

INFORMED COMMUNICATION EQUILIBRIUM ^{*}

Frédéric KOESSLER [†] Vasiliki SKRETA [‡]

June 2, 2026

Abstract

We consider a privately informed principal selecting a communication device to influence players' actions. In contrast to standard information design, the principal cannot commit ex-ante, and the device must elicit information from all players. We define an informed communication equilibrium (informed CE) as a perfect Bayesian equilibrium outcome of this informed communication design game. We show that the set of informed CE is the subset of communication equilibria (CE) that yield principal payoff vectors bounded below by an equilibrium payoff vector of the silent game, under some consistent interim beliefs. The ex-ante optimal CE may fail to be an informed CE. We provide sufficient conditions for a CE to be an informed CE and identify classes of games in which CE and informed CE coincide.

KEYWORDS: communication equilibrium, informed principal, correlated equilibrium, mediation.

JEL CLASSIFICATION: C72; D82.

^{*}We thank the editor and the anonymous referees for their very helpful comments, and Alex Bloedel, Françoise Forges, Johannes Hörner, Marie Laclau, Elliot Lipnowski, Andrés Salamanca, and Tristan Tomala for useful discussions and comments. We are particularly grateful to Daniel Clark for providing a counterexample to a conjecture that, in sender-receiver games, an ex-ante optimal communication equilibrium is an informed communication equilibrium. Frédéric Koessler gratefully acknowledges the support of the French National Research Agency under grant ANR SCOCOS (ANR-25-CE26-0626).

[†]HEC Paris - CNRS, GREGHEC, 1 rue de la Libération, 78351, Jouy-en-Josas France. koessler@hec.fr.

[‡]UT Austin, UCL and CEPR. vskreta@gmail.com.

1 Introduction

Communication can shape outcomes in games. The concept of communication equilibrium (CE), introduced by Myerson (1982, 1986) and Forges (1986), characterizes Nash equilibrium outcomes when players are endowed with arbitrary communication possibilities. A communication device collects reports from players and provides private recommendations, potentially correlating actions and enabling coordination beyond what unmediated communication achieves. This framework has been applied to settings ranging from conflict resolution to organizational design and market platforms (Goltsman et al., 2009; Ganguly and Ray, 2023; Hörner et al., 2015; Salamanca, 2024; Corrao, 2023).

A natural question is: what happens when an interested party with private information strategically selects the device? For example, platform operators, organizational leaders or diplomats design communication structures to serve their objectives and possess private information affecting their incentives. This paper studies *informed communication design*: a privately informed principal selects a communication device after observing their type. Our setting bridges two literatures. From communication equilibrium, we inherit the device’s power to correlate actions and elicit information. From the informed principal problem (Myerson, 1983), we inherit the credibility constraint that devices must be an equilibrium choice for the principal at the interim stage, after observing their type.

We provide a complete characterization of equilibrium outcomes arising when the principal has observed relevant information prior to choosing an arbitrary communication device. The device itself must be designed to *elicit* truthful information from all parties. Our concept of *informed communication equilibrium* (informed CE) characterizes the CE that survive deviations to arbitrary communication devices chosen by a privately informed, strategic principal. The communication device specifies to whom players are able to talk, the precision of the messages and instructions players receive, whether messages are private, public or semi-public and so forth. Players can lie and are free to follow or not the device’s instructions. An important feature of our model is that the principal’s device is *exclusive*: players cannot communicate outside it. This assumption fits environments where the principal controls the communication infrastructure (platforms, organizations, legal mediation).

We consider a base game with incomplete information and $n + 1$ players: a privately informed principal and n privately informed agents. The type profile is drawn from a common prior with full support, allowing for correlated types. Each player’s payoff depends arbitrarily on the profiles of actions and types. Sender-receiver games are a special case with a single informed principal, no action choice, and a single uninformed agent. An informed CE is de-

defined as a perfect Bayesian equilibrium outcome of the following informed communication design game: Nature draws the types of the principal and the agents. The informed principal then selects a communication device which maps players' reports to a distribution over private messages. After observing the device, all players submit confidential reports. Then, they receive private messages and, finally, choose actions.

Our main result (Theorem 1) shows that the set of informed CE coincides with the set of CE that yield payoff vectors which are interim individually rational for the principal. The set of interim individually rational payoff vectors, denoted by INTIR, consists of the principal's payoff vectors bounded below by an equilibrium payoff vector of the silent game (the base game without communication) for some consistent beliefs about the principal's type. The characterization in Theorem 1 offers a tractable method for identifying all perfect Bayesian equilibrium outcomes of the informed communication design game without solving any signaling or alternative communication game. For example, informed CE coincides with CE for all base games in which the worst payoff for each principal type is an equilibrium payoff in the silent game, as in Hörner et al. (2015).

In sender-receiver games, we show that all CE payoff vectors lie in the convex hull of INTIR, so CE and informed CE coincide when INTIR is convex, as in settings with transparent motives (Chakraborty and Harbaugh, 2010; Lipnowski and Ravid, 2020; Corrao and Dai, 2023). Importantly, while prior work in the mediation design literature (Myerson, 1982; Salamanca, 2021; Corrao, 2023; Corrao and Dai, 2023; Salamanca, 2024) focuses on ex-ante optimal CE, we show that such equilibria may fail to be informed CE. In a binary-action multi-receiver example, the principal's payoff vector corresponding to the ex-ante optimal communication device lies outside the INTIR set, prompting some principal types to deviate to alternative communication devices. This example illustrates the conflict between ex-ante optimality and interim credibility.

Literature. This paper contributes to and connects the literature on communication in games and informed principal problems. There is a long and influential literature in economics studying strategic information transmission by informed players via cheap talk messages (Crawford and Sobel, 1982; Forges, 1990b; Aumann and Hart, 2003; Chakraborty and Harbaugh, 2010; Blume and Board, 2013; Lipnowski and Ravid, 2020). We build on the correlated equilibrium (Aumann, 1974, 1987) and communication equilibrium (Myerson, 1982, 1986; Forges, 1986, 1990a) literatures. We depart by assuming a privately informed principal strategically selects the communication device. Informed communication design refines communication equilibrium by accounting for the principal's in-

centives at the device selection stage. Some communication equilibria fail this refinement because they yield the principal less than they could secure by choosing silence, that is, a trivial device followed by Nash equilibrium play in the resulting game.

A related line of research studies the implementation of communication equilibria through unmediated cheap talk among all players (see, e.g., Āzacis et al., 2025 and references therein). In that literature, players jointly determine the information structure without a mediated device. Depending on assumptions about the number of players and access to correlation devices, the entire set of communication equilibria may be implementable. Carmona and Laohakunakorn (2026) study complete information games in which every player can propose a correlation device, with one device randomly selected to generate the private messages. By contrast, our framework features a single informed principal who selects the communication device to affect play in arbitrary Bayesian games.

Our framework is a special case of the informed principal problem introduced by Myerson (1983), but differs in that there are no contractible decisions. Most of the literature on informed principals focuses on contractible decisions and abstracts from non-contractible actions (Maskin and Tirole, 1990, 1992; Mylovanov and Tröger, 2012, 2014; Koessler and Skreta, 2016). Work that allows for non-contractible actions typically considers single-agent settings with transfers (Wagner et al., 2015; Mekonnen, 2021; Clark, 2023, 2022). Koessler and Skreta (2023) study a general multi-agent setting without contractible decisions, but their informed principal selects devices that operate directly on verifiable information, without requiring elicitation as is the case in the information design literature (Forges, 1993; Kamenica and Gentzkow, 2011; Bergemann and Morris, 2016; Taneva, 2019). Information is also verifiable in De Clippel and Minelli (2004), who study a two-person bargaining problem with an informed principal. Like us, Salamanca (2025) studies interim communication design, but in a special setting (a two-state sender-receiver game with quadratic loss functions) and focuses on equilibrium refinements such as those proposed by Myerson (1983) and, more recently, by Mylovanov and Tröger (2025).

To our knowledge, this is the first paper to characterize the full set of perfect Bayesian equilibrium outcomes in a general informed principal problem. The absence of contractible decisions and the reliance on elicited information make strategic deviations by the principal to arbitrary communication devices tractable. In settings with contractible actions (Myerson, 1983; Maskin and Tirole, 1990; Mylovanov and Tröger, 2012; Koessler and Skreta, 2016) or verifiable information (Koessler and Skreta, 2023), sustaining equilibrium requires off-path beliefs and continuation equilibria tailored to the specific deviation. For instance, if a deviating mechanism reveals verifiable information that can-

not be credibly ignored (as in Koessler and Skreta 2023), off-path beliefs must be carefully constructed to render the deviation unprofitable. In our setting, by contrast, any deviation can be met with a babbling continuation equilibrium under beliefs independent of the deviation. This is because players can ignore the device's messages and play any Nash equilibrium of the silent game. This canonical response is available precisely because the device makes non-binding action recommendations, rather than contractible decisions or verifiable information that players cannot credibly disregard, and it drives our tractable characterization.

2 Model

Base game. A *base game* is a finite Bayesian game with $n+1$ players. Player 0 is the *principal*; players $i = 1, \dots, n$ are the *agents*. Let $I = \{1, \dots, n\}$ be the set of agents and $I_0 = \{0, 1, \dots, n\}$ the set of all players. For each $i \in I_0$, the non-empty finite set of actions available to player i is A_i (the principal can also choose an action). An action profile is denoted by $a = (a_0, a_1, \dots, a_n) \in A := \times_{i \in I_0} A_i$. The principal is privately informed about their type $t \in T$, and each agent i is privately informed about their type $\theta_i \in \Theta_i$, where T and each Θ_i are non-empty and finite. The agents' type profile is $\theta = (\theta_1, \dots, \theta_n) \in \Theta := \times_{i \in I} \Theta_i$.

The common prior is $p \in \Delta(T \times \Theta)$ and has full support. For each $t \in T$, write $p(t) := \sum_{\theta \in \Theta} p(t, \theta)$ for the marginal on T . Agent i 's belief about (t, θ_{-i}) conditional on θ_i is

$$p_i(t, \theta_{-i} \mid \theta_i) = \frac{p(t, \theta)}{\sum_{\tilde{t}, \tilde{\theta}_{-i}} p(\tilde{t}, \theta_i, \tilde{\theta}_{-i})}.$$

The principal's belief about $\theta = (\theta_i, \theta_{-i})$ conditional on t is

$$p_0(\theta \mid t) = \frac{p(t, \theta)}{p(t)}.$$

For every action profile $a \in A$ and type profile $(t, \theta) \in T \times \Theta$, the principal's payoff is denoted by $u(a, t, \theta)$, and agent i 's payoff by $v_i(a, t, \theta)$.

The base game described above is denoted by

$$G = \langle I_0, T, (\Theta_i)_{i \in I}, p, (A_i)_{i \in I_0}, u, (v_i)_{i \in I} \rangle.$$

A strategy profile of G is a profile of type-contingent mixed actions $(\sigma_i)_{i \in I_0}$, with $\sigma_0 : T \rightarrow$

$\Delta(A_0)$ and $\sigma_i : \Theta_i \rightarrow \Delta(A_i)$ for each $i \in I$. An *outcome* of G is a mapping $\mu : T \times \Theta \rightarrow \Delta(A)$. A Nash equilibrium outcome of G is an outcome induced by some Nash equilibrium strategy profile of G .

Communication devices and equilibria. A communication device is a mapping $\nu : R \rightarrow \Delta(M)$, where $R = \times_{i \in I_0} R_i$ and $M = \times_{i \in I_0} M_i$. Here, R_i is a non-empty finite set of reports that player i can send to the device, and M_i is a non-empty finite set of messages the device can send to player i . A *direct* communication device is a mapping $\nu : T \times \Theta \rightarrow \Delta(A)$, i.e., a device in which each player's report corresponds to their type ($R_0 = T$ and $R_i = \Theta_i$ for all $i \in I$), and each message is an action recommendation ($M_i = A_i$ for all $i \in I_0$). For notational simplicity we only consider direct communication devices in the rest of the paper. None of our results are affected by this restriction.¹

Let $G(\nu)$ be the base game extended by ν . In $G(\nu)$, the principal sends a report $r_0 \in T$ as a function of t , and each agent i sends a report $r_i \in \Theta_i$ as a function of θ_i . The profile of messages $m = (m_0, \dots, m_n) \in A$ is drawn according to $\nu(m_0, \dots, m_n \mid r_0, \dots, r_n)$, and each player i privately observes m_i and chooses an action $a_i \in A_i$. Players' payoffs depend only on the type and action profile, as in the base game G .

A *reporting strategy* for the principal is denoted by $\tau_0 : T \rightarrow \Delta(T)$, and a reporting strategy for agent i is denoted by $\tau_i : \Theta_i \rightarrow \Delta(\Theta_i)$ for $i \neq 0$, where $\tau_0(r_0 \mid t)$ denotes the probability that the principal sends report $r_0 \in T$ when their type is $t \in T$, and similarly for the agents. A reporting strategy is *truthful* if it assigns probability one to the report equal to the player's type. An *action strategy* for the principal is denoted by $\gamma_0 : T \times T \times A_0 \rightarrow \Delta(A_0)$, and an action strategy for agent i is denoted by $\gamma_i : \Theta_i \times \Theta_i \times A_i \rightarrow \Delta(A_i)$ for $i \neq 0$, where $\gamma_0(a_0 \mid t, r_0, m_0)$ denotes the probability that the principal chooses action $a_0 \in A_0$ when their type is $t \in T$, their report is r_0 and their message is m_0 , and similarly for the agents. An action strategy is *obedient* if it assigns probability one to the action equal to the player's message. A strategy profile of $G(\nu)$ is a profile of reporting and action strategies, denoted by $(\gamma, \tau) = (\gamma_i, \tau_i)_{i \in I_0}$ and called a *participation strategy profile*.²

We write $\tau(r \mid t, \theta) := \tau_0(r_0 \mid t) \prod_{i \in I} \tau_i(r_i \mid \theta_i)$ and $\gamma(a \mid m, r, t, \theta) := \gamma_0(a_0 \mid t, r_0, m_0) \prod_{i \in I} \gamma_i(a_i \mid \theta_i, r_i, m_i)$. Given ν and the participation strategy profile $(\gamma, \tau) = (\gamma_i, \tau_i)_{i \in I_0}$, the principal's

¹See Koessler and Skreta (2025) for a previous version of the paper without this restriction.

²As in Myerson (1982, 1983), the extended game $G(\nu)$ is treated as a simultaneous-move game over reporting and action strategies. Under our full-support assumption on p , modeling $G(\nu)$ as a multistage game in extensive form and imposing sequential rationality and consistent beliefs in such an extensive form of $G(\nu)$ would not affect the set of equilibrium outcomes of $G(\nu)$ (see Gerardi and Myerson, 2007, Sugaya and Wolitzky, 2021).

expected payoff in $G(v)$ when their type is t is

$$W_0(v, \gamma, \tau | t) = \sum_{\theta \in \Theta} \sum_{r \in T \times \Theta} \sum_{m \in A} \sum_{a \in A} p_0(\theta | t) \tau(r | t, \theta) v(m | r) \gamma(a | m, r, t, \theta) u(a, t, \theta),$$

and agent i 's expected payoff for type θ_i is

$$W_i(v, \gamma, \tau | \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{t \in T} \sum_{r \in T \times \Theta} \sum_{m \in A} \sum_{a \in A} p_i(t, \theta_{-i} | \theta_i) \tau(r | t, \theta) v(m | r) \gamma(a | m, r, t, \theta) v_i(a, t, \theta).$$

A profile of participation strategies (γ^*, τ^*) is a Nash equilibrium of $G(v)$ iff

$$\begin{aligned} W_0(v, \gamma^*, \tau^* | t) &\geq W_0(v, (\gamma_0, \gamma_{-0}^*), (\tau_0, \tau_{-0}^*) | t), & \text{for all } t \in T \text{ and all } (\gamma_0, \tau_0), \\ W_i(v, \gamma^*, \tau^* | \theta_i) &\geq W_i(v, (\gamma_i, \gamma_{-i}^*), (\tau_i, \tau_{-i}^*) | \theta_i), & \text{for all } i \in I, \theta_i \in \Theta_i, \text{ and all } (\gamma_i, \tau_i). \end{aligned} \quad (1)$$

A Nash equilibrium outcome of $G(v)$ is an outcome $\mu = \gamma \circ v \circ \tau$, with

$$\mu(a | t, \theta) = \sum_{r \in T \times \Theta} \sum_{m \in A} \tau(r | t, \theta) v(m | r) \gamma(a | m, r, t, \theta) \text{ for every } a \in A, t \in T, \theta \in \Theta,$$

induced by some Nash equilibrium participation strategy profile (γ, τ) of $G(v)$.

An outcome $\mu : T \times \Theta \rightarrow \Delta(A)$ is a *communication equilibrium (CE)* of G iff it is a Nash equilibrium outcome of $G(v)$ for some v . By the revelation principle (Myerson, 1982; Forges, 1986; Myerson, 1986), μ is a communication equilibrium iff it is a *truthful and obedient* equilibrium outcome of the extended game $G(\mu)$. Hence, a communication equilibrium can be defined canonically as follows:

Definition 1. An outcome $\mu : T \times \Theta \rightarrow \Delta(A)$ is a *communication equilibrium (CE)* iff

$$\sum_{t \in T} \sum_{\theta_{-i} \in \Theta_{-i}} p_i(t, \theta_{-i} | \theta_i) \sum_{a \in A} \mu(a | t, \theta) v_i(a, t, \theta) \geq \sum_{t \in T} \sum_{\theta_{-i} \in \Theta_{-i}} p_i(t, \theta_{-i} | \theta_i) \sum_{a \in A} \mu(a | t, \theta_{-i}, \theta'_i) v_i((\delta_i(a_i), a_{-i}), t, \theta),$$

for all $i \in I$, $\theta_i, \theta'_i \in \Theta_i$, and $\delta_i : A_i \rightarrow A_i$, and

$$\sum_{\theta \in \Theta} p_0(\theta | t) \sum_{a \in A} \mu(a | t, \theta) u(a, t, \theta) \geq \sum_{\theta \in \Theta} p_0(\theta | t) \sum_{a \in A} \mu(a | t', \theta) u((\delta_0(a_0), a_{-0}), t, \theta),$$

for all $t, t' \in T$ and $\delta_0 : A_0 \rightarrow A_0$.

Ex-ante optimal communication equilibrium. The interim expected payoff of principal type t from outcome μ is denoted by

$$U(\mu | t) = \sum_{\theta \in \Theta} p_0(\theta | t) \sum_{a \in A} \mu(a | t, \theta) u(a, t, \theta),$$

and $(U(\mu | t))_{t \in T}$ is the principal's *payoff vector*. The principal's ex-ante expected payoff from outcome μ is then

$$\sum_{t \in T} p(t) U(\mu | t).$$

Definition 2. An outcome $\mu : T \times \Theta \rightarrow \Delta(A)$ is an *ex-ante optimal communication equilibrium* iff it is a CE that maximizes $\sum_{t \in T} p(t) U(\mu | t)$ among all CE.

Finding an ex-ante optimal CE is a linear programming problem and is the primary object of interest in the mediation design literature (Myerson, 1982, Mitusch and Strausz, 2005, Goltsman et al., 2009, Salamanca, 2021, 2024, Ganguly and Ray, 2023, Corrao and Dai, 2023). The next section introduces communication design *games* in which the communication device is selected strategically by the principal.

3 Communication design game

Ex-ante communication design. Before studying the informed principal game in which the principal is privately informed when selecting their communication device, we first consider the simpler ex-ante communication design game $\Gamma^E(G)$ in which the principal selects the communication device ν *before* learning t . That is, the principal publicly chooses ν , then learns t , and then $G(\nu)$ is played. A strategy for the principal in $\Gamma^E(G)$ consists of a communication device ν^* and a participation strategy $(\gamma_0(\nu), \tau_0(\nu))$ for every ν . For every agent i , a strategy of agent i in $\Gamma^E(G)$ is a participation strategy $(\gamma_i(\nu), \tau_i(\nu))$ for every ν .³

We study the perfect Bayesian equilibria of $\Gamma^E(G)$. In this benchmark, since the principal is uninformed in the device-selection stage, a perfect Bayesian equilibrium is simply a subgame-perfect equilibrium: a Nash equilibrium strategy profile (ν^*, γ, τ) of $\Gamma^E(G)$ such that $(\gamma(\nu), \tau(\nu))$ is a Nash equilibrium of $G(\nu)$ for every ν .

We define an ex-ante communication equilibrium of G as a perfect Bayesian equilibrium

³As mentioned in the introduction, notice that we implicitly assume that players cannot communicate outside the device ν selected by the principal. This “exclusive communication” assumption is implicit in the definition of the extended game $G(\nu)$. It is natural when the principal controls the communication infrastructure or when institutional factors preclude side communication.

outcome of $\Gamma^E(G)$, that is, an outcome $\mu = \gamma(\nu^*) \circ \nu^* \circ \tau(\nu^*)$ induced by a perfect Bayesian equilibrium strategy profile (ν^*, γ, τ) of $\Gamma^E(G)$. By the revelation principle we can replace any Nash equilibrium choice of ν^* by the principal with the corresponding outcome μ , with truthful and obedient strategies along the equilibrium path. This leads to the following definition.

Definition 3. An outcome μ is an *ex-ante communication equilibrium (ex-ante CE)* of G iff

- μ is a communication equilibrium of G
- For every communication device ν , there exists a Nash equilibrium outcome $\underline{\mu}$ of $G(\nu)$ such that $\sum_{t \in T} p(t)U(\mu | t) \geq \sum_{t \in T} p(t)U(\underline{\mu} | t)$.

Informed communication design. Consider now the informed communication design game $\Gamma(G)$ in which the principal selects the communication device ν *after* learning t , and then the base game extended by ν is played. The main differences from the ex-ante communication design game are as follows:

- The communication device can depend on the principal type t ;
- Agents' *beliefs* (on and off path) may differ from the prior, because the device may depend on t ;
- The no deviation condition for the principal at the device selection stage is not an ex-ante condition but an *interim* condition, one for each type t .

Hence, $\Gamma(G)$ is a signaling (informed-principal) game. A strategy profile in $\Gamma(G)$ is given by (ν, γ, τ) , where (γ, τ) is a device-dependent participation strategy profile as before, and $\nu = (\nu(t))_{t \in T}$ is a device selection strategy for the principal, where $\nu(t) : T \times \Theta \rightarrow \Delta(A)$ for every t . To define a perfect Bayesian equilibrium of $\Gamma(G)$ we have to specify players' beliefs about others' types for each realization of a communication device ν and a belief-consistency condition. In what follows, we follow the definition of belief consistency in Myerson (1983), which is an appropriate definition in signaling games with a continuum of signals (the set of possible communication devices the principal can choose from) and is standard in the informed-principal literature (see, e.g., Mylovanov and Tröger, 2025 and references therein).

Belief consistency. For every ν , $i \in I$ and $(t, \theta) \in T \times \Theta$, denote by $\tilde{p}_i(t, \theta_{-i} \mid \theta_i, \nu)$ agent i 's posterior belief and by $\tilde{p}_0(\theta \mid t, \nu)$ the principal's posterior. Given strategy profile (ν, γ, τ) , this belief system is *consistent* if the following hold. For every ν on the equilibrium path ($\nu = \nu(t)$ for some t), posteriors are obtained by Bayes' rule. For instance, if $\nu(t)$ does not depend on t , then players' posteriors equal the prior. For every ν off path ($\nu(t) \neq \nu$ for all t), there exists $Q \in \Delta(T)$ such that

$$\tilde{p}_i(t, \theta_{-i} \mid \theta_i, \nu) = p_i^*(t, \theta_{-i} \mid \theta_i, Q) := \frac{p_i(t, \theta_{-i} \mid \theta_i) Q(t)}{\sum_{\tilde{t}, \tilde{\theta}_{-i}} p_i(\tilde{t}, \tilde{\theta}_{-i} \mid \theta_i) Q(\tilde{t})}, \quad \tilde{p}_0(\theta \mid t, \nu) = p_0(\theta \mid t). \quad (2)$$

The vector $Q \in \Delta(T)$ is called a *normalized-likelihood vector* (Myerson, 1983): $Q(t)$ represents the relative likelihood that a principal selecting ν is of type t . This can be formalized by discretizing the space of communication devices and letting $\varepsilon(\nu \mid t)$ be the infinitesimal probability that principal t selects (trembles to) ν , so that $Q(t) = \varepsilon(\nu \mid t) / (\sum_{\tilde{t}} \varepsilon(\nu \mid \tilde{t}))$ and (2) follows by Bayes' rule. This formulation embodies two consistency properties: (i) players choose strategies independently ("no signaling what you don't know" Fudenberg and Tirole, 1991), which follows from $\varepsilon(\nu \mid t)$ not depending on θ ; and (ii) players with identical information hold identical beliefs, which follows from $\varepsilon(\nu \mid t)$ not depending on i so that all agents use a common Q to update their beliefs. For example, if $Q(t) = 1/|T|$ for every t , then agents' posterior beliefs coincide with their prior beliefs (passive beliefs); if $Q(t) = 1$ for some t , agents' posteriors assign probability one to principal type t regardless of their type. When types are independently distributed across players, all agents have the same belief about the principal's type and maintain their prior beliefs about other agents' types. The principal's posterior remains $p_0(\theta \mid t)$ since choosing ν reveals no information about agents' types.

For every ν and Q , let $G_Q(\nu)$ be the *base game extended by ν given Q* , defined as $G(\nu)$ except that each agent's conditional belief $p_i(t, \theta_{-i} \mid \theta_i)$ is replaced by $p_i^*(t, \theta_{-i} \mid \theta_i, Q)$, while the principal's belief remains $p_0(\theta \mid t)$. A strategy profile (ν, γ, τ) is a perfect Bayesian equilibrium of $\Gamma(G)$ iff it is a Nash equilibrium of $\Gamma(G)$ and, for every ν , there exists $Q \in \Delta(T)$ such that $(\gamma(\nu), \tau(\nu))$ is a Nash equilibrium of $G_Q(\nu)$. The outcome induced by such a strategy profile is given by $\mu : T \times \Theta \rightarrow \Delta(A)$, where

$$\mu(a \mid t, \theta) = \sum_{r \in T \times \Theta} \sum_{m \in A} \tau(\nu(t))(r \mid t, \theta) \nu(t)(m \mid r) \gamma(\nu(t))(a \mid m, r, t, \theta).$$

As shown by Myerson (1983), it is without loss of generality for equilibrium outcomes to focus on a "pooling" strategy for the principal at the device selection stage, i.e., $\nu(t) = \nu(t') = \nu$ for all t, t' (principle of inscrutability). The principal's choice of device then conveys no in-

formation, and all needed communication from the principal can be built into the device ν itself. Agents thus have less information (hence fewer possible deviations) while the same outcome is induced. In addition, by the revelation principle, one can set $\nu = \mu$ and require truthful and obedient strategies along the equilibrium path, so μ is a communication equilibrium of G . These simplifications yield the following definition of a perfect Bayesian equilibrium outcome of $\Gamma(G)$, which we call an informed communication equilibrium:

Definition 4. An outcome $\mu : T \times \Theta \rightarrow \Delta(A)$ is an *informed communication equilibrium* (informed CE) of G iff

1. μ is a communication equilibrium of G ;
2. For every communication device ν , there exist $Q \in \Delta(T)$ and a Nash equilibrium outcome $\underline{\mu}$ of $G_Q(\nu)$ such that $U(\mu | t) \geq U(\underline{\mu} | t)$ for every t .

4 Characterization and special cases

We characterize ex-ante and informed communication equilibria for general Bayesian games, then examine uninformed principals and sender-receiver games as special cases.

4.1 General characterization

The next proposition shows that a communication equilibrium is an ex-ante communication equilibrium iff it is ex-ante individually rational, i.e., the ex-ante payoff of the principal is weakly above their ex-ante payoff at some Nash equilibrium of the base game G .

Definition 5. The set of *ex-ante individually rational payoff vectors* of G , denoted by EXIR, is the set of all payoff vectors $(U(t))_{t \in T} \in \mathbb{R}^T$ such that there exists a Nash equilibrium outcome $\underline{\mu}$ of the base game G with

$$\sum_{t \in T} p(t)U(t) \geq \sum_{t \in T} p(t)U(\underline{\mu} | t).$$

Proposition 1. An outcome μ is an ex-ante communication equilibrium of G iff μ is a communication equilibrium of G and $(U(\mu | t))_{t \in T} \in \text{EXIR}$.

Proof. We have to show that condition 2 of the definition of an ex-ante CE (Definition 3) is satisfied iff $(U(\mu | t))_{t \in T} \in \text{EXIR}$. For the “only if” direction, consider a communication device ν such that the same message profile is sent regardless of the report profile. For

such a device, the set of Nash equilibrium outcomes of the extended $G(\nu)$ trivially coincides with the set of Nash equilibrium outcomes of the base game G . Hence, by definition, $(U(\mu | t))_{t \in T} \in \text{EXIR}$. For the “if” direction, it suffices to show that every Nash equilibrium outcome of G is the outcome of a “babbling” Nash equilibrium of $G(\nu)$ for every ν . Fix any ν , and consider a Nash equilibrium $(\sigma_i)_{i \in I_0}$ of G . In the extended game $G(\nu)$, define the “babbling” participation strategy profile (γ, τ) as follows:

- $\tau_0(t) = \tau_0(t')$ for all $t, t' \in T$,
- $\tau_i(\theta_i) = \tau_i(\theta'_i)$ for all $\theta_i, \theta'_i \in \Theta_i$ and $i \in I$,
- $\gamma_0(t, r_0, m_0) = \sigma_0(t)$ for all $t \in T$, $r_0 \in T$, and $m_0 \in A_0$,
- $\gamma_i(\theta_i, r_i, m_i) = \sigma_i(\theta_i)$ for all $\theta_i \in \Theta_i$, $r_i \in \Theta_i$, $m_i \in A_i$, and $i \in I$.

Clearly, (γ, τ) is a Nash equilibrium of $G(\nu)$ and induces the same outcome as $(\sigma_i)_{i \in I_0}$. ■

Our next result, Theorem 1, establishes that informed CE are the CE of G that are individually rational for the principal in the following interim sense: the principal’s payoff vector is bounded below by a payoff vector of the base game where agents hold some consistent beliefs about the principal’s type. Let G_Q be the *base game given Q* , defined as G except that each agent’s conditional belief $p_i(t, \theta_{-i} | \theta_i)$ is replaced by $p_i^*(t, \theta_{-i} | \theta_i, Q)$, while the principal’s belief remains $p_0(\theta | t)$.

Definition 6. The set of *interim individually rational payoff vectors* of G , denoted by INTIR, is the set of all payoff vectors $(U(t))_{t \in T} \in \mathbb{R}^T$ such that there exist $Q \in \Delta(T)$ and a Nash equilibrium outcome $\underline{\mu}$ of G_Q such that

$$U(t) \geq U(\underline{\mu} | t), \quad \text{for every } t \in T.$$

Notice that all Nash equilibrium payoff vectors of the base game G are both ex-ante and interim individually rational.

Theorem 1. An outcome μ is an informed communication equilibrium of G iff μ is a communication equilibrium of G and $(U(\mu | t))_{t \in T} \in \text{INTIR}$.

Proof. The proof mirrors that of Proposition 1. We have to show that condition 2 of the definition of an informed CE (Definition 4) is satisfied iff $(U(\mu | t))_{t \in T} \in \text{INTIR}$. For the “only if” direction, if the communication device ν is constant, then the set of Nash equilibrium

outcomes of $G_Q(v)$ coincides with the set of Nash equilibrium outcomes of G_Q , so $(U(\mu | t))_{t \in T} \in \text{INTIR}$. For the “if” direction, again, it is immediate that every Nash equilibrium outcome of G_Q is the outcome of a babbling Nash equilibrium of $G_Q(v)$. ■

EXIR reflects control force ex-ante (the principal selects their preferred device at the ex-ante stage). INTIR reflects a control force at the interim stage where each principal type selects their preferred device and a signaling force that allows agents’ beliefs to vary with the device chosen by the principal. There is no signaling force shaping EXIR because beliefs remain the same for all devices on and off the equilibrium path.

To isolate the control force we first study an uninformed principal where there is no signaling force and the ex-ante and the interim control force coincide. Then we proceed to study sender-receiver games, which are the simplest class of games where there is a signaling force and the ex-ante and interim control forces differ.

4.2 Uninformed principal

When the principal is uninformed ($|T| = 1$), we have $\text{EXIR} = \text{INTIR}$ and clearly the sets of ex-ante and informed communication equilibria coincide. In particular, if the base game is a *complete information game* ($|T| = |\Theta_i| = 1$ for every i), then a CE is simply a correlated equilibrium of the base game (Aumann, 1974, 1987), and the set of individually rational payoffs for the principal consists of all payoffs weakly above their worst Nash equilibrium payoff in G .

Corollary 1. *If the base game G is a complete information game, then an outcome μ is an informed communication equilibrium of G iff μ is a correlated equilibrium of G such that the principal’s payoff is weakly above their worst Nash equilibrium payoff in G .*

Corollary 1 establishes that even under complete information, the set of informed CE outcomes need not collapse to that of CE outcomes, as the principal’s discretion over the mediation device secures a payoff guarantee via the silent game; indeed, the next example shows that the set of informed CE may be strictly contained in the set of correlated equilibria.

Example 1 (Uninformed principal). Consider the two-player complete information base game G in Table 1 where the principal is the row player and the agent is the column player. The unique Nash equilibrium of G is (d, d) with payoff 4 for the principal. The correlated equilibrium μ shown in the table yields payoff 3 for the principal, which is below

the worst Nash equilibrium payoff of G . Therefore, μ is not an informed/ex-ante CE, as $3 \notin \text{INTIR} = \text{EXIR} = [4, +\infty)$.

	a	b	c	d
a	0, 0	2, 4	4, 2	3, 3
b	4, 2	0, 0	2, 4	3, 3
c	2, 4	4, 2	0, 0	3, 3
d	3, 3	3, 3	3, 3	4, 4

$$\mu = \begin{pmatrix} 0 & \frac{1}{6} & \frac{1}{6} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} & 0 \\ \frac{1}{6} & \frac{1}{6} & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Table 1: Complete information game and correlated equilibrium that is not an informed/ex-ante CE. Rows correspond to actions of the principal and columns to actions of the agent.

4.3 Sender-receiver games

We now examine sender-receiver games, a well-studied and important special case of our model. The base game G is a *sender-receiver game* if there are two players, with the sender being the only privately informed player and the receiver being the only player who takes an action. If the principal is the receiver, then the sets of informed CE and CE coincide because every CE is better than the babbling outcome for the receiver, so the interesting case is when the principal is the sender. Under this convention, $|\Theta_1| = |A_0| = 1$ and $A = A_1$ and we drop the subscript for the agent.

For this setting, every off-path posterior belief $q \in \Delta(T)$ for the receiver can be made consistent in $\Gamma(G)$ by choosing the normalized-likelihood vector Q defined by

$$Q(t) = \frac{q(t)}{p(t) \sum_{\tilde{t}} \frac{q(\tilde{t})}{p(\tilde{t})}}.$$

Indeed, with this Q , we have $p_1^*(t \mid Q) = p(t)Q(t) / (\sum_{\tilde{t}} p(\tilde{t})Q(\tilde{t})) = q(t)$ for every t . We extend players' payoff functions linearly to mixed actions in the usual way: for each $y \in \Delta(A)$ and $t \in T$, let $u(y, t) = \sum_a y(a)u(a, t)$ and $v(y, t) = \sum_a y(a)v(a, t)$. The set of mixed actions that are optimal for the receiver given belief q is

$$Y(q) := \arg \max_{y \in \Delta(A)} \sum_t q(t)v(y, t).$$

Then, a Nash equilibrium outcome of G_Q has $\mu(t) = y$ for all t , where $y \in Y(q)$. The set INTIR can be rewritten as

$$\text{INTIR} = \bigcup_{q \in \Delta(T)} \{(u(y, t))_t : y \in Y(q)\} + \mathbb{R}_+^T.$$

Hence we have:

Corollary 2. *If the base game G is a sender-receiver game, then an outcome μ is an informed communication equilibrium of G iff μ is a communication equilibrium of G and there exist $q \in \Delta(T)$ and $y \in Y(q)$ such that $U(\mu | t) \geq u(y, t)$ for every $t \in T$.*

When the principal has private information, informed and ex-ante CE are generally not comparable because INTIR imposes interim constraints (type by type) but allows off-path beliefs to vary through Q , while EXIR imposes only one ex-ante constraint but fixes beliefs at the prior. Example 2 shows that ex-ante control force can be stronger than the signaling and the interim control force combined whereas the opposite holds in Example 3.

Example 2 (Informed CE not ex-ante CE). Consider the sender-receiver game in Table 2 with prior $p(t_1) = \frac{1}{2}$. This example has transparent motives: the principal's payoff does not depend on their type. The set of CE payoff vectors is the diagonal line connecting the fully revealing payoff vector, $(0,0)$, and the non-revealing payoff vector, $(1,1)$, and INTIR is the entire positive quadrant, so the sets of CE and informed CE coincide. The ex-ante individual rationality constraint EXIR is given by $U(t_1) + U(t_2) \geq 2$, so the unique ex-ante CE is the non-revealing outcome.

	a_1	a_2	a_3
t_1	0, 3	1, 2	0, 0
t_2	0, 0	1, 2	0, 3

Table 2: Sender-receiver game for Example 2. Rows correspond to the sender's types (chosen by nature) and columns to the receiver's actions.

Example 3 (Ex-ante CE not informed CE). Consider the sender-receiver game in Table 3 with prior $p(t_1) = \frac{1}{2}$. Here, an outcome μ is an ex-ante CE if and only if it is a CE with $U(t_1) + U(t_2) \geq 4$. Consider the CE given by

$$\begin{aligned} \mu(a^1 | t_1) &= 1 - \alpha & \mu(a^5 | t_1) &= \alpha \\ \mu(a^1 | t_2) &= \alpha & \mu(a^5 | t_2) &= 1 - \alpha \end{aligned}$$

with $\alpha = \frac{4}{5}$. This yields $(U(t_1), U(t_2)) = (1, 4)$, which satisfies the ex-ante constraint and is therefore an ex-ante CE. However, as shown in Figure 1, this payoff vector lies outside INTIR and is therefore not an informed CE payoff vector.

receiver's obedience constraint, $\delta_a \in Y(q_a)$ for every $a \in \bar{A}$, where $q_a = (q_a(t))_t$ is defined by

$$q_a(t) = \frac{\mu(a | t)p(t)}{\sum_{\tilde{t}} \mu(a | \tilde{t})p(\tilde{t})}.$$

Hence, for every $a \in \bar{A}$, the vector $(u(a, t))_t$ belongs to INTIR, and the convex combination $(\sum_a \mu(a | \bar{t})u(a, t))_t$ belongs to $\text{co}(\text{INTIR})$. From the principal's truth-telling constraint under μ , we also have, for every t ,

$$U(\mu | t) \geq \sum_a \mu(a | \bar{t})u(a, t).$$

Since $\text{co}(\text{INTIR})$ is an upper set, we get $(U(\mu | t))_t \in \text{co}(\text{INTIR})$. ■

Proposition 2 follows from Theorem 1 and the preceding Lemma:

Proposition 2. *If the base game G is a sender-receiver game and INTIR is convex, then the set of informed communication equilibria of G coincides with the set of communication equilibria of G .*

Example 3 above is a sender-receiver game with non-convex INTIR where there are CE that are not informed CE. Proposition 2 applies directly to sender-receiver games with *transparent motives*, i.e., games in which the principal's payoff does not depend on their type: $u(a, t) = u(a)$ for all a and t .

Corollary 3. *If the base game G is a sender-receiver game with transparent motives, then the sets of CE and informed CE coincide.*

Proof. If the base game G is a sender-receiver game with transparent motives, we have $(U(t))_t \in \text{INTIR}$ if and only if

$$U(t) \geq \min_{q \in \Delta(T)} \min_{y \in Y(q)} u(y)$$

for every $t \in T$, so INTIR is convex and every CE is an informed CE by Proposition 2. ■

Proposition 2 does not extend to games in which multiple players take actions. For instance, in the two-player game of Example 1, INTIR is convex but some CE payoff vectors are not in INTIR. Similarly, Proposition 2 does not extend to games with a privately informed receiver, even under transparent motives.⁴

⁴Such a counterexample is available upon request.

5.2 Worst outcome and pessimal belief

Moving beyond sender-receiver games, we now identify conditions under which all CE payoff vectors lie in INTIR, ensuring that CE and informed CE coincide. Suppose there exist a pessimal likelihood vector $Q \in \Delta(T)$ and a worst action profile $\underline{a} \in A$ such that the constant outcome that assigns probability one to $\underline{a}$ is a Nash equilibrium outcome of G_Q and

$$u(\underline{a}, t, \theta) \leq u(a, t, \theta) \quad \text{for every } t \in T, a \in A, \theta \in \Theta.$$

Under this condition, all CE outcomes trivially generate payoff vectors in INTIR, so CE and informed CE coincide. To see this, note that for every CE μ and every $t \in T$, we have

$$U(\mu | t) = \sum_{\theta, a} p_0(\theta | t) \mu(a | t, \theta) u(a, t, \theta) \geq \sum_{\theta} p_0(\theta | t) u(\underline{a}, t, \theta),$$

where the right-hand side is the payoff from the worst outcome under Q , placing $(U(\mu | t))_t$ in INTIR.

The assumption of a worst outcome is strong but arises naturally in several important classes of games. Bargaining or peace negotiation games often feature a disagreement outcome such as war or negotiation failure. In such settings, if there exists a likelihood vector Q and a corresponding equilibrium in which all players choose actions leading to the worst outcome (e.g., war) for the principal regardless of type, then every CE is an informed CE. This applies directly to the conflict resolution framework studied by Hörner et al. (2015), where the war outcome provides a uniform lower bound across all type profiles.⁵

6 Which CE are informed CE?

As illustrated in Example 1 and Example 3, some CE yield payoff vectors outside INTIR and therefore are not informed CE. More broadly, correlated devices can facilitate coordinated behavior that hurts the principal, so that equilibrium payoffs under correlation can fall strictly below those attained at a Nash equilibrium.⁶ This is natural in environments where several agents have aligned incentives against the principal and can use correlation to coordinate. In such settings, the CE set may contain payoff vectors that violate interim individual

⁵More generally, coordination games exhibiting strategic complementarities may have extremal Nash equilibria that are worst for the principal. Supermodular games, which include oligopoly games with complementary goods and games with network effects, have a natural order on equilibria and often admit pessimal equilibria.

⁶See, e.g., Rudov et al. (2025).

rationality.

We identify classes of CE that are guaranteed to be informed CE as well. We first study CE that are undominated for the principal, and then turn to CE that can be implemented through a public, unmediated cheap-talk stage from the principal to the agents.

Pareto-frontier and ex-ante optimal CE. A natural candidate is a CE payoff vector on the Pareto frontier for the principal's types. The next observation shows that at least one such payoff vector must be an informed CE payoff vector.

Observation 1. At least one CE payoff vector on the Pareto frontier for the principal's types is an informed CE payoff vector.

Proof. Every Nash equilibrium outcome of the base game G is a CE with payoff vector in INTIR, so the intersection of CE payoff vectors with INTIR is a non-empty set. Since INTIR is an upper set, any CE payoff vector in INTIR is weakly dominated by some Pareto-frontier CE payoff vector that also lies in INTIR. Hence at least one Pareto-frontier CE payoff vector lies in INTIR and is therefore an informed CE by Theorem 1. ■

Observation 1 implies that if the set of Pareto-frontier CE payoff vectors is a singleton, then it corresponds to an ex-ante optimal CE that is also an informed CE. We now show that this property is satisfied in the class of games with generalized transparent motives, defined as follows.⁷

Definition 7 (Generalized transparent motives). A base game G has *generalized transparent motives* if the principal's payoff and beliefs are independent of their own type: for all $t, t' \in T$, $\theta \in \Theta$, and $a \in A$,

$$u(a, t, \theta) = u(a, t', \theta) \quad \text{and} \quad p_0(\theta | t) = p_0(\theta | t').$$

Observation 2. If G has generalized transparent motives, then the CE payoff vector on the Pareto frontier for the principal's types is unique and is an informed CE payoff vector.

Proof. Under generalized transparent motives, the principal's truth-telling constraints at a CE μ imply $U(\mu | t) = U(\mu | t')$ for all $t, t' \in T$. Hence, CE payoff vectors are completely ordered and there is a unique Pareto-frontier CE payoff vector. Therefore, by Observation 1, it is an informed CE. ■

⁷We thank a referee for suggesting this result in games with generalized transparent motives.

Moving beyond games with generalized transparent motives, are all CE payoff vectors on the principal's Pareto frontier informed CE payoff vectors? Is the ex-ante optimal CE necessarily an informed CE? In general the answer is no, even in sender-receiver games.⁸ The next example gives a binary-action game with three uninformed agents in which the principal's best informed CE yields a strictly lower ex-ante payoff than the ex-ante optimal CE, even though INTIR is convex. In that example, the unique informed CE payoff vector lies on the principal's Pareto frontier by Observation 1, while the ex-ante optimal CE payoff vector lies outside INTIR (see Figure 2).

Example 4 (Ex-ante optimal CE not informed CE). The principal considers promoting a new technology to three agents, each of whom chooses whether to adopt the new technology (action 1) or stick with their current technology (action 0). The principal holds private information that determines the suitability of adoption for each agent and the principal's payoff from each adoption configuration. The principal's type space is binary, $T = \{t_1, t_2\}$, with prior $p(t_1) = \frac{3}{4}$.

The principal's payoff is:

$$u(a_1, a_2, a_3, t_1) = \begin{cases} 3 & \text{if } a_1 = 1 \text{ and } (a_2, a_3) \neq (1, 0), \\ 0 & \text{if } a_1 = 0 \text{ for all } (a_2, a_3), \\ -c & \text{if } (a_1, a_2, a_3) = (1, 1, 0), \end{cases}$$

$$u(a_1, a_2, a_3, t_2) = \begin{cases} 0 & \text{if } a_1 = 1 \text{ for all } (a_2, a_3), \\ 1 & \text{if } a_1 = 0 \text{ and } (a_2, a_3) \neq (1, 0), \\ -c & \text{if } (a_1, a_2, a_3) = (0, 1, 0), \end{cases}$$

where $c > 0$.

The principal of type t_1 receives a payoff of 3 when agent 1 adopts and 0 when agent 1 does not adopt, regardless of what the other agents do—except in the case where all but agent 3 adopt, i.e., when the action profile is $(1, 1, 0)$. In that case, the principal incurs a loss of $-c$ because the technology is precarious for principal type t_1 when agents 1 and 2 adopt without agent 3 adopting. The principal of type t_2 receives a payoff of 0 when agent 1 adopts and 1 when agent 1 does not adopt, regardless of the actions of the other agents—except when only agent 2 adopts, i.e., when the action profile is $(0, 1, 0)$. In that case, the principal incurs a loss of $-c$ because the technology is precarious for principal type t_2 when only agent 2 adopts, without agents 1 or 3 adopting.

⁸We are very grateful to Daniel Clark for providing a three-type, twenty-action sender-receiver counterexample to one of our previous conjectures. Details are available upon request.

Let q denote the common belief assigned to t_1 . Each agent acts solely based on their belief q that the principal is of type t_1 . Assume each agent's payoff is such that agent 1 strictly prefers to adopt if and only if $q < \frac{1}{2}$, agent 2 if and only if $q > \frac{3}{4}$, and agent 3 if and only if $q > \frac{7}{10}$. The induced adoption profile is:

$$\begin{cases} (1, 0, 0) & \text{if } 0 \leq q < \frac{1}{2}, \\ (0, 0, 0) & \text{if } \frac{1}{2} < q < \frac{7}{10}, \\ (0, 0, 1) & \text{if } \frac{7}{10} < q < \frac{3}{4}, \\ (0, 1, 1) & \text{if } \frac{3}{4} < q \leq 1. \end{cases}$$

Crucially, the action profiles $(1, 1, 0)$ and $(0, 1, 0)$, which generate a payoff of $-c$ for the principal, are never played with positive probability regardless of q . A vector $(U(t_1), U(t_2))$ is in INTIR if and only if $U(t_1) \geq 0$, $U(t_2) \geq 0$, and $U(t_2) \geq 1 - \frac{1}{3}U(t_1)$, as illustrated in Figure 2. Note that INTIR is convex.




Figure 2: Payoff sets in Example 4. The set of CE payoff vectors is the blue area. The ex-ante optimal CE payoff vector at $(1, 0)$ lies outside INTIR (gray area) and is therefore not an informed CE, despite INTIR being convex.

Without a communication device, i.e., if agents act solely based on their prior $p(t_1) = \frac{3}{4}$, the equilibrium action profiles are $(0, 1, 1)$ and $(0, 0, 1)$, yielding a payoff of 0 for the principal under t_1 and 1 under t_2 , for an expected payoff of $\frac{1}{4}$. By contrast, a communication device allows the principal to fine-tune personalized recommendations based on their reported type.

The ex-ante optimal CE payoff vector is unique and given by $(1, 0)$. To see this, first

observe that such a payoff is obtained with the following CE:

$$\mu((1,1,1) | t_1) = \frac{1}{3}, \quad \mu((0,1,0) | t_1) = \frac{2}{3}, \quad \mu((1,1,0) | t_2) = 1.$$

If the principal lies, reporting t_2 when it is t_1 , the outcome $(1,1,0)$ results in a $-c$ payoff. Similarly, if t_2 pretends to be t_1 , with probability $\frac{2}{3}$ the outcome $(0,1,0)$ is realized, also triggering a $-c$ payoff. Hence, μ satisfies the principal's truth-telling constraints. Agents' obedience also holds. Agent 1 believes t_1 with probability $\frac{1}{2}$ when recommended $a_1 = 1$, and more than $\frac{1}{2}$ when recommended $a_1 = 0$, consistent with their threshold. Agent 2 believes t_1 with probability $\frac{3}{4}$ when recommended $a_2 = 1$, which equals their cutoff, so obedience holds (weakly). Agent 3, when recommended $a_3 = 1$, is certain of t_1 ; when recommended $a_3 = 0$, the belief is $\frac{2}{3}$, which is below their threshold $\frac{7}{10}$, validating their obedience as well.

Second, consider a fictitious relaxed information design problem of the principal, assuming the principal can choose an information structure ex-ante and can fully control the actions of agents 2 and 3 conditional on the type of the principal and the action of agent 1. In this relaxed problem, agent 1 plays action $a_1 = 1$ if $q < \frac{1}{2}$ and action $a_1 = 0$ if $q > \frac{1}{2}$. The maximum payoff vector of the principal is therefore $(3,0)$ if $q < \frac{1}{2}$ and $(0,1)$ if $q > \frac{1}{2}$. The value function of the principal in this relaxed problem is:

$$\begin{cases} 3q & \text{if } q \leq \frac{1}{2}, \\ 1 - q & \text{if } q > \frac{1}{2}. \end{cases}$$

Using standard concavification techniques from Bayesian persuasion, we obtain that the unique solution to this relaxed problem is to split the prior $p = \frac{3}{4}$ into the posteriors $q = \frac{1}{2}$ and $q = 1$ for agent 1. The induced payoff vector for the principal is $(1,0)$, the same as the CE payoff vector above. Hence, this payoff vector is the unique ex-ante optimal CE payoff vector. As can be seen in Figure 2, it is not in INTIR, so it is not an informed CE by Theorem 1, even though INTIR is convex.

This example illustrates how a principal can exploit communication devices to optimally leverage heterogeneity in adoption thresholds. While the coordination value of such devices has been well understood since Myerson (1982), Forges (1986), and Myerson (1991), it is less clear whether these gains are attainable when the principal chooses the device after learning their type. In fact, Yilankaya (1999), Maskin and Tirole (1990), and Mylovanov and Tröger (2014), among others, underscore that ex-ante optimal mechanisms are equilibrium outcomes of the informed-principal game. By contrast, this example shows that the ex-ante optimal CE assigns a very low payoff to type t_2 , so a principal of type t_2 deviates by choos-

ing silence. Thus, the ex-ante optimal CE need not be an informed CE.

In the previous example, the ex-ante optimal CE cannot be obtained with cheap talk from the principal to the uninformed agents. We proceed to establish that for every base game G all CE arising through public cheap talk from the principal to the agents are informed CE. In particular, this implies that if an ex-ante optimal CE can be implemented through public cheap talk from the principal to the agents, then it is an informed CE.

Unmediated public communication. We show that every equilibrium outcome of the base game G augmented with *public, unmediated* (cheap-talk) communication from the principal to the agents yields a principal payoff vector in INTIR, and hence constitutes an informed CE. An immediate implication is that whenever the ex-ante optimal CE can be implemented through unmediated public cheap talk, it is necessarily an informed CE. It is important to emphasize that Proposition 3 relies on the assumption that the only communication channel is public and unmediated from the principal to the players.

Proposition 3. *If μ is a Nash equilibrium outcome of the base game G augmented by public unmediated communication from the principal, then μ is an informed communication equilibrium.*

Proof. Public unmediated communication corresponds to a (non-direct) communication device v^{PU} with $|R_i| = 1$ for all $i \in I$, $M_i = R_0$ for all $i \in I_0$, and $v^{PU}((r_0, \dots, r_0) \mid r_0, r_{-0}) = 1$ for every $r_0 \in R_0$. In what follows we drop the subscript 0. Consider a Nash equilibrium of $G(v^{PU})$. Let $\tau : T \rightarrow \Delta(R)$ denote the principal's reporting strategy at that equilibrium, and let μ be the corresponding outcome. For every $r \in \text{supp } \tau$, let $Q_r \in \Delta(T)$ be defined by $Q_r(t) = \tau(r \mid t) / (\sum_{\tilde{r} \in T} \tau(r \mid \tilde{r}))$. Observe that, for every $t \in T$,

$$U(\mu \mid t) = \sum_{r \in R} \tau(r \mid t) U(\mu_r \mid t),$$

where, for every $r \in \text{supp } \tau$, μ_r is a Nash equilibrium outcome of G_{Q_r} . Because τ is an equilibrium reporting strategy of $G(v^{PU})$ we have:

$$U(\mu \mid t) \geq U(\mu_r \mid t) \quad \forall t \in T, \forall r \in \text{supp } \tau.$$

Now pick any $r \in \text{supp } \tau$ and let $Q = Q_r$ and $\underline{\mu} = \mu_r$. We have $U(\mu \mid t) \geq U(\underline{\mu} \mid t)$ for every t , and $\underline{\mu}$ is a Nash equilibrium of G_Q , so $(U(\underline{\mu} \mid t))_t \in \text{INTIR}$. Since μ is a CE and $(U(\mu \mid t))_t \in \text{INTIR}$, Theorem 1 implies that μ is an informed CE. ■

Proposition 3 implies that every outcome of the unmediated cheap talk version of a sender-receiver game à la Crawford and Sobel (1982) lies in INTIR. Hence, in such games, if the ex-ante optimal CE can be achieved via such unmediated cheap talk then it is an informed CE. Kamenica and Lin (2024) identify conditions under which, in a sender-receiver game, ex-ante commitment to an experiment has no value for the principal relative to unmediated cheap talk. Specifically, they show that in finite sender-receiver games and generically over payoffs and priors, the principal attains the optimal persuasion payoff via cheap talk if and only if there exists a deterministic optimal experiment. Hence, in such finite and generic sender-receiver games, Proposition 3 implies that if the optimal experiment is deterministic, then the corresponding outcome is an informed CE.

References

- AUMANN, R. J. (1974): “Subjectivity and Correlation in Randomized Strategies,” *Journal of Mathematical Economics*, 1, 67–96.
- (1987): “Correlated Equilibrium as an Expression of Bayesian Rationality,” *Econometrica*, 55, 1–18.
- AUMANN, R. J. AND S. HART (2003): “Long Cheap Talk,” *Econometrica*, 71, 1619–1660.
- ĀZACIS, H., M. LACLAU, AND P. VIDA (2025): “Unmediated communication in games with (in)complete information: the 4-player case,” *Working paper*, halshs-04895364v3.
- BERGEMANN, D. AND S. MORRIS (2016): “Bayes correlated equilibrium and the comparison of information structures in games,” *Theoretical Economics*, 11, 487–522.
- BLUME, A. AND O. BOARD (2013): “Language Barriers,” *Econometrica*, 81, 781–812.
- CARMONA, G. AND K. LAOHAKUNAKORN (2026): “Privately designed correlated equilibrium,” *Games and Economic Behavior*, 157, 237–252.
- CHAKRABORTY, A. AND R. HARBAUGH (2010): “Persuasion by Cheap Talk,” *American Economic Review*, 100, 2361–2382.
- CLARK, D. (2022): “The informed principal with agent moral hazard,” *Working Paper*.
- (2023): “Contracting with private information, moral hazard, and limited commitment,” Tech. rep., Working Paper.
- CORRAO, R. (2023): “Mediation Markets: The Case of Soft Information,” Tech. rep., Working Paper.
- CORRAO, R. AND Y. DAI (2023): “The Bounds on Mediated Communication,” Tech. rep., Working Paper.

- CRAWFORD, V. P. AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50, 1431–1451.
- DE CLIPPEL, G. AND E. MINELLI (2004): “Two-person bargaining with verifiable information,” *Journal of Mathematical Economics*, 40, 799–813.
- FORGES, F. (1986): “An Approach to Communication Equilibria,” *Econometrica*, 54, 1375–1385.
- (1990a): “Equilibria with Communication in a Job Market Example,” *Quarterly Journal of Economics*, 105, 375–398.
- (1990b): “Universal Mechanisms,” *Econometrica*, 58, 1341–1364.
- (1993): “Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information,” *Theory and Decision*, 35, 277–310.
- FUDENBERG, D. AND J. TIROLE (1991): “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory*, 53, 236–260.
- GANGULY, C. AND I. RAY (2023): “Simple mediation in a cheap-talk game,” *Games*, 14, 47.
- GERARDI, D. AND R. B. MYERSON (2007): “Sequential equilibria in Bayesian games with communication,” *Games and Economic Behavior*, 60, 104–134.
- GOLTSMAN, M., J. HÖRNER, G. PAVLOV, AND F. SQUINTANI (2009): “Mediation, Arbitration and Negotiation,” *Journal of Economic Theory*, 144, 1397–1420.
- HÖRNER, J., M. MORELLI, AND F. SQUINTANI (2015): “Mediation and peace,” *The Review of Economic Studies*, 82, 1483–1501.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KAMENICA, E. AND X. LIN (2024): “Commitment and Randomization in Communication,” *arXiv preprint arXiv:2410.17503*.
- KOESSLER, F. AND V. SKRETA (2016): “Informed Seller with Taste Heterogeneity,” *Journal of Economic Theory*, 165, 456–471.
- (2023): “Informed information design,” *Journal of Political Economy*, 131, 3186–3232.
- (2025): “Informed communication equilibrium,” *hal-05124399*.
- LIPNOWSKI, E. AND D. RAVID (2020): “Cheap talk with transparent motives,” *Econometrica*, 88, 1631–1660.
- MASKIN, E. AND J. TIROLE (1990): “The principal-agent relationship with an informed principal: The case of private values,” *Econometrica*, 379–409.
- (1992): “The principal-agent relationship with an informed principal, II: Common values,” *Econometrica*, 1–42.

- MEKONNEN, T. (2021): “Informed principal, moral hazard, and limited liability,” *Economic Theory Bulletin*, 9, 119–142.
- MITUSCH, K. AND R. STRAUZ (2005): “Mediation in Situations of Conflict and Limited Commitment,” *Journal of Law, Economics, and Organization*, 21, 467–500.
- MYERSON, R. B. (1982): “Optimal Coordination Mechanisms in Generalized Principal-Agent Problems,” *Journal of Mathematical Economics*, 10, 67–81.
- (1983): “Mechanism design by an informed principal,” *Econometrica*, 1767–1797.
- (1986): “Multistage Games with Communication,” *Econometrica*, 54, 323–358.
- (1991): *Game Theory, Analysis of Conflict*, Harvard University Press.
- MYLOVANOV, T. AND T. TRÖGER (2012): “Informed principal problems in generalized private values environments,” *Theoretical Economics*, 7, 465–488.
- (2014): “Mechanism Design by an Informed Principal: Private Values with Transferable Utility,” *The Review of Economic Studies*, 81, 1668–1707.
- (2025): “Neo-Optimum: A Unifying Solution to the Informed-Principal Problem,” *mimeo*.
- RUDOV, K., F. SANDOMIRSKIY, AND L. YARIV (2025): “Extreme Equilibria: The Benefits of Correlation,” in *Proceedings of the 26th ACM Conference on Economics and Computation*, 897–897.
- SALAMANCA, A. (2021): “The value of mediated communication,” *Journal of Economic Theory*, 192, 105191.
- (2024): “Biased Mediators in Conflict Resolution,” *American Law and Economics Review*, ahead014.
- (2025): “Mediation Design by an Informed Party,” *Working paper*.
- SUGAYA, T. AND A. WOLITZKY (2021): “The revelation principle in multistage games,” *The Review of Economic Studies*, 88, 1503–1540.
- TANEVA, I. (2019): “Information design,” *American Economic Journal: Microeconomics*, 11, 151–185.
- WAGNER, C., T. MYLOVANOV, AND T. TRÖGER (2015): “Informed-principal problem with moral hazard, risk neutrality, and no limited liability,” *Journal of Economic Theory*, 159, 280–289.
- YILANKAYA, O. (1999): “A note on the seller’s optimal mechanism in bilateral trade with two-sided incomplete information,” *Journal of Economic Theory*, 87, 267–271.