneously. The action of the receiver is the only one that is payoff-relevant, and the sender's payoff does not depend on the state. That is, the sender's payoff at time t is $u_S(\tilde{a}_t)$ and the receiver's payoff at time t is $u_R(\omega, \tilde{a}_t)$. The actions are observed every period, and players have perfect recall. The players' overall payoffs are defined as the liminf of the expected time average of the one-period payoffs. That is, $U_S := \liminf_{T\to\infty} \mathbb{E}[\frac{1}{T}\sum_{t=1}^{T} u_S(\tilde{a}_t)]$ and $U_R := \liminf_{T\to\infty} \mathbb{E}[\frac{1}{T}\sum_{t=1}^{T} u_R(\omega, \tilde{a}_t)]$. This is the transparent-motive case of the repeated games of *pure information transmission* as defined in Forges (2020), and we denote it as $\Gamma_\infty(p)$.

The correlated equilibria of $\Gamma_\infty(p)$ are defined similarly, and Forges (1985) shows that the set of correlated equilibrium payoffs of this game $\mathcal{C}(\Gamma_\infty(p))$ coincides with the set of communication equilibrium payoffs of the basic decision problem $\mathcal{M}(DP_0(p))$. Therefore, the highest sender's expected payoff induced by a correlated equilibrium of $\Gamma_\infty(p)$ is the same as the sender's payoff in a sender's preferred communication equilibrium $\mathcal{V}_{MD}(p)$. Moreover, Lemma 2 and 4 of Habu et al. (2021) imply that every sender's Nash-equilibrium payoff of $\Gamma_\infty(p)$ corresponds to a sender's payoff of a one-stage cheap talk equilibrium.

I Other Extensions

The full-dimensionality condition Our main characterizations on the strict value of elicitation and mediation rely on the full-dimensionality condition at the prior (see Definition 3 and Lemma 1). This condition holds for almost every prior in finite games and, at every binary prior such that no disclosure is suboptimal under cheap talk.[54] However, it is more restrictive when we consider games with infinitely many actions and more than two states. Closing the gap between our sufficient and necessary condition for $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$ in Theorem 3 when the full-dimensionality condition does not hold remains an open problem. A promising route might be the following. Suppose that the full-dimensionality condition fails at p. Then, the largest dimension of the support of a cheap-talk optimal distribution τ^* of beliefs at p is $k < n-1$. We can redefine the state space $\tilde{\Omega}$ to be equal to the extreme points of the convex hull of supp(τ^*). This would also require redefining the receiver's prior belief and the sender's indirect payoff correspondence. The full-dimensionality condition holds in this redefined cheap talk environment and our characterizations can be applied. The drawback of this approach is that the new environment depends on the exact cheap talk solution τ^* considered. We leave a more detailed analysis of this issue for future research.

[54] Recall also the sufficient condition we derived in Proposition 3 for the multidimensional moment-measurable case.

Beyond transparent motives The main analysis focused on the case of the state-independent sender's payoff function. Without this assumption, it is still possible to express the Honesty constraint purely in terms of the unconditional distribution of beliefs. Suppose that the sender's indirect payoff at state ω and the receiver's posterior μ is uniquely given by $V(\mu, \omega)$. It is easy to show (see for example Doval and Skreta (2024)) that the truth-telling constraint can be written as

$$\int V(\mu, \omega) \left(\frac{\mu(\omega)}{p(\omega)} - \frac{\mu(\omega')}{p(\omega')} \right) \mathrm{d}\tau(\mu) \geq 0 \qquad \forall \omega, \omega' \in \Omega. \tag{13}$$

These are $n(n-1)$ moment constraints. Therefore, the optimal mediation problem is still linear in τ, and the same techniques of Proposition 1 can be applied to derive the sender's optimal value under mediation and show that there exists an optimal mediation plan with no more than n^2 signals. It would be more challenging to extend our remaining results. In Appendix G.1, we show via example that Theorem 2 may fail with state-dependent sender's payoff. We leave the formal analysis of the general state-dependent case for future research.[55]

Multiple receivers and private communication Our analysis can be immediately extended to the case with multiple receivers interacting in a game conditional on some public information, that is, the mediator sends the same message to all the receivers. In this case, the indirect payoff correspondence $\mathbf{V}(\mu)$ collects all possible expected sender's expected payoff across all the correlated equilibria of the game the receivers play conditional on public belief μ. This correspondence is still upper hemi-continuous and therefore all our results extend to this case.

Instead, if the mediator can privately communicate with every single receiver, then the analysis would be considerably more challenging.[56] However, some of our results can be relatively easily extended in the intermediate case where communication is private but the receivers do not interact in the game but rather solve an isolated decision problem, and the payoff of the sender is additively separable with respect to the profile of receiver's actions. This case would be trivial under standard Bayesian persuasion: the sender can just solve multiple different single-receiver Bayesian persuasion problems. This is not the case for a

[55] When $V(\mu, \omega) = \tilde{V}(\mu)b(\omega) + a(\omega)$ for some continuous function $\tilde{V}(\mu)$, strictly positive vector $b \in \mathbb{R}^n_{++}$, and arbitrary vector $a \in \mathbb{R}^n$ all our results apply as written. This immediately follows from the fact that $b(\omega)$ and $a(\omega)$ drop from (TT) and from the sender's unconditional expected payoff due to (zeroCov). Observe that the sender here can also be said to have "transparent motives" because the sender's preferences at different states are positive affine transformations of each other.

[56] Even without the truth-telling constraint, the analysis of the standard information design problem is complicated by the fact that potentially all the higher-order beliefs of the receivers matter. See, for example, Mathevet et al. (2020) for a belief-based analysis of the information-design problem with multiple receivers interacting in a game.

mediator who must elicit information from the sender, even if they maximize the sender's payoff. The reason is that the truth-telling constraint will not be separable with respect to the receiver's posterior beliefs.

In particular, in Appendix G.1, we show by example that already in the intermediate setting described above, the mediation trilemma fails: with private communication, a mediator can achieve the first-best Bayesian persuasion value whilst strictly improving on cheap talk, and this is true even under transparent motives.

J Infinite State Space

In this appendix, we show our main theorems extend to the case when Ω is a compact metric space. The other parts of the model are the same as Section 2. By Revelation Principle, we focus on the CE outcomes $\pi \in \Delta(\Omega \times A)$ that satisfy Consistency, Obedience, and Honesty. Here, Honesty means for p-almost all $\omega \in \Omega$, $\mathbb{E}_{\pi^\omega}[u_S(a)] = \max_{\omega' \in \Omega} \mathbb{E}_{\pi^{\omega'}}[u_S(a)]$.

We define the indirect value correspondence $\mathbf{V} : \Delta(\Omega) \rightrightarrows \mathbb{R}$ as in Section 3, which is upper hemi-continuous, compact, convex, and non-empty valued, and the upper (lower) envelopes are denoted as $\overline{V}$ ($\underline{V}$). As Definition 1, we say a distribution $\tau \in \Delta(\Delta(\Omega))$ and a measurable function $V : \Delta(\Omega) \to \mathbb{R}$ are induced by some CE outcome $\pi \in \Delta(\Omega \times A)$ if $\tau = (\phi^\pi)_\# \operatorname{marg}_A \pi$ and $V(\mu) = \mathbb{E}_\pi[u_S | \phi^\pi(a) = \mu]$, where $\phi^\pi : A \to \Delta(\Omega)$ is defined by $\phi^\pi(a) = \pi^a$.

Theorem 1 in the main text can be extended as follows:

Theorem 1*. If a distribution $\tau \in \Delta(\Delta(\Omega))$ and a measurable function $V : \Delta(\Omega) \to \mathbb{R}$ are induced by some CE outcome, then it satisfies

(i) Consistency*:

$$\int_{\Delta(\Omega)} \mu \, \mathrm{d}\tau(\mu) = p;$$

(ii) Obedience*: For τ-almost all $\mu \in \Delta(\Omega)$, $V(\mu) \in \mathbf{V}(\mu)$;

(iii) Honesty*:

$$\int_{\Delta(\Omega)} V(\mu)(\mu - p) \, \mathrm{d}\tau(\mu) = \mathbf{0}.$$

Conversely, if (τ, V) satisfies (i), (ii), and (iii), then there exists a CE outcome $\pi \in \Delta(\Omega \times A)$ such that $\mathbb{E}_\tau[V] = \mathbb{E}_\pi[u_S]$.

Proof. Suppose $\tau \in \Delta(\Delta(\Omega))$ and $V : \Delta(\Omega) \to \mathbb{R}$ are induced by some communication equilibrium outcome $\pi \in \Delta(\Omega \times A)$. For every $h \in C(\Omega)$,

$$\int_{\Delta(\Omega)} \langle h, \mu \rangle \, \mathrm{d}\tau(\mu) = \int_A \langle h, \phi^\pi(a) \rangle \, \mathrm{d} \operatorname{marg}_A \pi(a) = \int_A \langle h, \pi^a \rangle \, \mathrm{d} \operatorname{marg}_A \pi(a)$$
$$= \int_{\Omega \times A} h(\omega) \, \mathrm{d}\pi(\omega, a) = \langle h, p \rangle.$$

where $\mathbb{I}$ denotes the indicator function. The first equality is by $\tau = (\phi^\pi)_\# \operatorname{marg}_A \pi$, the second equality is by definition, the third one is by the law of iterated expectations, and the last one is by Consistency of π. Hence, τ satisfies Consistency*.

Since V is induced by π, $V(\mu)$ is the conditional expectation of u_S with respect to $\operatorname{marg}_A \pi$, conditional on $\phi(a) = \mu$. Note that by Obedience, π is supported on $a \in A^*(\mu)$ only, where $A^*(\mu) = \operatorname{argmax}_{a \in A} \mathbb{E}_\mu[u_R(\omega, a)]$ is nonempty-compact-valued and weakly measurable by the measurable maximum theorem (Aliprantis and Border, 2006, Theorem 18.19). Therefore, $V(\mu) \in [\min_{a \in A^*(\mu)} u_S(a), \max_{a \in A^*(\mu)} u_S(a)]$ and V is measurable, so Obedience* is satisfied.

By Honesty of π and the fact that u_S does not depend on ω, we have $\mathbb{E}_{\pi^\omega}[u_S] = \mathbb{E}_\pi[u_S]$ for any $\omega \in \Omega$. For any $h \in C(\Omega)$,

$$\int_{\Omega \times A} u_S(a) h(\omega) \, \mathrm{d}\pi(\omega, a) = \int_\Omega h(\omega) \mathbb{E}_{\pi^\omega}[u_S] \, \mathrm{d}p(\omega) = \langle h, p \rangle \mathbb{E}_\pi[u_S] = \langle h, p \rangle \int_{\Delta(\Omega)} V(\mu) \, \mathrm{d}\tau(\mu),$$

where the first equality is by iterated expectation and Consistency, the second one follows from Honesty, and the last one is by the fact that (τ, V) is induced by π. We also have

$$\int_{\Omega \times A} u_S(a) h(\omega) \, \mathrm{d}\pi(\omega, a) = \int_A u_S(a) \langle h, \pi^a \rangle \, \mathrm{d} \operatorname{marg}_A \pi(a)$$
$$= \int_A \mathbb{E}[u_S(a) \langle h, \pi^a \rangle \mid \pi(a) = \mu] \, \mathrm{d} \operatorname{marg}_A \pi(a) = \int_A V(\phi(a)) \langle h, \phi^\pi(a) \rangle \, \mathrm{d} \operatorname{marg}_A \pi(a)$$
$$= \int_{\Delta(\Omega)} V(\mu) \langle h, \mu \rangle \, \mathrm{d}\tau(\mu),$$

where the first two equalities are by iterated expectation, the third one follows from the fact that V is induced by π, and the last equality is by $\tau = (\phi^\pi)_\# \operatorname{marg}_A \pi$. Therefore,

$$\int_{\Delta(\Omega)} V(\mu) \langle h, \mu - p \rangle \, \mathrm{d}\tau(\mu) = 0$$

for every $h \in C(\Omega)$, so Honesty* holds.

Next, we show by construction that for any $\tau \in \Delta(\Delta(\Omega))$ and $V \in \mathbf{V}$ that satisfy

Consistency* and Honesty*, there exists a communication equilibrium outcome π with $\mathbb{E}_\tau[V] = \mathbb{E}_\pi[u_S]$. Since $V \in \mathbf{V}$, by Lemma 2 of Lipnowski and Ravid (2020), there exists a measurable $\lambda : \Delta(\Omega) \to \Delta(A)$ such that for all $\mu \in \Delta(\Omega)$, $\lambda(\mu) \in \operatorname{argmax}_{\alpha \in \Delta(A)} \mathbb{E}_\mu[u_R(\alpha, \omega)]$ is a mixed best response for the receiver with posterior μ, and $V(\mu) = \int_A u_S(a) \, \mathrm{d}\lambda(\mu)(a)$.

Given two measurable spaces $(S, \mathcal{S})$ and $(T, \mathcal{T})$, a probability kernel from S to T is a mapping $\nu : S \times \mathcal{T} \to \mathbb{R}_+$ such that $\nu(\cdot, B)$ is $\mathcal{S}$-measurable for fixed $B \in \mathcal{T}$ and $\nu(s, \cdot)$ a probability measure on T for fixed s. By Online Appendix 3 of Kamenica and Gentzkow (2011), there exists a probability kernel σ from Ω to $\Delta(\Omega)$ that induces τ by Consistency*. That is, $p \otimes \sigma \in \Delta(\Omega, \Delta(\Omega))$ has a regular conditional probability such that $(p \otimes \sigma)^\mu = \mu$ almost surely and $\tau = \operatorname{marg}_{\Delta(\Omega)}(p \otimes \sigma)$. Define a probability kernel κ from $\Delta(\Omega)$ to A such that $\kappa(\mu, \cdot) = \lambda(\mu)$. As λ is measurable, κ is a well-defined probability kernel. The composition $\sigma \otimes \kappa$ is a kernel from Ω to A that satisfies $\sigma \otimes \kappa(\omega, D) = \int_{\Delta(\Omega)} \int_A \mathbb{I}[a \in D] \, \mathrm{d}\kappa(\mu, a) \, \mathrm{d}\sigma(w, \mu)$ for every $\omega \in \Omega$ and Borel $D \subseteq A$. Let $\pi = p \otimes \sigma \otimes \kappa \in \Delta(\Omega \times A)$, we show this is the desired CE outcome. First, note that π satisfies Consistency by construction, as $\pi(W, A) = \int_W \int_{\Delta(\Omega)} \lambda(\mu)(A) \, \mathrm{d}\sigma(\omega, \mu) \, \mathrm{d}p(\omega) = \int_W \sigma(\omega, \Delta(\Omega)) \, \mathrm{d}p(\omega) = p(W)$ for any Borel $W \subseteq \Omega$.

For any function $u : \Omega \times A \to \mathbb{R}$,

$$\begin{aligned} \mathbb{E}_\pi[u] &= \int_\Omega \int_{\Delta(\Omega)} \int_A u(\omega, a) \, \mathrm{d}\kappa(\mu, a) \, \mathrm{d}\sigma(\omega, \mu) \, \mathrm{d}p(\omega) \\ &= \int_{\Delta(\Omega)} \int_\Omega \int_A u(\omega, a) \, \mathrm{d}\lambda(\mu)(a) \, \mathrm{d}(p \otimes \sigma)^\mu(\omega) \, \mathrm{d}\tau(\mu) \\ &= \int_{\Delta(\Omega)} \mathbb{E}_\mu \left(\int_A u(\omega, a) \, \mathrm{d}\lambda(\mu)(a) \right) \mathrm{d}\tau(\mu), \end{aligned} \tag{14}$$

where the first equality follows from construction of π, the second and the third are implied by the fact that τ is induced by σ.

To see Obedience, take any measurable $\tilde{a} : A \to \Delta(A)$, by definition of λ, we have $\mathbb{E}_\mu \left(\int_A u(\omega, a) \, \mathrm{d}\lambda(\mu)(a) \right) \geq \mathbb{E}_\mu \left(\int_A u(\omega, \tilde{a}) \, \mathrm{d}\lambda(\mu)(a) \right)$ for any $\mu \in \Delta(\Omega)$. Taking expectation with respect to τ, we have $\int u_R(\omega, a) \, \mathrm{d}\pi(\omega, a) \geq \int u_R(\omega, \tilde{a}) \, \mathrm{d}\pi(\omega, a)$, and π satisfies Obedience by Lemma 3.

By construction, $\sigma \otimes \kappa$ is a regular conditional probability of π given ω. Hence,

$$\mathbb{E}_{\pi^\omega}[u_S] = \int_{\Delta(\Omega)} \left(\int_A u_S(a) \, \mathrm{d}\lambda(\mu)(a) \right) \mathrm{d}\sigma(\omega, \mu) = \int_{\Delta(\Omega)} V(\mu) \, \mathrm{d}\sigma(\omega, \mu).$$

By (14), for every Borel $W \subseteq \Omega$,

$$\int_{\Omega\times A} u_S(a)\mathbb{I}[\omega \in W]\,\mathrm{d}\pi(\omega, a) = \int_{\Delta(\Omega)} V(\mu)\mu(W)\,\mathrm{d}\tau(\mu) = p(W)\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau(\mu) = p(W)\mathbb{E}_\pi[u_S],$$

where the second equality follows from Honesty*. This then implies that

$$\int_W (\mathbb{E}_{\pi^\omega}[u_S] - \mathbb{E}_\pi[u_S])\,\mathrm{d}p(\omega) = 0$$

for every Borel $W \subseteq \Omega$. Hence, $\mathbb{E}_{\pi^\omega}[u_S] = \mathbb{E}_\pi[u_S]$ p-almost surely, π satisfies Honesty. ■

As in the proof of Proposition 1, we may consider the auxiliary program (η-MD). Fix any sequence of feasible η_n that converges weakly to η, for every $h \in C(\Omega)$, $0 = \int\langle h, \mu - p\rangle\,\mathrm{d}\eta_n(\mu, s) \to \int\langle h, \mu - p\rangle\,\mathrm{d}\eta(\mu, s)$ and $0 = \int s\langle h, \mu - p\rangle\,\mathrm{d}\eta_n(\mu, s) \to \int s\langle h, \mu - p\rangle\,\mathrm{d}\eta(\mu, s)$. So the feasibility set of the auxiliary program is compact, and hence the attainment part in Proposition 1 extends.

To extend Theorem 2 to the general case, we assume that the upper envelope $\overline{V}$ is Lipschitz on $\Delta(\Omega)$.[57] Kolotilin et al. (2024) show that this is a sufficient condition for the dual attainment of Bayesian persuasion.

Theorem 2*. Suppose $\overline{V}$ is Lipschitz on $\Delta(\Omega)$. Then $\mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p)$ if and only if $\mathcal{V}_{BP}(p) = \mathcal{V}_{CT}(p)$.

Proof. It suffices to show the only if direction. Take $\tau \in \Delta(\Delta(\Omega))$ and $V \in \mathbf{V}$ that solves the mediation problem. Since $\mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p)$, $V = \overline{V}$ almost surely with respect to τ. By Corollary 1 and 2 of Kolotilin et al. (2024), there exists a Lipschitz $f : \Omega \to \mathbb{R}$ such that $\overline{V}(\mu) \leq \langle f, \mu\rangle$ for all $\mu \in \Delta(\Omega)$ and $\overline{V}(\mu) = \langle f, \mu\rangle$ for all $\mu \in \mathrm{supp}(\tau)$. By the truth-telling constraint, $\int \overline{V}(\mu)\langle f, \mu - p\rangle\,\mathrm{d}\tau(\mu) = 0$. It follows that $\int \overline{V}(\mu)^2\,\mathrm{d}\tau(\mu) = \int \overline{V}(\mu)\langle f, \mu\rangle\,\mathrm{d}\tau(\mu) = \langle f, p\rangle \int \overline{V}(\mu)\,\mathrm{d}\tau(\mu) = (\int \overline{V}(\mu)\,\mathrm{d}\tau(\mu))^2$. Therefore, $\overline{V}$ is τ-almost surely constant, implying that τ is also feasible under cheap talk. ■

Next, we show the first statement of Theorem 3 can be extended. We first introduce an expanded cheap talk hull

$$\tilde{H} = \{\mu \in \Delta(\Omega) : \exists\alpha > 1 \text{ such that } (\alpha p + (1-\alpha)\mu, \overline{V}_{CT}(p)) \in \mathrm{Gr}(\mathbf{V}_{CT})\}.$$

Theorem 3*. If there exists $\mu \in \tilde{H}$ and $\lambda \in (0, 1)$ such that $\overline{V}_{CT}(\lambda\mu + (1-\lambda)p) > \overline{V}_{CT}(p) > \underline{V}_{CT}(\mu)$, then $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$.

[57] We endow $\Delta(\Omega)$ with the Kantorovich-Rubinstein norm that metrizes the weak* topology.

Proof. Take $\mu \in \tilde{H}$ and $\lambda \in (0,1)$ that satisfies the assumption. Note that $\tilde{H}$ is convex and $p \in \tilde{H}$, any point in $\text{co}\{p, \mu\}$ is also in $\tilde{H}$. The same construction as in the proof of Theorem 3 leads to the desired result. ■

To extend Theorem 4, we first extend our definition of moments. Let T be a continuous linear map from $\Delta(\Omega)$ to a locally convex space. We say T is multidimensional if $X = T(\Delta(\Omega))$ has a dimension strictly larger than 1. We assume that $\mathbf{V} = V$ is singleton-valued and that $V(\mu) = v(T(\mu))$ for some continuous strictly quasiconvex $v : X \to \mathbb{R}$.

Theorem 4*. If $\tilde{H}(p) = \Delta(\Omega)$ and T is multidimensional, then exactly one of these cases holds:

(1) $\max V = \mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p) = \mathcal{V}_{CT}(p) > V(p)$;

(2) $\max V > \mathcal{V}_{BP}(p) > \mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p) > V(p)$.

Proof. As T is multidimensional and v is strictly quasiconvex, Corollary 6 of Lipnowski and Ravid (2020) implies no disclosure is suboptimal under cheap talk. By Theorem 2* and the fact that $\tilde{H}(p) = \Delta(\Omega)$, $\mathcal{V}_{BP}(p) = \mathcal{V}_{MD}(p)$ if and only if $\{V > \mathcal{V}_{CT}(p)\} = \emptyset$. If $\max V = \mathcal{V}_{CT}(p)$, then (1) trivially holds. It suffices to show that $\max V > \mathcal{V}_{CT}(p)$ implies (2).

To see $\max V > \mathcal{V}_{BP}(p)$, note that otherwise there exists a Bayes-plausible τ^* such that $V(\mu) = \max V$ for all $\mu \in \text{supp}(\tau^*)$, which then implies τ^* is also feasible under cheap talk, contradiction.

To show $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$, we follow the same construction as in the proof of Theorem 4. In particular, there exists $\mu_1, \mu_2 \in D_+$ such that $\hat{\mu} \coloneqq \frac{1}{2}\mu_1 + \frac{1}{2}\mu_2 \in D_-$, where $D_+ = \{V > \overline{V}_{CT}(p)\}$ and $D_- = \{V < \overline{V}_{CT}(p)\}$. By continuity, D_+ is open, and so is $\text{co}\, D_+$. Therefore, there exists $\lambda \in (0,1)$ such that $\lambda\hat{\mu} + (1-\lambda)p \in \text{co}\, D_+$. We may partially extend Lemma 6 to show that for every $\mu \in \Delta(\Omega)$, $\mu \in \text{co}\, D_+$ implies $\overline{V}_{CT}(\mu) > \overline{V}_{CT}(p)$ and $\mu \in \text{co}\, D_-$ implies $\overline{V}_{CT}(\mu) < \overline{V}_{CT}(p)$. Therefore, $\overline{V}_{CT}(\lambda\hat{\mu} + (1-\lambda)p) > \overline{V}_{CT}(p) > \overline{V}_{CT}(\hat{\mu})$. Theorem 3* implies $\mathcal{V}_{MD}(p) > \mathcal{V}_{CT}(p)$. ■