The Structure of Joint Posterior Beliefs*

David Lagziel[†] Ehud Lehrer[‡]

March 25, 2026

Abstract

We analyze information structures in environments with privately informed agents and a partially informed mediator who provides public signals. Necessary and sufficient conditions, termed *internal* and *external consistency*, are derived for the rationalizability of joint posteriors. Analyzing the geometry of the feasible set of posteriors shows that informational constraints alone can generate non-convexity, thereby limiting the applicability of standard concavification techniques. We also characterize the convexity of the feasible set in the single-agent setting. Applications to belief depolarization, information design, Bayesian persuasion, and potential games are discussed.

Journal of Economic Literature classification numbers: C72, D82, D83.

Keywords: joint posteriors, posterior likelihood, graph of information, internal consistency, external consistency, convexity of posteriors

*For their valuable comments, the authors wish to thank participants of the University of Toronto Economic Theory Workshop, the Durham University Economics Seminar, the Adam Smith Business School Micro theory seminar of Glasgow University, INSEAD EPS seminar, the Western University Economics seminar, the Tel Aviv University Game Theory Seminar and the Rationality Center Game Theory Seminar. Lagziel acknowledges the support of the Israel Science Foundation, Grant #2074/23. Lehrer acknowledges the support of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project Number 461570745.

†Department of Economics, Ben-Gurion University of the Negev, Beer-Sheba 8410501, Israel. E-mail: davidlag@bgu.ac.il. Visiting Professor, Department of Economics, University of Toronto, 2025–2027.

‡Economics Department, Durham University, Durham DH1 3LB, UK. E-mail: ehud.m.lehrer@durham.ac.uk.

1 Introduction

What is the role of a mediator in reaching an agreement, such as a peace treaty, between two opposing sides? A natural objective is to persuade both parties to accept the agreement. Yet this form of persuasion is subtle as the mediator must shape a *joint perception* shared by both sides, ensuring that each finds the agreement acceptable. Our paper begins with this observation.

We study a setting with privately informed players and a partially informed external information provider (a *mediator*), all defined on a finite state space with a common prior. Each player and the mediator are endowed with a partition of the state space. The mediator sends a *public* signal according to a stochastic rule that is measurable with respect to its information. Upon observing the signal, players update their beliefs by Bayes' rule, generating a *joint posterior profile* across players (and across states at which the signal can occur). Our central question is a feasibility one: given a joint posterior profile, when does there exist a mediator-compatible public signal that induces it? We answer this by providing necessary and sufficient conditions for implementability.

This implementability concern is central in economics and game theory, where Bayesian updating is the cornerstone of belief revision under uncertainty. In strategic environments with incomplete information, players use signals to update their beliefs about the underlying state of the world, often relying on publicly observed information. While it might be possible to construct posterior beliefs through arbitrary signaling mechanisms, we focus on the more subtle problem of whether such posteriors can be implemented under the restriction that signals must be measurable with respect to the mediator's limited knowledge. Moreover, the mediator's partition of the state space may differ from those of the players, creating the novel possibility that players have common knowledge of events that the mediator does not know.

Our main tool is the *graph of information*, following the framework of Rodrigues-Neto (2009), which encodes the state space and players' informational partitions in a concise form. On this graph, we define a *posterior likelihood function* φ that specifies the ratio of posterior probabilities across *adjacent states* (i.e., states lying in the same information set of some player), thus providing a compact representation of Bayesian updating. We then ask under what conditions such a function can be rationalized by a mediator-generated signal.

While the graph provides a convenient representation of informational constraints, the core of our analysis lies in understanding how these constraints shape the implementability of belief updates. The function φ is particularly valuable in our set-up because the likelihood ratios it encodes are shared by all players who cannot distinguish between adjacent states. As a result, the information carried by φ is adequate for our analysis, rendering the graph-based model both concise and sufficiently informative.

Our main theorem shows that a joint posterior profile is implementable by a mediator's public signal if and only if these likelihood ratios satisfy two multiplicative consistency conditions: *internal consistency* (a cycle condition within each common knowledge component) and *external consistency* (a loop condition generated by the mediator pooling states across components). These conditions mirror and extend earlier consistency notions in the literature (notably Rodrigues-Neto (2009) and Hellman and Samet, 2012), while introducing the role of a third-party mediator.

We then use the aforementioned characterization to study the geometry of the feasible set, showing that it *need not be convex* even in the absence of incentive constraints or restrictions on the number of signals, thus limiting the direct applicability of concavification arguments. A characterization of convexity of the feasible set in the single-player case is provided, as well. The restrictive nature of convexity that arises from this characterization is rather striking, so that even in standard persuasion models where receivers have a private type (as in, e.g., Kolotilin et al. (2017) and Guo and Shmaya (2019), among many others), convexity fails. Finally, we show how the same graph-based tools inform problems in information design and belief depolarization, and provide an additive analogue relevant for potential games.

Our results provide a bridge between abstract information structures and concrete Bayesian updating. We show that φ can arise from Bayesian posteriors induced by a signal that the mediator releases. This opens the door to interpreting the mediator as a generator of Blackwell experiments and leads to new insights about the implementability of distributions over states and the coherence of beliefs across players. In doing so, we extend the theory of common priors and beliefs, clarify the conditions under which players' posteriors can be coherently derived from shared signals, and provide a graph-theoretic approach to understanding the flow of information in multi-agent systems.

A Negotiation Game: motivating example. To motivate our model and results, consider a game with two players, indexed by $i = 1, 2$. Each player has two available actions: *attack* (denoted A) and *compromise* (denoted C). The set of states is $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, endowed with a uniform common prior. The players hold asymmetric information: player 1's partition is $\mathcal{P}_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$, and player 2's partition is $\mathcal{P}_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3, \omega_4\}\}$. Thus, when either ω_1 or ω_2 is realized, player 2 learns the state with probability 1, while player 1's posterior is uniform over $\{\omega_1, \omega_2\}$. Similarly, when ω_3 or ω_4 is realized, then player 1 learns the state with probability 1, while player 2's posterior is uniform over $\{\omega_3, \omega_4\}$. This example bears some resemblance to the framework of Hörner et al. (2015), which was extended by Özyurt and Zeng (2025) to incorporate a privately informed mediator into the original game.

The payoffs of the game are presented in Figure 1. Let G_i denote the payoff matrix when state ω_i is realized. In G_1 and G_3, player 1's dominant strategy is A and player 2's dominant strategy is C, whereas in G_2 and G_4 the dominant strategies are reversed.

		Player 2 A	**Player 2** C
Player 1	A	$(2,-5)$	$(2,-4)$
Player 1	C	$(-1,-1)$	$(0,0)$

Given that $\omega \in \{\omega_1, \omega_3\}$

		Player 2 A	**Player 2** C
Player 1	A	$(-5,2)$	$(-1,-1)$
Player 1	C	$(-4,2)$	$(0,0)$

Given that $\omega \in \{\omega_2, \omega_4\}$

Figure 1: Two payoff matrices with Player 1 (row player) and Player 2 (column player).

The interpretation of the game is straightforward. Each player can either attack the other or compromise by agreeing to a peace treaty. In states ω_1 and ω_3, player 1 holds the superior attacking position, whereas in states ω_2 and ω_4 player 2 holds this advantage.[1] Nevertheless, from a social perspective it is optimal to reach a peace agreement, as it maximizes the aggregate payoff. Such a treaty, however, requires a joint concession by both players.

Consider now the players' equilibrium behavior. If the realized state is either ω_1 or ω_2, then player 2 is fully informed while player 1's posterior is $\left(\frac{1}{2}, \frac{1}{2}, 0, 0\right)$. In this case, given that player 2's dominant strategies are C in ω_1 and A in ω_2, player 1's optimal action is A. On the other hand, if the realized state is either ω_3 or ω_4, then player 1 is fully informed while player 2's posterior is $\left(0, 0, \frac{1}{2}, \frac{1}{2}\right)$, and player 2's optimal action is A. Hence, in every equilibrium of the game, the action profile (C, C) is never played.

We now introduce the role of a mediator. The mediator also possesses private information, represented by the partition $F = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. Figure 2 depicts the players' and the mediator's information structures. The mediator's task is to design a signaling mechanism that persuades both players to accept a peace treaty in equilibrium, with positive probability. Such persuasion, however, can never succeed with probability 1. Namely, whenever a player is fully informed and holds the superior attacking position (i.e., in states ω_2 or ω_3), that player will necessarily choose A.

If the mediator fully discloses its information (i.e., reveals the realized F-cell), both players become perfectly informed about the state, and the peaceful outcome (C, C) is never

[1] A joint attack leads to the worst aggregate outcomes, $(2,-5)$ or $(-5,2)$, though still favorable for the player with the superior position. If one player attacks while the other compromises, payoffs reflect the attacking player's advantage.

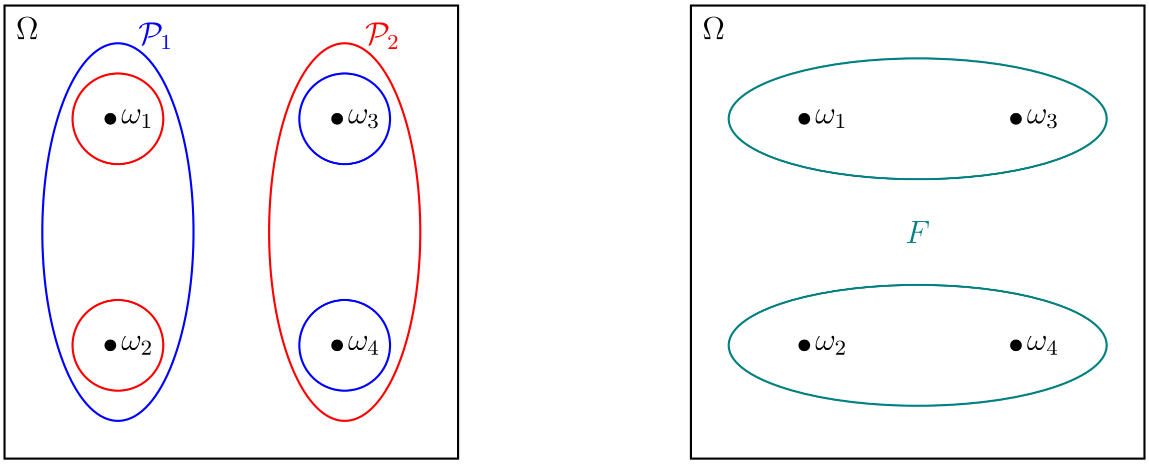

(a) The players' information

(b) The mediator's information

Figure 2: On the left, Figure (a) illustrates the information structure of player 1 (blue) and player 2 (red). On the right, Figure (b) portrays the information structure of the mediator (teal).

sustained in equilibrium. By contrast, suppose the mediator could induce the posterior $\left(\frac{1}{3}, \frac{2}{3}, 0, 0\right)$ for player 1 while player 2 remains fully informed. Then player 1 is indifferent between A and C: action A yields $\frac{1}{3}\cdot 2+\frac{2}{3}\cdot(-5) = -\frac{8}{3}$, and action C yields $\frac{1}{3}\cdot 0+\frac{2}{3}\cdot(-4) = -\frac{8}{3}$. Hence, there exists an equilibrium in which (C, C) is played with positive probability. A similar argument applies to player 2. If the mediator induces the posterior $\left(0, 0, \frac{2}{3}, \frac{1}{3}\right)$ for player 2 while player 1 remains fully informed, then player 2 is indifferent between A and C. Once again, this creates an equilibrium in which a peace agreement (the socially optimal outcome) is reached with positive probability.

This raises a natural question: can the mediator design a signaling strategy that simultaneously induces these posteriors? For instance, is it possible to generate a posterior $\left(\frac{1}{3}, \frac{2}{3}, 0, 0\right)$ for player 1 when the realized state lies in $\{\omega_1, \omega_2\}$, while at the same time generating a posterior $\left(0, 0, \frac{2}{3}, \frac{1}{3}\right)$ for player 2 when the realized state lies in $\{\omega_3, \omega_4\}$? In this example, the answer is negative. More generally, this question captures the central problem of the paper.

Why, in this specific example, does the mediator fail to generate the stated joint posterior profile? The obstacle is measurability: any signaling rule based on the partition F must assign the same signal distribution in ω_1 and ω_3, and likewise in ω_2 and ω_4. Consequently, once the mediator induces a posterior of the form $(p, 1-p, 0, 0)$ for player 1, this necessarily translates into a posterior of $(0, 0, p, 1-p)$ for player 2, and vice versa.

The Characterization. We provide necessary and sufficient conditions for a mediator to implement a feasible joint posterior, given the players' information and the mediator's partition. This characterization captures the inherent limitations of such implementation, expressed in terms of likelihood ratios across specific states.[2]

To illustrate, consider the sequence $(\omega_1, \omega_2, \omega_4, \omega_3)$. The event $\{\omega_1, \omega_2\}$ forms a *common knowledge component* (CKC), that is, a minimal set on which all players can agree.[3] Hence, ω_1 and ω_2 lie in the same CKC and are linked through the players' knowledge. In contrast, ω_2 and ω_4 are not in the same CKC but belong to the same information set of the mediator. Next, ω_4 and ω_3 again fall within the same CKC. Finally, ω_3 belongs to the same mediator information set as ω_1. Thus, the sequence $(\omega_1, \omega_2, \omega_4, \omega_3)$ alternates between steps within a CKC (via the players' information) and steps across CKCs induced by the mediator's coarser information. This sequence is an example of an *F-loop* (see Section 2.4), in which the mediator's information cyclically connects distinct CKCs of the players.

Our characterization establishes that, for a joint posterior to be feasible, the product of likelihood ratios along every loop must equal one. For instance, under the posterior $\left(\frac{1}{3}, \frac{2}{3}, 0, 0\right)$ the likelihood ratio of ω_1 to ω_2 is $\varphi(\omega_1, \omega_2) = \frac{1/3}{2/3} = \frac{1}{2}$. Under the posterior $\left(0, 0, \frac{2}{3}, \frac{1}{3}\right)$ the likelihood ratio of ω_4 to ω_3 is $\varphi(\omega_4, \omega_3) = \frac{1/3}{2/3} = \frac{1}{2}$. In this example, the feasibility condition fails for the proposed posteriors, but it does hold for $(p, 1-p, 0, 0)$ and $(0, 0, p, 1-p)$. This property, referred to later as *external consistency*, is a new constraint created by mediator measurability, and it underpins our general characterization.

The second key property in our characterization, termed *internal consistency*, imposes a similar condition within each CKC of the players. To formalize this, we use the notion of a *F-cycle*: a closed path of states contained in a CKC, where each pair of adjacent states is connected by an atom of one of the players' partitions, except for a single pair connected by an atom of the mediator's partition (denoted by F). Internal consistency requires that the product of likelihood ratios along every such F-cycle equals one.

Beyond our feasibility characterization, an additional and intriguing question concerns the optimality of the mediator's signaling function. Suppose that the mediator seeks to maximize the probability that the outcome (C, C) is played. To this end, the mediator can employ multiple signals to generate feasible posteriors of the forms $(p, 1-p, 0, 0)$ and $(0, 0, p, 1-p)$ for various values of p, each inducing the desired profile with different probabilities. In a broader setting, however, this becomes a more complex problem: a necessary preliminary step is to identify which posteriors are feasible. In other words, before addressing the question of optimality, one must first determine the set of feasible posteriors, and this is the focus of the current research.

[2]Notably, the likelihood ratio also plays a central role in the negotiation game of Hörner et al. (2015); see the definition of λ therein and the subsequent analysis.

[3]Formally, each player's information generates this event; equivalently, a CKC is a minimal non-empty subset that is measurable with respect to the σ-algebra of all players.

Our final result elaborates on this point and provides an illustrative example. Specifically, we examine the properties of the feasible set of joint posterior beliefs. A key question is whether this set is convex, since convexity plays a crucial role in concavification arguments as in Aumann and Maschler (1995) and Kamenica and Gentzkow (2011). Using our characterization, we demonstrate that the set of joint posterior beliefs *need not be convex*, and we characterize such convexity in the single-player set-up. We prove that the feasible set is convex if and only if every F-loop is *non-informative*, meaning that it comprises a unique partition element of F.[4] Although similar forms of non-convexity have been reported in the literature,[5] our result is fundamentally different, as it arises solely from informational (i.e., measurability) constraints, rather than from limitations on the number of signals, incentive-compatibility constraints, or other exogenous restrictions.

Related literature. Several strands of the literature inform our analysis. Since the foundational work of Harsanyi (1967–1968), much effort has been devoted to formalizing and understanding the informational structures that underlie such environments. A key insight in this literature is that players' information can be represented as partitions over a finite state space, and the relationships among these partitions encode the possible flow and structure of information in the game. Harsanyi (1967–1968) provides the basis for belief-based reasoning in strategic settings.

Aumann's framework for knowledge and common knowledge (see Aumann, 1976) formalized the use of partitions to represent agents' information. Building on this foundation, our paper extends the idea by introducing an external mediator that generates signals constrained by its own partition. A central concept in our analysis is the CKC, which is rooted in Aumann's original formulation.

The question of whether a joint posterior originates from a *common prior* has attracted significant attention across various settings. This inquiry has been studied extensively in contexts ranging from finite state spaces (see, e.g., Morris, 1994), to compact state spaces (e.g., Feinberg, 2000; Heifetz, 2006), and to countable state spaces (see Lehrer and Samet, 2014). In the present framework, we assume that players initially share a common prior and acquire information through their individual partitions of the state space, subsequently updating their beliefs via Bayesian conditioning.

The novel aspect of our model is the introduction of an external information source, the mediator. Given the observed joint posterior profile of the players, we inquire whether such a profile can be rationalized as arising from a known common prior, augmented by addi-

[4]The notion of non-informative loops is used in Lagziel et al. (2025) to compare information providers.

[5]See, e.g., Kosenko (2020); Ball and Espín-Sánchez (2021); Candogan and Strack (2023); Herings et al. (2024); Chen et al. (2025); see broader discussion in the literature review.

tional public information disseminated by the mediator, which itself may be only partially informative to the players. In other words, we ask: *does the joint posterior stem from the interaction between a common prior and an external, symmetric informational input?*

The conditions we impose, internal and external consistency, are inspired by the cycle-based consistency concepts introduced by Rodrigues-Neto (2009) and Hellman and Samet (2012), though our focus is on implementability via constrained signals rather than belief structures.

At a conceptual level, our mediator corresponds to a restricted Blackwell experiment Blackwell (1951), with implementability shaped by the mediator's coarse knowledge. In this sense, our work complements the Bayesian persuasion framework of Kamenica and Gentzkow (2011), where a sender optimally selects signals to influence beliefs; we instead ask when a given belief structure can be realized at all under informational constraints.

Motivated by the question of aggregating experts' opinions and building on the results of Kellerer (1961), Strassen (1965), and Gutmann et al. (1991), the work of Dawid et al. (1995) provides a characterization of joint posterior beliefs for the case of two agents. This was extended by Arieli et al. (2021), who characterize the set of feasible distributions of joint posterior beliefs that can arise among multiple agents in a binary state space, given a Blackwell experiment that provides different (potentially asymmetric) information to each agent. Their characterization is closely related to the no-trade literature in that it identifies constraints on belief distributions that are consistent with a common prior and Bayesian updating, even when agents receive heterogeneous private signals.[6] More recently, Herings et al. (2024) studied a similar question, but derived a different characterization while accounting for an arbitrary finite state space.

Independently of Arieli et al. (2021), Ziegler (2020) studies a setting in which a mediator provides private signals to two receivers without committing to a common information structure. The mediator therefore chooses information structures robustly, maximizing expected payoff against the worst-case interpretation consistent with Bayesian rationality. In doing so, Ziegler derives necessary feasibility constraints on the joint distribution of receivers' posteriors, conditions that coincide with those of Arieli et al. (2021) as both necessary and sufficient in the case of two agents.

Our analysis distinguishes itself from existing literature in four fundamental ways, which collectively define a novel research agenda. First, we adopt a setting in which players receive private signals from a fixed information structure regarding an unknown state that is not

[6] See also the follow-up paper and review of the no-trade history in Morris (2020), as well as the study of Burdzy and Pitman (2020), which follows Dawid et al. (1995), to derive probabilistic bounds on the polarization of posteriors in a two-agent setting.

necessarily binary, allowing for a richer initial belief space. Second, the mediator is only partially informed about the true state, and their own information structure is fixed, which shifts the focus from optimal information design to the mechanics of communication under given constraints. Our characterization is thus framed in terms of the information structures of both the players and the mediator. Third, the mediator utilizes a public communication channel, a deliberate simplification that allows us to focus on the impact of generating common knowledge, rather than the complexities of personalized private signals. Finally, and most critically, our primary focus is on characterizing the specific joint posteriors that can be induced by the mediator, rather than the distribution over these posteriors. We then build upon this to define and characterize the entire set of implementable joint posteriors, investigating the conditions under which any element of this set can be implemented through a stochastic public signal, a major departure from the standard distributional analysis in this field.

These differences become most apparent in our final result, which shows that convexity (of the set of joint posterior beliefs) can fail purely due to information constraints. While instances of non-convexity are well known, they typically stem from different sources. For example, Herings et al. (2024) obtain non-convexity of feasible posterior distributions under limited message alphabets, whereas our framework imposes no such restriction. In Candogan and Strack (2023), non-convexity arises from incentive-compatibility constraints with privately informed receivers, while our analysis is entirely payoff-free. Chen et al. (2025) demonstrate non-convex feasible regions under support-size caps (K-signaling), whereas in our setting, constraints emerge from likelihood-ratio products on F-loops, without any bound on the support. Kosenko (2020) visualize non-convex feasible-posterior shapes under garbling, while we provide necessary and sufficient graph-based conditions, identifying precisely when such non-convexity arises from a partially informed mediator and privately informed players. Finally, Ball and Espín-Sánchez (2021) establish non-convexity when the set of experiments is restricted, whereas in our framework, the restriction is due to the measurability of kernels given the mediator's information.

Our characterization also clarifies when convexity can and cannot be expected. In Section 4 we provide a sharp convexity characterization for the single-player case, and we show by example that non-convexity can persist in richer environments, including settings with multiple privately informed players, even when the mediator is fully informed. Together, these findings build the case that non-convexity can arise purely from informational constraints, independent of payoffs, support, or signal limitations.

Finally, our results contribute to the broader literature on higher-order beliefs (see the review by Geanakoplos, 1994), by characterizing which belief patterns, encoded in posterior likelihoods, can emerge from shared public signals constrained by a third party's limited information.

The structure of the paper. The paper is organized as follows. Section 2 presents the model and defines joint beliefs, posterior likelihoods, the graph of information, and the notions of cycles and loops. Section 3 introduces internal and external consistency and proves the key graph-theoretic result (Proposition 1) underlying our main characterization. Section 3.2 states the main implementability result (Theorem 1). Section 4 studies the convexity of the feasible set. Section 5 discusses extensions and applications, including multiple signals, persuasion, polarization, and potential games.

2 Preliminaries

Let $N = \{1, 2, \ldots, n\}$ denote the set of players, and let Ω be a non-empty, finite state space with a strictly positive common prior μ. Each player $i \in N$ has a finite partition $\mathcal{P}_i$ of Ω, representing player i's information. For any state $\omega \in \Omega$, we denote by $\mathcal{P}_i(\omega)$ the element of the partition $\mathcal{P}_i$ that contains ω. A Common Knowledge Component (CKC), typically denoted by $C \subseteq \Omega$, is a minimal non-empty subset of states that is measurable with respect to the σ-algebra of every player (see Aumann, 1976). The notation $C(\omega)$ refers to the CKC that contains the state ω.

Let F be a partition of Ω belonging to an agent outside the set N, which we refer to as the *mediator*. As before, $F(\omega)$ denotes the information set of F that contains ω. The mediator may provide additional information to the players beyond their private information. For this purpose, the mediator uses a public signaling function $\tau : \Omega \to \Delta(S)$, where S is a finite set of signals and $\Delta(S)$ is the set (simplex) of distributions over S. Let $\tau(s \mid \omega)$ denote the probability that the public signal $s \in S$ is observed given state ω. Note that the signaling function is measurable with respect to (henceforth, w.r.t.) F.[7] Formally, τ is an F-measurable stochastic kernel:

(i) For each state $\omega \in \Omega$, $\tau(\cdot \mid \omega)$ is a probability distribution over a finite set of signals S.

(ii) For all $\omega' \in F(\omega)$ and $s \in S$, we have $\tau(s \mid \omega') = \tau(s \mid \omega)$.

The signaling function τ is also known as a Blackwell experiment (see Blackwell, 1951), so the mediator can be viewed as a generator of Blackwell experiments.

[7]Measurability w.r.t. F means that if ω and ω' are indistinguishable from the mediator's perspective, namely, in case $\omega' \in F(\omega)$, then $\tau(s \mid \omega') = \tau(s \mid \omega)$.

2.1 Bayesian updating and joint posterior beliefs

When the realized state is ω, player i is informed of $\mathcal{P}_i(\omega)$. Assuming there exists $\omega' \in \mathcal{P}_i(\omega)$ such that $\tau(s|\omega') > 0$, upon observing s, player i updates his belief using Bayes' rule. The updated belief is given by the following distribution.

$$\mu_{\tau,i}(\omega' \mid \mathcal{P}_i(\omega), s) = \frac{\mu(\omega')\tau(s \mid \omega')}{\sum_{\omega'' \in \mathcal{P}_i(\omega)} \mu(\omega'')\tau(s \mid \omega'')}. \tag{1}$$

In particular, $\mu_{\tau,i}(\cdot \mid \mathcal{P}_i(\omega), s)$ is a probability distribution whose support is a subset of $\mathcal{P}_i(\omega)$. For notational convenience, we set $\mu_{\tau,i}(\omega' \mid \mathcal{P}_i(\omega), s) = 0$ for every ω', whenever $\tau(s \mid \mathcal{P}_i(\omega)) = 0$.

When the realized state may vary, we define the *joint posterior associated with* τ *and* s as the set of posterior profiles across all states and players that can be generated by s:

$$\boldsymbol{\mu}_{\tau,s} = \Big\{ \big(\mu_{\tau,i}(\cdot \mid \mathcal{P}_i(\omega), s)\big)_{i \in N} \ : \ \omega \in \Omega \text{ with } \tau(s \mid \omega) > 0 \Big\}. \tag{2}$$

This joint posterior records not only each player's beliefs about the true state but also their beliefs about others' beliefs and higher-order beliefs, making it a central object for analyzing equilibrium behavior.

2.2 Joint beliefs

To allow for even greater generality, the following definition of a *joint belief* requires neither a signal nor a mediator. It will be used throughout the paper to define the joint profile of beliefs that a mediator can potentially generate.

Definition 1. *A* joint belief *is a map* $\mathrm{JB} : \Omega \times N \to \Delta(\Omega) \cup \{\mathbf{1}_\emptyset\}$, *where* $\mathbf{1}_\emptyset$ *is the zero vector, satisfying:*

(*i*) Common realized state*: if* $\mathrm{JB}(\omega, i)(\omega) > 0$ *for some* i*, then* $\mathrm{JB}(\omega, j)(\omega) > 0$ *for all* $j \in N$*.*

(*ii*) Information-set support*: if* $\mathrm{JB}(\omega, i) \neq \mathbf{1}_\emptyset$*, then the support of* $\mathrm{JB}(\omega, i)$ *is a subset of* $\mathcal{P}_i(\omega)$*, and* $\mathrm{JB}(\omega, i)(\omega) > 0$*.*

(*iii*) Information-set invariance*: for every* $\omega' \in \mathcal{P}_i(\omega)$*, if* $\mathrm{JB}(\omega, i) \neq \mathbf{1}_\emptyset$ *and* $\mathrm{JB}(\omega', i) \neq \mathbf{1}_\emptyset$*, then* $\mathrm{JB}(\omega', i) = \mathrm{JB}(\omega, i)$*.*

(*iv*) Likelihood-ratio agreement*: for every i, j, ω and $\omega' \in \mathcal{P}_i(\omega) \cap \mathcal{P}_j(\omega)$,*

$$\frac{\mathrm{JB}(\omega, i)(\omega)}{\mathrm{JB}(\omega', i)(\omega')} = \frac{\mathrm{JB}(\omega, j)(\omega)}{\mathrm{JB}(\omega', j)(\omega')},$$

assuming the denominators are nonzero.

When well-defined (i.e., when $\mathrm{JB}(\omega, i) \neq \mathbf{1}_\emptyset$), $\mathrm{JB}(\omega, i)$ stands for the belief distribution that player i assigns to Ω when the set $\mathcal{P}_i(\omega)$ is realized. The corresponding probability of state ω' is denoted by $\mathrm{JB}(\omega, i)(\omega')$.

Several remarks are in order. First, the conditions in Definition 1 are necessary for a joint belief to be induced as a joint posterior belief by a public signal. Indeed, they follow from the basic properties of Bayesian inference given the existence of a common prior and a mediator's public signal.[8] For example, a common realized state (condition (i) above) must be satisfied whenever players consider the same probability space, and in cases without a common prior, likelihood-ratio agreement (condition (iv) above) need not be satisfied. Second, we include the indicator $\mathbf{1}_\emptyset$ to capture cases where the mediator's signal has zero likelihood at ω; in such cases, we set $\mathrm{JB}(\omega, i) = \mathbf{1}_\emptyset$ for all such states. Note that $\tau(s \mid \omega') = 0$ for some $\omega' \in \mathcal{P}_i(\omega)$ does not by itself make the posterior in Eq. (1) ill-defined; it only implies that $\mu_{\tau,i}(\omega' \mid \mathcal{P}_i(\omega), s) = 0$. Third, the definition hinges on the players' fixed partitions, primarily through the information-set support and invariance requirements (conditions (ii) and (iii) above), but it does not relate to the mediator's partition. Moreover, these two conditions apply to each player *at the individual level*, rather than across players.

Given a JB and a common prior μ, a natural question is whether there exists a signaling function τ and a signal s that jointly induce this JB as a posterior. We approach this question by examining posterior likelihood ratios. This builds on a key observation that, although posterior updating is typically player-dependent (as the denominator reflects player i's information partition), the following likelihood ratios are not:

$$\frac{\mu_{\tau,i}(\omega \mid \mathcal{P}_i(\omega), s)}{\mu_{\tau,i}(\omega' \mid \mathcal{P}_i(\omega'), s)} = \frac{\mu(\omega)\tau(s \mid \omega)}{\mu(\omega')\tau(s \mid \omega')}, \quad \forall \omega, \omega' \in \mathcal{P}_i(\omega). \tag{3}$$

In particular, whenever ω and ω' lie in the same information set (i.e., $\omega' \in \mathcal{P}_i(\omega)$, or equivalently $\mathcal{P}_i(\omega) = \mathcal{P}_i(\omega')$) and both conditional probabilities are positive, this ratio is well defined and independent of the player. This also clarifies the two constraints imposed in Definition 1. We explore this idea further in the following section.

To formally define the problem, fix a JB and a common prior μ, and let Ω_+ denote the set of states such that $\mathrm{JB}(\omega, i)(\omega) > 0$ for some player i, and hence for all players. In particular,

[8] A broader discussion of this point is provided in connection with Eq. (3) below.

we restrict attention to the states that are assigned positive probability. Normalizing μ to this set, we obtain the distribution $\mu(\cdot \mid \Omega_+)$.[9]

For a pair $\omega, \omega' \in \Omega_+$ such that $\omega' \in \mathcal{P}_i(\omega)$ for some i, consider the ratio $\frac{\mathrm{JB}(\omega,i)(\omega)}{\mathrm{JB}(\omega',i)(\omega')}$. We ask whether there exists a signaling function τ of the mediator and a signal s that induce the same likelihood ratios. Specifically, does there exist a signal s and a function τ such that

$$\frac{\mathrm{JB}(\omega,i)(\omega)}{\mathrm{JB}(\omega',i)(\omega')} = \frac{\mu(\omega \mid \Omega_+)\tau(s \mid \omega)}{\mu(\omega' \mid \Omega_+)\tau(s \mid \omega')} = \frac{\mu(\omega)\tau(s \mid \omega)}{\mu(\omega')\tau(s \mid \omega')}, \tag{4}$$

for every player i and every $\omega' \in \mathcal{P}_i(\omega)$? Note that the right-hand side of Eq. (4) does not depend on player i, but only on whether ω and ω' lie in the same information set for some player. To further discuss this question and see why likelihood ratios are sufficient statistics for generating joint posteriors, consider the following example.

Example 1. There are five states and two players such that $\mathcal{P}_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4, \omega_5\}\}$, $\mathcal{P}_2 = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}, \{\omega_5\}\}$, and $F = \{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3, \omega_5\}\}$. Figure 3 illustrates the knowledge structures of the players as well as that of the mediator.

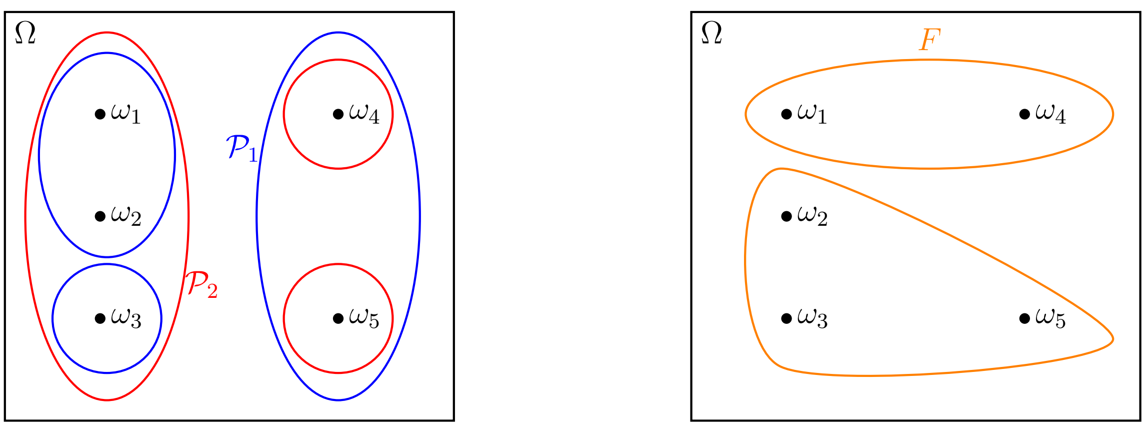

(a)
The players' information

(b)
The mediator's information

Figure 3: On the left, Figure (a) illustrates the information structure of player 1 (blue) and player 2 (red). On the right, Figure (b) portrays the information structure of the mediator (orange).

Starting with a basic set-up and in the absence of any mediator, one may ask whether there exists a common prior that induces the joint posterior given in Table 1.

[9]We discuss the case where some states are assigned zero probability in Section 3.2.

	Player 1	Player 2
ω_1	(1/2,1/2,0,0,0)	(1/3,1/3,1/3,0,0)
ω_2	(1/2,1/2,0,0,0)	(1/3,1/3,1/3,0,0)
ω_3	(0,0,1,0,0)	(1/3,1/3,1/3,0,0)
ω_4	(0,0,0,1/2,1/2)	(0,0,0,1,0)
ω_5	(0,0,0,1/2,1/2)	(0,0,0,0,1)

Table 1: A JB; player-specific posteriors conditional on the realized states.

Indeed, such a prior exists: the uniform distribution over Ω induces exactly these posterior beliefs. Most importantly, in this example there are two states, ω_1 and ω_2, that lie in the same information set for both players (in contrast to the four-state example presented in the introduction).

As a consequence, the ratio between the probabilities assigned to these states (namely, $\frac{1/2}{1/2}$ for player 1 and $\frac{1/3}{1/3}$ for player 2) is identical across players. This property persists even when a mediator provides additional information, as illustrated in the next table, and it constitutes a cornerstone of our characterization.

Suppose now that the common prior is uniform. Does there exist a public signal that the mediator could reveal so as to induce the JB given in Table 2?

	Player 1	Player 2
ω_1	(1/3,2/3,0,0,0)	(1/5,2/5,2/5,0,0)
ω_2	(1/3,2/3,0,0,0)	(1/5,2/5,2/5,0,0)
ω_3	(0,0,1,0,0)	(1/5,2/5,2/5,0,0)
ω_4	(0,0,0,1/3,2/3)	(0,0,0,1,0)
ω_5	(0,0,0,1/3,2/3)	(0,0,0,0,1)

Table 2: The JB conditional on the mediator's signal, under a uniform prior.

The answer is yes. Consider a signaling function τ and a signal s such that $\tau(s \mid \omega_i) = 1/5$ for $i \in \{1,4\}$ and $\tau(s \mid \omega_i) = 2/5$ for $i \in \{2,3,5\}$. Since $\tau(s \mid \omega_1) = \tau(s \mid \omega_4)$ and $\tau(s \mid \omega_2) = \tau(s \mid \omega_3) = \tau(s \mid \omega_5)$, the function τ is F-measurable.

Now modify Table 2 by replacing player 1's posteriors at ω_4 and ω_5 with $(0,0,0,1/5,4/5)$. The resulting table still defines a joint belief (Definition 1), but it is *not* implementable by any F-measurable public signal.

The reason is that if, under ω_1, the posterior of player 1 is $\left(\frac{1}{3},\frac{2}{3},0,0,0\right)$, then due to the F-measurability restriction, the ratio

$$\frac{\tau(s \mid \omega_1)}{\tau(s \mid \omega_2)} \quad \text{must match} \quad \frac{\tau(s \mid \omega_4)}{\tau(s \mid \omega_5)}.$$

However, in the former case, the ratio is $\frac{1}{2}$, while in the latter it is $\frac{1}{4}$.

A hierarchy of beliefs.

A joint belief specifies each player's posterior at every state, and therefore encodes not only first-order beliefs but also higher-order beliefs, which are central for equilibrium analysis. For example, in Table 1, when ω_4 is realized player 1 assigns probability $1/2$ to each of ω_4 and ω_5. Hence player 1 assigns probability $1/2$ to the event that player 2 is certain the state is ω_4 and probability $1/2$ to the event that player 2 is certain the state is ω_5. This illustrates how the state-by-state specification of posteriors captures a *hierarchy of beliefs*.

This example also highlights why likelihood ratios are the right object for our analysis: any F-measurable public signal must preserve relative likelihoods of states within each information set, and these ratios will become the algebraic constraints in our implementability characterization.

2.3 The graph of information and PL functions

In this section, we present two key elements for our analysis: the graph of information and posterior likelihood functions. Starting with the former, we represent the agents' private information by a directed graph, following Rodrigues-Neto (2009). The graph provides a compact way to record which pairs of states are *directly comparable* for at least one player (i.e., states that a player may confound within an information set). This will later allow us to express Bayesian feasibility restrictions as simple multiplicative constraints along particular chains in the graph.

Formally, let $G = (V, E)$ be a directed graph with vertex set $V = \Omega$. For two states $\omega, \omega' \in \Omega$, we include a directed edge $(\omega, \omega') \in E$ (denoted $\omega \to \omega'$) if there exists a player $i \in N$ such that $\omega' \in \mathcal{P}_i(\omega)$. Since each $\mathcal{P}_i$ is a partition, whenever $\omega \to \omega'$ we also have $\omega' \to \omega$. We keep both directions because likelihood ratios will be defined on directed edges. See below the illustrative graph for Example 1.

Let $\twoheadrightarrow$ denote the transitive closure of $\to$: we write $\omega \twoheadrightarrow \omega'$ if there exists a finite sequence of states $\omega = \omega_1, \omega_2, \ldots, \omega_m = \omega'$ such that $\omega_k \to \omega_{k+1}$ for all $k < m$. Because edges come in both directions, $\twoheadrightarrow$ is an equivalence relation; its equivalence classes coincide with the common knowledge components (CKCs). In particular, G decomposes into connected components, each corresponding to a CKC.

The second central object in our characterization is a function that records *local* likelihood ratios (comparisons across edges).

Definition 2 (Posterior likelihood (PL) function)**.** *A positive function φ defined over E is* *a* posterior likelihood *function if it satisfies the condition*[10]

$$\varphi(\omega, \omega') = \frac{1}{\varphi(\omega', \omega)}. \tag{5}$$

Intuitively, $\varphi(\omega, \omega')$ is the likelihood ratio of ω relative to ω' as perceived by any player who cannot distinguish between them. Later, a joint belief (together with a prior) will induce such an φ on E via Bayes' rule. We typically ignore self edges (as $\omega \in \mathcal{P}_i(\omega)$ trivially), but if needed, set $\varphi(\omega, \omega) = 1$.

Example 1, continued. The graph of information corresponding to the model described in Example 1, omitting arrows from a state to itself, is given in Figure 4. Note that the likelihood ratio of the probabilities of ω_1 and ω_2 is $1/2$ (recall Table 2 above). This is true for both players since ω_1 and ω_2 belong to the same information sets of the two. We therefore obtain that $\varphi(\omega_1, \omega_2) = 1/2$, and the other values of the PL function are:

$$\varphi(\omega_2, \omega_3) = 1, \quad \varphi(\omega_3, \omega_1) = 2, \quad \text{and} \quad \varphi(\omega_4, \omega_5) = 1/2, \tag{6}$$

while preserving the relation in Eq. (5).

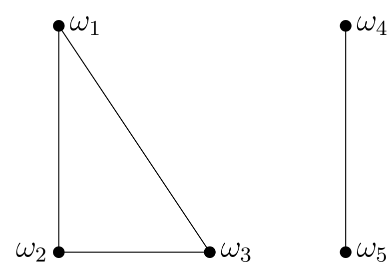

Figure 4: The graph of information. In this example, there are two connected components, each corresponding to a CKC.

2.4 Cycles and information loops

The mediator's partition F may be coarser than the players' common-knowledge partition: the mediator may pool together states that all players (collectively) do distinguish. This mismatch is captured by two graph-theoretic objects: F-cycles and F-loops. We introduce them here so that the subsequent consistency section can focus purely on restrictions on φ rather than on the combinatorics of the environment.

An F-cycle is a chain of player-edges that "closes" because the terminal state lies in the same mediator cell as the initial state.

[10]We abuse notation and use $\varphi(\omega, \omega')$ instead of $\varphi((\omega, \omega'))$.

Definition 3 (F-cycle)**.** *An* F-cycle *is a finite sequence of states* $\omega_1, \omega_2, \ldots, \omega_{m+1}$ *such that*

$$(\omega_i, \omega_{i+1}) \in E \quad \textit{for } i = 1, \ldots, m, \qquad \textit{and} \qquad \omega_{m+1} \in F(\omega_1).$$

Since $\omega_1 \twoheadrightarrow \omega_{m+1}$, every F-cycle is contained in a single CKC; the mediator enters only through the closing condition $\omega_{m+1} \in F(\omega_1)$. When F is discrete (singleton cells), an F-cycle reduces to an ordinary directed cycle $\omega_{m+1} = \omega_1$.

Next, we define information loops that capture the mediator's ability to connect distinct CKCs through its own coarser information. The concept of a loop plays a pivotal role in Lagziel et al. (2025) to provide conditions such that one mediator dominates another, in an extension of Blackwell's work on the comparison of experiments (see Blackwell, 1951).

Definition 4 (F-loop)**.** *An* F-loop *is a finite sequence of pairs* $\big((\omega_i, \bar{\omega}_i)\big)_{i=1}^{m}$ *such that, for each* $i = 1, \ldots, m$*:*

(i) $\omega_i \twoheadrightarrow \bar{\omega}_i$ *(the pair lies in the same CKC);*

(ii) $\bar{\omega}_i \not\twoheadrightarrow \omega_{i+1}$ *(successive pairs lie in different CKCs);*

(iii) $\bar{\omega}_i \in F(\omega_{i+1})$*, with the convention* $\omega_{m+1} := \omega_1$*.*

Condition (i) means that each pair $(\omega_i, \bar{\omega}_i)$ lies within a single CKC. Condition (ii) forces the loop to move across CKCs. Condition (iii) explains *how* it moves: the mediator cannot distinguish $\bar{\omega}_i$ from the next state ω_{i+1} because they lie in the same F-cell. For instance, the example given in the introduction depicts a situation where the mediator cannot distinguish between states ω_1 and ω_3, both in different CKCs, and between states ω_2 and ω_4, although the two players commonly distinguish between these states. This forms an F-loop of $((\omega_1, \omega_2), (\omega_4, \omega_3))$.

The graph, cycles, and loops introduced above are purely structural objects determined by the players' information partitions and the mediator's partition F. In Section 3.1, we translate these objects into algebraic restrictions on the posterior likelihood function φ by evaluating products of likelihood ratios along them. Along an F-cycle we multiply φ directly on successive edges, whereas along an F-loop we must compare states within a CKC that may not be adjacent; we do so by extending φ along paths inside each CKC (internal consistency will guarantee that this extension is well defined). These consistency conditions will then be used in our characterization of Bayesian implementability.

3 Consistency conditions and main result

In this section we introduce two restrictions, referred to as *internal* and *external* consistency, that characterize when a posterior-likelihood function φ can be generated by an F-measurable signaling function. Section 3.1 develops the graph-theoretic characterization, and Section 3.2 applies it to Bayesian updating to obtain our main implementability theorem.

3.1 Internal and external consistency

We start with a purely structural question that depends only on the players' partitions $(\mathcal{P}_i)_i$ and the mediator's partition F. Fix a PL function φ on the edge set E. Does there exist an F-measurable, strictly positive function $f : \Omega \to \mathbb{R}_{++}$ such that

$$\varphi(\omega, \omega') = \frac{f(\omega)}{f(\omega')} \qquad \text{for every } (\omega, \omega') \in E?$$

To address this question, recall the notions of an F-cycle and an F-loop from Section 2.4. We use these objects to formulate two consistency conditions, which Proposition 1 shows are necessary and sufficient. This result will later allow us to test whether a joint belief JB is implementable: we map JB into its induced PL function and verify the consistency conditions (see Eq. (4) and Remark 1).

The first property, *internal consistency*, parallels Definition 2 in Rodrigues-Neto (2009), where cycles are defined solely in terms of the players' information, independently of the mediator.[11]

Definition 5. *Let* $(\omega_i, \omega_{i+1})_{i=1}^{m}$ *be an* F*-cycle. Then,* Internal consistency [**INC**] *holds if*

$$\prod_{i=1}^{m} \varphi(\omega_i, \omega_{i+1}) = 1. \tag{7}$$

Hellman and Samet (2012) refer to the left-hand side of Eq. (7) (when F is trivial) as the *type ratio of a chain*. Note that for every $(\omega_1, \omega_2) \in E$, the sequence $(\omega_1, \omega_2, \omega_1)$ is an F-cycle. Thus, by Eq. (5), a positive function φ defined on E is a PL function if and only if Eq. (7) holds for every such F-cycle. In particular, [**INC**] implies Eq. (5), and hence φ is indeed a PL function.

The following Lemma 1 uses [**INC**] to extend any φ in a consistent manner in every CKC. (The proof is relegated to the appendix; see Section A.1.)

[11] In specific cases, Hellwig (2013) simplifies the consistency test of Rodrigues-Neto (2009) for verifying whether a given set of players' posteriors is compatible with a common prior.

Lemma 1. *Assume* [**INC**]. *Then, φ can be extended to any pair (ω_1, ω_{m+1}) where $\omega_1 \twoheadrightarrow \omega_{m+1}$, and it holds that for any sequence of edges $\big((\omega_i, \omega_{i+1})\big)_{i=1}^m$ in E,*

$$\prod_{i=1}^{m} \varphi(\omega_i, \omega_{i+1}) = \varphi(\omega_1, \omega_{m+1}). \tag{8}$$

In Example 1, the sequence $\big((\omega_1, \omega_2), (\omega_2, \omega_3), (\omega_3, \omega_1)\big)$ forms an F-cycle, and indeed, in Table 2, the product of the corresponding values that φ assigns to these edges does equal 1 (see Eq. (6) above).

Remark 1. Lemma 1 *implies that if φ satisfies* [**INC**]*, then for each connected component C of G the ratios encoded by φ determine a probability distribution on C. Specifically,*

$$\mu_\varphi(\omega' \mid C) = \frac{\varphi(\omega', \omega)}{\sum_{\omega'' \in C} \varphi(\omega'', \omega)}, \tag{9}$$

where ω is any reference state in C (the expression is independent of this choice).

Similar to the notion of an F-cycle and the internal consistency property, we employ the F-loop to define external consistency as follows.

Definition 6. *Assume* [**INC**] *holds so that φ extends as in* Lemma 1*, and let $\big((\omega_i, \bar{\omega}_i)\big)_{i=1}^m$ be an F-loop. Then,* External Consistency [**EXC**] *holds if*

$$\prod_{i=1}^{m} \varphi(\omega_i, \bar{\omega}_i) = 1. \tag{10}$$

The [**EXC**] condition has substantive content only when the mediator lacks knowledge of the players' common knowledge. Equivalently, if for every ω we have $F(\omega) \subseteq C(\omega)$, then no F-loops arise and [**EXC**] becomes vacuous. An extreme example is an all-knowing mediator, who knows the exact identity of the realized state, i.e., $F(\omega) = \{\omega\}$ for every $\omega \in \Omega$.

We revisit the introductory example to illustrate the role of [**EXC**] in our setting. Suppose the mediator aims to induce a JB in which player 1's belief is $(\frac{1}{3}, \frac{2}{3}, 0, 0)$ and player 2's belief is $(0, 0, \frac{3}{4}, \frac{1}{4})$. The corresponding posterior likelihood ratios are $\varphi(\omega_1, \omega_2) = \frac{1}{2}$ and $\varphi(\omega_4, \omega_3) = \frac{1}{3}$. The product of ratios along the loop is therefore $\varphi(\omega_1, \omega_2)\varphi(\omega_4, \omega_3) = \frac{1}{6} \neq 1$, and as previously stated, this JB is indeed infeasible.

The technical backbone of our analysis, stated in Proposition 1, establishes that internal and external consistency, w.r.t. a given PL function φ, are necessary and sufficient conditions for the existence of an F-measurable function that reproduces the likelihood ratios encoded by φ. (The proof is relegated to the appendix; see Section A.2.)

Proposition 1. *Fix φ. There exists an F-measurable and strictly positive function f defined on V such that, for every $(\omega, \omega') \in E$,*

$$\varphi(\omega, \omega') = \frac{f(\omega)}{f(\omega')}, \tag{11}$$

if and only if [**INC**] *and* [**EXC**] *hold for the extension provided by* Lemma 1.

Once Proposition 1 establishes a characterization of a given PL function, we can employ it in Section 3.2 to show how one can take a JB, translate it into a PL function, and then replicate this JB for a given mediator.

As stated in Remark 1, when φ satisfies [**INC**] it induces a distribution $\mu_\varphi(\cdot \mid C)$ on each connected component C. Section 3.2 shows that if φ also satisfies [**EXC**], then the mediator can choose an F-measurable signaling function τ and a signal s such that the posterior distribution over Ω given s admits the decomposition

$$\mu_\tau(\cdot \mid s) = \sum_C \mu(C \mid s)\, \mu_\varphi(\cdot \mid C), \tag{12}$$

where the sum ranges over all connected components of G, and

$$\mu(C \mid s) = \frac{\sum_{\omega \in C} \tau(s \mid \omega)\mu(\omega)}{\sum_{\omega \in \Omega} \tau(s \mid \omega)\mu(\omega)}.$$

In simple terms, $\mu_\tau(\cdot \mid s)$ is the posterior distribution over Ω after observing the message s, sent by the mediator.

Remark 2. *The preceding discussion has focused on the case of public signals, where all players observe the same realization from the information structure. Evidently, if a signaling function induces the joint posteriors of the entire group N, then it also induces the joint posteriors of any subgroup $I \subseteq N$. Consequently,* Proposition 1 *implies that if the corresponding* PL *function φ_{JB} satisfies the consistency conditions for N players, it will also do so for any subset of N.*

3.2 Main characterization: implementation to Bayesian updating

In this section, we apply Proposition 1 to the context of Bayesian updating. Suppose all players share a common prior μ, and let JB denote a joint belief, as defined in Definition 1. Can this joint belief be induced by a signaling function τ of the mediator and a signal s?

We employ Proposition 1 to answer this question. To see this, let us focus first on the subset $\Omega_+ \subseteq \Omega$ of states ω for which[12] $\mathrm{JB}(\omega, \mathrm{i})(\omega) > 0$ for some player i. Consider the graph $G_+ = (\Omega_+, E_+)$, which is the restriction of the original graph $G = (\Omega, E)$ to the subset Ω_+. Using Eq. (4), define a new PL function φ_{JB} on the edge set E_+ as follows:

$$\varphi_{\mathrm{JB}}(\omega, \omega') = \frac{\mathrm{JB}(\omega, i)(\omega)}{\mathrm{JB}(\omega', i)(\omega')} \cdot \frac{\mu(\omega')}{\mu(\omega)}, \quad \forall\ (\omega, \omega') \in E_+\ , \omega' \in \mathcal{P}_i(\omega). \tag{13}$$

Since both μ and $\mathrm{JB}(\omega, \mathrm{i})(\omega)$ are strictly positive on Ω_+, φ_{JB} is strictly positive as well. Note that the common realized state constraint (condition (i) in Definition 1), ensures that $\mathrm{JB}(\omega, j)(\omega) > 0$ for every player j and every $\omega \in \Omega_+$. Accordingly, the information-set invariance requirement (condition (iii) in Definition 1) ensures that $\mathrm{JB}(\omega, i) = \mathrm{JB}(\omega', i)$, as long as both are not the zero vector. Moreover, the likelihood-ratio agreement (condition (iv) in Definition 1) ensures that the right-hand side is invariant to the chosen player i, so φ_{JB} is a well-defined PL function. Given that **[INC]** holds, we can extend φ_{JB} to every pair (ω, ω') in a given CKC as done in Lemma 1. So, whenever **[INC]** holds, we henceforth consider the extended PL function.

Recall that $\boldsymbol{\mu}_{\tau,s}$ (see Eq. (2) above) denotes the joint posterior associated with τ and s. As an implication of Proposition 1, the following theorem states that internal and external consistency (w.r.t. φ_{JB}) are necessary and sufficient conditions for the existence of an F-measurable signaling function τ and a signal s, such that $\boldsymbol{\mu}_{\tau,s} = \mathrm{JB}$.[13] Other than the two consistency conditions, we also require that Ω_+ is F-measurable, because any JB that assigns zero mass to some states and positive mass to other states within the same mediator information set is impossible to implement. (The proof is relegated to the appendix; see Section A.3.)

Theorem 1 (Main Characterization: Implementable Joint Posteriors)**.** *Fix all agents' partitions, a prior, and a joint belief* JB*. Then there exists an F-measurable signaling function τ and a signal s, such that $\boldsymbol{\mu}_{\tau,s} = \mathrm{JB}$ on Ω_+ if and only if*

- *(i) Ω_+ is measurable w.r.t. F; and*
- *(ii) φ_{JB} satisfies conditions* **[INC]** *and* **[EXC]** *in G_+.*[14]

Note that condition (i) is necessary because the prior μ is strictly positive. If the mediator sends a given signal at some $\omega \in \Omega_+$, then the same signal must also be sent at every

[12] Recall that $\mathrm{JB}(\omega, \mathrm{i})$ is a distribution while $\mathrm{JB}(\omega, \mathrm{i})(\omega)$ is a probability.

[13] To be clear, the equivalence holds for every player and every state in Ω_+, specifically, $\forall \omega \in \Omega_+$ and $\forall i \in N$, we have $\mu_{\tau,i}(\cdot|\mathcal{P}_i(\omega), s) = \mathrm{JB}(\omega, i)$.

[14] **[EXC]** is satisfied using the extension provided by Lemma 1, ensured by **[INC]**.

$\omega' \in F(\omega)$, with the same probability. Consequently, if $\omega' \notin \Omega_+$, the posterior probability assigned to ω' cannot be zero unless the prior probability of ω' was already zero.

The economic implication of Theorem 1 is as follows. Suppose a function φ is given. By Remark 1, φ induces a probability distribution over each connected component, and thus determines the posterior beliefs of all players. Theorem 1 identifies conditions under which there exists a signaling mechanism, implemented by the mediator, that induces these posterior beliefs, as described in Eq. (1).

4 Non-convexity of the feasible set

The key idea of *concavification* originates from the work of Aumann and Maschler (1995) on repeated games with incomplete information and, more recently, from the Bayesian persuasion literature. It builds on the "splitting lemma" which traces back to Blackwell (1951), relying on the ability to "split" the prior through signaling, under the crucial property that any convex combination of feasible posteriors is itself feasible. In other words, the set of feasible posteriors is convex. As it turns out, this property fails in more general informational environments, such as those analyzed here.

To understand why convexity may break down, consider the consistency constraints in Theorem 1. The external consistency condition builds on the function φ_{JB} defined in Equation (13). Given either an F-loop or an F-cycle, take two distinct feasible joint posteriors that individually satisfy the consistency constraints. For convexity to hold, their average should also satisfy these constraints. However, the constraints are *nonlinear* in the posterior probabilities, since they arise from a product of likelihood ratios. So the average may fail to sustain a necessary feasibility condition.

Recall that Definition 1 allows $\mathrm{JB}(\omega, i)(\omega)$ to take the value $\mathbf{1}_\emptyset$, which does not represent a probability distribution. Consequently, the set of joint beliefs is not convex in the standard sense: if one joint belief is a valid distribution and another is $\mathbf{1}_\emptyset$, their convex combination is not a well-defined belief. We therefore restrict our attention to cases where, conditional on the state and player, the support of the two joint beliefs are in the same partition elements. Formally, we adopt the following definition of convexity:

Definition 7. *The feasible set is* convex *if for any two feasible joint beliefs* JB_1 *and* JB_2 *with matching information-supports (i.e.,* $\mathrm{JB}_1(\omega, i) = \mathbf{1}_\emptyset \iff \mathrm{JB}_2(\omega, i) = \mathbf{1}_\emptyset$ *for all* (ω, i)*), the joint belief* $\lambda \mathrm{JB}_1 + (1-\lambda)\mathrm{JB}_2$ *is feasible for all* $\lambda \in [0, 1]$.

Non-convexity arises not only in multiplayer environments. Even in a single-player setting, when the mediator and the player possess different information, the induced set of posteriors may be non-convex.[15] To illustrate how convexity fails, consider the following example with four states and a non-uniform prior.

Example 2. Fix $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, endowed with a prior of $\mu(\omega_i) = 0.2$ for every $i = 1, 2, 3$ and $\mu(\omega_4) = 0.4$. Player 1's partition is $\mathcal{P}_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$, and the mediator's partition is $F = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. So there exists an F-loop of $((\omega_1, \omega_2)(\omega_4, \omega_3))$.

Fix the F-measurable signaling function τ where $\tau(s_1|\{\omega_1, \omega_3\}) = \frac{1}{3} = 1 - \tau(s_2|\{\omega_1, \omega_3\})$ and $\tau(s_1|\{\omega_2, \omega_4\}) = \frac{2}{3} = 1 - \tau(s_2|\{\omega_2, \omega_4\})$. The induced posteriors are given in Table 3(a). Now consider the average of the two posteriors given in Table 3(b), and let us see whether external consistency holds:

$$\frac{\mathrm{JB}(\omega_1,1)(\omega_1)}{\mathrm{JB}(\omega_2,1)(\omega_2)} \cdot \frac{\mu(\omega_2)}{\mu(\omega_1)} \cdot \frac{\mathrm{JB}(\omega_4,1)(\omega_4)}{\mathrm{JB}(\omega_3,1)(\omega_3)} \cdot \frac{\mu(\omega_3)}{\mu(\omega_4)} = \frac{1/2}{1/2} \cdot \frac{1/5}{1/5} \cdot \frac{13/20}{7/20} \cdot \frac{1/5}{2/5} = \frac{13}{14} \neq 1.$$

Hence, according to Theorem 1, the average posterior is not feasible and therefore the feasible set related to μ is not convex.

	post$(\cdot \mid s_1)$	post$(\cdot \mid s_2)$
ω_1	(1/3,2/3,0,0)	(2/3,1/3,0,0)
ω_2	(1/3,2/3,0,0)	(2/3,1/3,0,0)
ω_3	(0,0,1/5,4/5)	(0,0,1/2,1/2)
ω_4	(0,0,1/5,4/5)	(0,0,1/2,1/2)

(a)

	Av. post $= \frac{1}{2}(\mathrm{post}(\cdot \mid \mathrm{s}_1) + \mathrm{post}(\cdot \mid \mathrm{s}_2))$
ω_1	(1/2,1/2,0,0)
ω_2	(1/2,1/2,0,0)
ω_3	(0,0,7/20,13/20)
ω_4	(0,0,7/20,13/20)

(b)

Table 3: Figure (a) depicts the player's posterior belief given the mediator's signaling function τ, and Figure (b) presents the average over the two posterior beliefs.

Next, one can devise an example with a fully informed mediator and a uniform prior that still does not admit convexity. Fix $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, endowed with a uniform prior, a fully informed mediator, while $\mathcal{P}_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ and $\mathcal{P}_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. So there exists a cycle that goes through all four states. Now consider two signals, the first s_1 is completely non-informative and the second is $\tau(s_2|\omega_1) = \tau(s_2|\omega_2) = \frac{1}{6}$, $\tau(s_2|\omega_3) = \frac{2}{6}$ and $\tau(s_2|\omega_4) = \frac{3}{6}$. One can easily compute the joint posteriors given each signal and their average, and use our characterization to verify that the outcome is a non-implementable joint posterior, although the mediator has full information.

We can also simplify the state space with an example that builds on three equally likely states, where $\mathcal{P}_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}\}$, $\mathcal{P}_2 = \{\{\omega_1, \omega_3\}, \{\omega_2\}\}$, and $F = \{\{\omega_1\}, \{\omega_2, \omega_3\}\}$.

[15] Kolotilin et al. (2017) and Guo and Shmaya (2019) provide insights into persuasion problems with privately informed receivers.

Take two signals where the first is completely non-informative and the second is $\tau(s_2|\omega_2) = \tau(s_2|\omega_3) = 2\tau(s_2|\omega_1)$. Again, a simple computation shows that the average of the two joint posterior beliefs is not implementable.

On the other hand, it is easy to show that the feasible set is convex in two benchmark cases of: (i) uninformed players; and (ii) a fully informed mediator with one player.[16] The question now arises: when is the feasible set convex? Our graph-of-information characterization provides tools to address this question, but it requires two additional notions: *pseudo-independence* and *informative loops*.

Definition 8. *Let $\mathcal{A} = \{A_1, \ldots, A_\ell\}$ and $\mathcal{B} = \{B_1, \ldots, B_k\}$ be two partitions of Ω. We say that $\mathcal{A}$ and $\mathcal{B}$ are* pseudo-independent *if, for every strictly positive prior μ, there exist constants $a_1, \ldots, a_\ell$ and $b_1, \ldots, b_k$ such that, for every i and j with $A_i \cap B_j \neq \emptyset$, $\mu(A_i \cap B_j) = a_i b_j$.*

Technically, pseudo-independence means that, for every prior μ, the incidence matrix $\big(\mu(A_i \cap B_j)\big)_{ij}$, which represents the joint distribution induced by μ over the two partitions, equals the matrix which is the outer product of two vectors, $(a_1, \ldots, a_\ell)$ and $(b_1, \ldots, b_k)$, except for the cells where $A_i \cap B_j = \emptyset$.

This definition establishes a structural condition on the geometry of the partitions. Because the factorization is required to hold for *every* prior, pseudo-independence describes a combinatorial property of the information structures themselves, independent of any specific belief system or probability measure.

Example 3. Consider Ω, μ, and $\mathcal{P}_1$ as in Example 2. For convenience, denote $A_1 = \{\omega_1, \omega_2\}$ and $A_2 = \{\omega_3, \omega_4\}$, while the mediator's partition is now $F = \{B_1, B_2, B_3\}$ where $B_1 = \{\omega_1, \omega_3\}, B_2 = \{\omega_2\}$, and $B_3 = \{\omega_4\}$. It is clear that $\mathcal{P}_1$ and F are not probabilistically independent with respect to μ.[17] However, for the constants $a_1 = a_2 = 1$, $b_1 = b_2 = 0.2$, and $b_3 = 0.4$, we have $\mu(A_i \cap B_j) = a_i b_j$ whenever $A_i \cap B_j \neq \emptyset$. For example, $\mu(A_1 \cap B_1) = \mu(\omega_1) = 1 \cdot 0.2 = 0.2$. Since we can provide a similar factorization for every prior μ, the partitions are pseudo-independent. It is important to note that when $A_i \cap B_j = \emptyset$, no condition is imposed.

Next, the definition of *informative loops* is taken from Lagziel et al. (2025) who uses it to compare mediators:

Definition 9. *An F-loop $\big((\omega_i, \bar{\omega}_i)\big)_{i=1}^m$ is* informative *if there exist an i such that ω_i and $\bar{\omega}_i$ belong to two different information sets of F.*

[16]Obviously, the set would also be convex in a binary state space, as any connected set is convex in one dimension.

[17]The two partitions are probabilistically independent with respect to μ if $\mu(A_i \cap B_j) = \mu(A_i)\mu(B_j)$ for every i, j.

An informative loop implies that, within a chain of states that players cannot distinguish, the mediator *can* distinguish some states. This capability allows the mediator to induce likelihood ratios that do not aggregate linearly, creating non-convexity.

In a single-player setting, the implication of $F(\omega_i) \neq F(\bar{\omega}_i)$ is that ω_i and $\bar{\omega}_i$ are indistinguishable to the single player, whereas the mediator can distinguish between them. In this sense, the mediator may convey additional information to the player along the loop. For this reason, the loop is called informative. For example, the F-loop $((\omega_1, \omega_2)(\omega_4, \omega_3))$ discussed in Example 2 is informative because ω_1 and ω_2 lie in different information sets of F. Had F been trivial (i.e., $F = \{\Omega\}$), so that the mediator knows nothing, the sequence $((\omega_1, \omega_2)(\omega_4, \omega_3))$ would still form an F-loop, but it would no longer be informative.

We are now ready to characterize the convexity of the feasible set. This characterization relies on two conditions. The first is a combinatorial condition on the mediator's and the player's partitions, ensuring pseudo-independent factorizations, and the second concerns the information (knowledge) available to the mediator along every loop. (The proof is relegated to the appendix; See Section A.4.)

Theorem 2. *Consider a player endowed with a partition $\mathcal{P}_1$ and a mediator whose partition is F. The following statements are equivalent.*

- *The feasible set is convex for every prior.*
- *$\mathcal{P}_1$ and F are pseudo-independent.*
- *There is no informative F-loop.*

Theorem 2 can be extended to show that, given an informative F-loop, the feasible set of posterior beliefs is not convex for almost every prior. In other words, a non-convex set of posterior beliefs is generic, in the presence of an informative F-loop. This result is given in the following Proposition 2. (The proof is relegated to the appendix; See Section A.5.)

Proposition 2 (Generic non-convexity under an informative F-loop)**.** *Consider a player endowed with a partition $\mathcal{P}_1$ and a mediator whose partition is F. If there exists an informative F-loop, then the set of priors for which the feasible set is convex has Lebesgue measure zero.*

Interestingly, there are various, quite standard persuasion models with an informative F-loop. For example, following Guo and Shmaya (2019), consider an online platform (Sender) that wants a customer (Receiver) to buy a product. The product has an unknown quality $q \in q_H, q_L$, where $0 < q_L < q_H < 1$, known only to the Sender, while the customer reads reviews and forms a private "type" t: either optimistic t_H (w.p. q) or pessimistic

q_L (w.p. $1-q$). Evidently, the optimistic type is easier to persuade about a high-quality product than the pessimistic one. The Receiver can either buy or not, and conditional on buying, his payoff depends on q (and 0 otherwise), while the Sender's payoff is positive only if the Receiver buys (and 0 otherwise). Thus, the state space comprises 4 states, where the Receiver's and the Sender's partitions are $\mathcal{P}_1 = \{\{q_H t_H, q_L t_H\}, \{q_H t_L, q_L t_L\}\}$ and $F = \{\{q_H t_H, q_H t_L\}, \{q_L t_H, q_L t_L\}\}$, respectively. This yields a 4-state, informative F-loop, so the feasible set is generically not convex.

5 Extensions and applications

5.1 A Mediator as a Blackwell-experiments generator

In the previous sections we provided necessary and sufficient conditions for a JB to be generated by a single F-measurable signaling function and a single signal. Now suppose we are given several beliefs, $\text{JB}_1, \ldots, \text{JB}_m$. As in Eq. (13), each JB_i induces a PL function φ_{JB_i}. Assume that each φ_{JB_i} satisfies the **[INC]** and **[EXC]** conditions, thereby enabling the application of Theorem 1. This means that for each i, there exists a signaling function τ_i and a signal s_i such that the posterior they induce coincides with $\mu_i = \mu_{\text{JB}_i}$, as in Eq. (12). Thus, each posterior is individually generated by a distinct signaling function and signal. Our question here is whether there exists a single signaling function τ that generates all of these posteriors with positive probability, and no others.

To formally introduce this question, denote

$$\mu_\tau(\omega \mid s) = \frac{\mu(\omega)\tau(s \mid \omega)}{\sum_{\omega' \in \Omega} \mu(\omega')\tau(s \mid \omega')}, \tag{14}$$

which is the posterior probability of ω given that the signal s has been generated by τ. The corresponding posterior distribution over Ω will be denoted by $\mu_\tau(\cdot \mid s)$. Note that the signaling function and the signal are constructed so as to retain all posterior probabilities consistent with the respective JBs and subject to the mediator's informational constraint. Once constructed, they induce a distribution over the state space obtained through standard Bayesian updating.

We then ask: under what conditions does there exist a signaling function τ, measurable with respect to F, such that the signals it produces with positive probability are $s_1, \ldots, s_m$ and

$$\big\{\mu_\tau(\cdot \mid s_1), \ldots, \mu_\tau(\cdot \mid s_m)\big\} = \big\{\mu_1, \ldots, \mu_m\big\}? \tag{15}$$

If the answer is affirmative, the mediator effectively serves as a generator of Blackwell experiments and, in particular, can generate the corresponding posteriors. To investigate this question, we begin with the following definition. A nonzero function $u : \Omega \to \mathbb{R}$ is called an *option* (or a *state-contingent claim*), as it specifies a monetary payoff for each possible state.

Definition 10. (i) *We say that the family of distributions* $\nu_1, \dots, \nu_k$ preserves positivity (PP) *w.r.t.* μ *if, for every option* u *such that* $\mathbb{E}_{\nu_i}[u] \geq 0$ *for every* i*, it follows that* $\mathbb{E}_{\mu}[u] \geq 0$. (ii) *We say that the family of distributions* $\nu_1, \dots, \nu_k$ strictly preserves positivity (SPP) *w.r.t.* μ *if, for every option* u *such that* $\mathbb{E}_{\mu}(u) = 0$ *and* $\mathbb{E}_{\nu_i}(u) \geq 0$ *for every* i*, it follows that* $\mathbb{E}_{\nu_i}(u) = 0$ *for every* i*.*

To better understand the notion of positivity preservation, suppose that μ is the prior belief a decision maker holds about the state space Ω, and let $\mu_1, \dots, \mu_m$ be the posteriors induced by observed signals. If the posteriors $\mu_1, \dots, \mu_m$ *do not* preserve positivity w.r.t. μ, then there exists a utility function assigning payoffs to states such that the expected utility under the prior μ is strictly lower than the expected utility under each posterior μ_i. This represents a case of time inconsistency: ex ante (i.e., under the prior), the option yields a negative expected reward, while ex post (i.e., conditional on any posterior), the same option yields a positive expected reward, regardless of the specific signal realized.

Strict positivity preservation w.r.t. μ means that the distributions $\mu_1, \dots, \mu_k$ are fully aligned with μ in the following sense: there is no option whose expected value is nonnegative under all μ_i, zero under μ, and strictly positive under at least one μ_i. Formally, if a function u has zero mean under μ, that is, μ regards it as "fair", and each μ_i weakly favors it (i.e., $\mathbb{E}_{\mu_i}[u] \geq 0$ for all i), then it must be that $\mathbb{E}_{\mu_i}[u] = 0$ for all i. In other words, for any option that is μ-neutral, if it is weakly favorable for all μ_i, then it is strictly favorable for none.

Observation 1. *To show that* SPP *w.r.t.* μ *implies* PP *w.r.t.* μ*, assume the former. Consider an option* u *such that* $\mathbb{E}_{\mu_i}[u] \geq 0$ *for every* i *and assume, contrary to* PP*, that* $\mathbb{E}_{\mu}[u] < 0$*. Define the option* $v = u - \mathbb{E}_{\mu}[u]$*. Then,* $\mathbb{E}_{\mu_i}[v] = \mathbb{E}_{\mu_i}[u] - \mathbb{E}_{\mu}[u] \geq 0$ *and* $\mathbb{E}_{\mu}[v] = \mathbb{E}_{\mu}[u] - \mathbb{E}_{\mu}[u] = 0$*. By* SPP*,* $\mathbb{E}_{\mu_i}[v] = 0$ *and thus* $\mathbb{E}_{\mu_i}[u] = \mathbb{E}_{\mu}[u] < 0$ *for every* i*, which contradict the assumption.*

We can now discuss the following Theorem 3 that extends previous results to a set of joint posteriors. Fix a family of PL functions $\varphi_{\mathrm{JB}_1}, \dots, \varphi_{\mathrm{JB}_m}$ (induced by the aforementioned JBs) that satisfy **[INC]** and **[EXC]**, and recall Eq. (12) which prescribes a distribution for every such PL function. In particular, φ_i corresponds to the posterior μ_i. The first part of the following Theorem 3 states that PP is a necessary and sufficient condition for the existence of a strategy τ whose signals generate posterior distributions, all of which are contained in

$\{\mu_1, \ldots, \mu_m\}$. The second part of the theorem uses SPP to characterize when there is a signaling function whose set of posterior coincides with $\{\mu_1, \ldots, \mu_m\}$. (The proof is given in Section A.6 in the appendix.)

Theorem 3. *Let μ be a common prior, and let $\varphi_{\text{JB}_1}, \ldots, \varphi_{\text{JB}_m}$ be* PL *functions that satisfy* **[INC]** *and* **[EXC]**. *For every* JB_i, *assume that Ω^i_+, defined as in* Theorem 1, *is measurable w.r.t. F. Then, there exists an F-measurable signaling function τ such that:*
(i) for every signal s generated by τ with positive probability, the posterior it induces belongs to the family $\{\mu_1, \ldots, \mu_m\}$ if and only if this family PP *w.r.t. μ;*
(ii) the set of posteriors it generates with positive probability coincides with $\{\mu_1, \ldots, \mu_m\}$ if and only if this family SPP *w.r.t. μ.*

One should note the construction underlying Theorem 3. First, a common prior and a set of JBs are fixed to satisfy the conditions of Theorem 1. Next, the implied strategies are employed to construct a set of posteriors over Ω that abstract from the players' private information. The necessary and sufficient conditions stated in Theorem 3 are then imposed on these posteriors, relative to the prior, to guarantee the existence of an F-measurable unified strategy. Once this strategy is applied, together with the common prior and the players' private information, it induces the desired JBs.

5.2 Fact-checking and depolarization procedures

The issue of polarization, either ideological or affective, has gained much concern in public discourse, and therefore attention both in general and in research.[18] Yet, there has been considerably less focus on methods of *reducing* polarization that naturally align with our framework, and can provide insights into this issue.

For that purpose, we interpret the mediator as a fact-checker (or some other verification / regulatory authority) operating on a news or social-media platform. The underlying state $\omega \in \Omega$ describes both the true content of a story and the relevant context. Users are partitioned into groups $i \in N$ with information partitions $\mathcal{P}_i$ ("echo chambers"), and the fact-checker has a partition F capturing the evidence it can access (e.g., databases, legal documents). Accordingly, the fact-checking policy is an F-measurable public experiment $\tau : \Omega \to \Delta(S)$ whose realizations $s \in S$ are labels such as "True", "Misleading", or "False".

[18]See, for instance, Nimark and Sundaresan (2019); Matvejka and Tabellini (2021); Levy (2021); Yuksel (2022); Perego and Yuksel (2022); Bowen et al. (2023); Ikan et al. (2025); Goldstein et al. (2025) among many others.

Given $p \in [0,1]$, we say that a strategy τ induces *p–consensus* if

$$D(\tau) = \mathbb{E}_{\mu,\tau}\left[\max_{i,j\in N} \|\big(\mu_{\tau,i}(\cdot \mid \mathcal{P}_i(\omega), s)\big) - \big(\mu_{\tau,j}(\cdot \mid \mathcal{P}_j(\omega), s)\big)\|\right] \leq 2(1-p),$$

where $\|\cdot\|$ is the l_1-norm and $D(\tau)$ is referred to as the *expected disagreement* given τ. Note that the case where players have a common prior, and no private information relates to $p = 1$, a situation we refer to as *a full consensus*.

Evidently, the case of full consensus is very demanding. It implies that all disagreements completely vanish, and intuitively, this would require an overwhelming amount of information on the side of the mediator. The following proposition captures this. (The proof is straightforward, thus omitted.)

Proposition 3. *The mediator can induce full consensus for every common prior if and only if $\mathcal{P}_i \vee F = \mathcal{P}_j \vee F$ for every $i, j \in N$.*[19]

The characterizing criterion is quite restrictive: it requires the mediator to bridge *any* informational gap between any two individuals. In many environments this is unrealistic. A potentially more reasonable and modest objective is to reduce polarization without necessarily eliminating it completely, that is, achieving p-consensus for some $p < 1$. In this case, our analysis and results become much relevant.

The mediator's problem in this section can be interpreted as choosing an F-measurable experiment τ to maximize consensus, or equivalently to minimize $D(\tau)$. Crucially, Theorem 2 and the consistency conditions [INC] and [EXC] characterize the *feasible* joint posteriors that can arise from such a τ, so $D(\tau)$ is minimized over a set that is typically non-convex and constrained by products of likelihood ratios along F-cycles and F-loops.

The next example illustrates that, under this objective, full revelation of the mediator's partition F is generally *not* optimal: a strictly less informative, noisy experiment can yield strictly lower expected disagreement.

Example 4. Consider a probability space with four states $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ with the prior $\mu = \left(\frac{1}{16}, \frac{4}{16}, \frac{5}{16}, \frac{6}{16}\right)$. There are two players and a mediator whose partitions are given by $\mathcal{P}_1 = \big\{\{\omega_1, \omega_3, \omega_4\},\ \{\omega_2\}\big\}$, $\mathcal{P}_2 = \big\{\{\omega_1\},\ \{\omega_2, \omega_3, \omega_4\}\big\}$, and $F = \big\{\{\omega_1, \omega_3\},\ \{\omega_2, \omega_4\}\big\}$, respectively.

Given that the mediator provides no information (a constant signal), each player simply conditions on $\mathcal{P}_i(\omega)$, and the resulting expected disagreement is approximately 0.848. However, under a fully revealing policy, the expected disagreement drops to 0.808, so full

[19] $\mathcal{P}_i \vee F$ is the coarsest common refinement of $\mathcal{P}_i$ and F, i.e., the partition generated by intersecting their cells.

revelation of F is helpful in reducing polarization relative to having no mediator. But the mediator could even do better. For example, consider a strictly less informative τ^{noisy} with two signals s_1, s_2:

$$\tau^{\text{noisy}}(s_1 \mid \omega_1) = \tau^{\text{noisy}}(s_1 \mid \omega_3) = 0, \qquad \tau^{\text{noisy}}(s_1 \mid \omega_2) = \tau^{\text{noisy}}(s_1 \mid \omega_4) = \tfrac{3}{4},$$

and $\tau^{\text{noisy}}(s_2 \mid \omega) = 1 - \tau^{\text{noisy}}(s_1 \mid \omega)$ for all ω. This experiment partially pools the two F-cells and is therefore strictly less informative (in the Blackwell sense) than the completely informative strategy. A direct calculation shows that the expected disagreement under τ^{noisy} is 0.775.

Thus, although full revelation of F reduces polarization relative to having no mediator, it is *strictly dominated* (for the purpose of minimizing expected disagreement) by a coarser, noisy F-measurable experiment. The reason is that fully revealing F interacts with the heterogeneous partitions $(\mathcal{P}_i)_i$ and, through the consistency constraints in Theorem 1, creates states at which players' posteriors diverge sharply. By partially pooling F-cells, the mediator can smooth these differences and achieve a higher level of approximate consensus.

Example 4 shows that once we move from full consensus to a more reasonable objective, the mediator faces a genuine information-design problem. The graph-of-information structure and the internal and external consistency conditions delineate the feasible joint posteriors, but within this feasible set, it is generally *not* optimal to reveal all available information: carefully designed garblings of F can reduce polarization more effectively.

5.3 A general framework for information design and persuasion

Our study provides a general information framework that could be applied in various settings within the realm of information design. For example, in a benchmark binary state-space setting, where players have no private information and with a fully informed mediator, Theorems 1 and 3 collapse to the standard splitting lemma: Bayes plausibility is the only feasibility constraint. This recovers the classical concavification approach of Aumann and Maschler (1995) and Kamenica and Gentzkow (2011) as a special case of our characterization.

In binary environments, our graph-based constraints are vacuous, and implementability reduces to the martingale conditions studied by Dawid et al. (1995), Arieli et al. (2021) and Morris (2020) for feasible joint posterior beliefs. Our Theorem 3 extends previous results to arbitrary finite state spaces and to mediators with partial information.

In addition, our framework subsumes the information structures used in the optimal disclosure in auctions literature. In models such as Eső and Szentes (2007), handicap auction and related work on optimal information disclosure in auctions, a seller observes a state (e.g., product quality or the bidders' value profile) and then commits to a disclosure rule before running a standard auction. Likewise, in Bergemann and Pesendorfer (2007) and subsequent analyses of optimal information disclosure in classic auctions, the seller's disclosure rule is explicitly modeled as an information structure or signal device mapping the realized state into messages seen by bidders. When the seller does not have perfect information, or bidders also hold private signals (as in more general information-design problems with privately informed agents, as in Candogan and Strack, 2023) our internal and external consistency conditions describe exactly which belief patterns about values can be generated by any feasible disclosure policy.

In our language, the underlying state is a realization ω, each bidder's ex ante information is a partition $\mathcal{P}_i$ (often trivial before the seller's message), and the seller/auctioneer is exactly a mediator whose knowledge is described by a partition F (typically the full, singleton partition when the seller is fully informed). The disclosure rule is an F-measurable Blackwell experiment, so for any such auction environment, the induced beliefs across bidders and states form a joint posterior to which Theorems 1 and 3 apply.

Evidently, these extensions imply that the mediator (from our framework) has some goal function, and this raises a crucial question about the concavification technique that is typically used in the literature. We discuss this issue in the two following sections.

5.4 Mediator's optimization

To address the question of mediator optimization, one should first specify the objective function of the mediator. In the leading example discussed above, one could consider a benevolent mediator who seeks, for instance, to maximize the peace chances, or the probability that the (C, C) outcome emerges in equilibrium. Another example of such an objective function could be to maximize social welfare, defined as a weighted average of the players' equilibrium expected payoffs.

Once the objective function is fully specified, one must consider the issue of multiple equilibria. Upon observing the signal issued by the mediator, all players update their beliefs and then play an incomplete-information game. Such a game typically has multiple equilibria, over which the mediator has no direct control. Any equilibrium from this set could emerge. Of course, one could adopt the one that maximizes the objective function, but this choice is arguably as arbitrary as any other.

Finally, once the objective function and the target equilibrium are fixed, the mediator must solve an optimization problem through optimal splitting, which may vary significantly from the ideas of concavification of Aumann and Maschler (1995) and Kamenica and Gentzkow, 2011. The divergence from concavification arises from the non-convexity of the feasible set, as studied in Section 4.

5.5 An additive interpretation and implications to potential games

The graph-theoretic constraints derived for Bayesian updating reflect a structural property applicable beyond information design. By applying a logarithmic transformation to the internal consistency condition, we obtain an additive constraint that characterizes *Potential Games* (Monderer and Shapley, 1996).

Consider a directed graph $G = (V, E)$ with edge weights $\rho : E \to \mathbb{R}$. We define *Additive Internal Consistency* **[INC-ADD]** as the requirement that for every cycle $(\omega_1, \ldots, \omega_{m+1} = \omega_1)$:

$$\sum_{i=1}^{m} \rho(\omega_i, \omega_{i+1}) = 0. \tag{16}$$

By applying the transformation $g = \log f$ to Proposition 1, it follows that **[INC-ADD]** is necessary and sufficient for the existence of a potential function $P : V \to \mathbb{R}$ such that $\rho(\omega, \omega') = P(\omega) - P(\omega')$.

This structure applies directly to strategic games. Let the vertices of the *game graph* be action profiles, with directed edges representing unilateral deviations. In this context, a cycle corresponds to a sequence of unilateral deviations that returns to the initial action profile.

We weight each deviation by the deviator's utility difference:

$$\rho\left(a, (a'_i, a_{-i})\right) := u_i(a'_i, a_{-i}) - u_i(a).$$

Corollary 1. *A strategic game Γ is a Potential Game if and only if the associated edge function ρ satisfies **[INC-ADD]** on every cycle of the game graph.*

This observation unifies the two frameworks: the mediator constructs a global belief system consistent with local likelihood ratios, while a potential game implies a global preference ordering consistent with local unilateral deviations. In both cases, the obstruction to "globalizing" local information is the non-vanishing of accumulated weights along cycles in the underlying graph.

6 Conclusion

This paper analyzes the implementation of joint posterior beliefs in environments where agents possess private information and a non-strategic mediator holds partial information. By modeling the informational constraints of the mediator as measurability restrictions with respect to a specific partition, we provide a necessary and sufficient characterization for the rationalization of joint beliefs.

The core theoretical contribution lies in the identification of two structural conditions: *Internal Consistency* (INC) and *External Consistency* (EXC). The study demonstrates that a joint posterior is implementable if and only if the product of likelihood ratios along every cycle (within a Common Knowledge Component) and every loop (across components connected by the mediator's information) equals one. This characterization effectively translates abstract measurability constraints into verifiable algebraic properties of the posterior likelihood function.

A significant implication of this analysis is the finding that the set of feasible joint posteriors is generally *non-convex*. Unlike standard Bayesian persuasion frameworks, such as Kamenica and Gentzkow (2011), where feasibility is determined solely by Bayes plausibility, this study shows that informational constraints alone can preclude convexity. We further characterize the specific conditions, namely pseudo-independence and the absence of informative loops, under which convexity is restored.

Finally, the paper extends these insights to broader economic applications. It generalizes the framework to the generation of multiple posteriors via the *preservation of positivity* condition. Furthermore, the developed graph-theoretic approach is shown to be relevant for diverse fields, offering a connection to potential games through additive internal consistency and providing insights into belief depolarization, where coarser information structures may outperform full revelation in minimizing disagreement.

References

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

Aumann, R. J. (1976): "Agreeing to Disagree," *The Annals of Statistics*, 4, 1236–1239, publisher: Institute of Mathematical Statistics.

Aumann, R. J. and M. Maschler (1995): *Repeated Games with Incomplete Information*, Cambridge, MA, USA: MIT Press.

Ball, I. and J.-A. Espín-Sánchez (2021): "Experimental Persuasion," .

Bergemann, D. and M. Pesendorfer (2007): "Information structures in optimal auctions," *Journal of Economic Theory*, 137, 580–609.

Blackwell, D. (1951): "Comparison of Experiments," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley, Calif.: University of California Press, 93–102.

Bowen, T. R., D. Dmitriev, and S. Galperti (2023): "Learning from Shared News: When Abundant Information Leads to Belief Polarization*," *The Quarterly Journal of Economics*, 138, 955–1000.

Burdzy, K. and J. Pitman (2020): "Bounds on the probability of radically different opinions," *Electronic Communications in Probability*, 25, 1–12, publisher: Institute of Mathematical Statistics and Bernoulli Society.

Candogan, O. and P. Strack (2023): "Optimal disclosure of information to privately informed agents," *Theoretical Economics*, 18, 1225–1269, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/TE5173.

Chen, Y., T. Lin, W. Tang, and J. Tucker-Foltz (2025): "Explainable Information Design," ArXiv:2508.14196 [cs].

Dawid, A. P., M. H. DeGroot, J. Mortera, R. Cooke, S. French, C. Genest, M. J. Schervish, D. V. Lindley, K. J. McConway, and R. L. Winkler (1995): "Coherent combination of experts' opinions," *Test*, 4, 263–313.

Eső, P. and B. Szentes (2007): "Optimal Information Disclosure in Auctions and the Handicap Auction," *The Review of Economic Studies*, 74, 705–731.

Feinberg, Y. (2000): "Characterizing Common Priors in the Form of Posteriors," *Journal of Economic Theory*, 91, 127–179.

Geanakoplos, J. (1994): "Chapter 40 Common knowledge," in *Handbook of Game Theory with Economic Applications*, Elsevier, vol. 2, 1437–1496.

Goldstein, N., D. Lagziel, and O. Raveh (2025): "Political Rational Inattention: A New Measure With an Application to Political Polarization," .

Guo, Y. and E. Shmaya (2019): "The Interval Structure of Optimal Disclosure," *Econometrica*, 87, 653–675, publisher: [Wiley, The Econometric Society].

Gutmann, S., J. H. B. Kemperman, J. A. Reeds, and L. A. Shepp (1991): "Existence of Probability Measures with Given Marginals," *The Annals of Probability*, 19, 1781–1797, publisher: Institute of Mathematical Statistics.

Harsanyi, J. C. (1967–1968): "Games with Incomplete Information Played by 'Bayesian' Players, Parts I–III," *Management Science*, 14, 159–182, 320–334, and 486–502, part I: The Basic Model; Part II: Bayesian Equilibrium Points; Part III: The Basic Probability Distribution of the Game.

Heifetz, A. (2006): "The Positive Foundation of the Common Prior Assumption," *Games and Economic Behavior*, 56, 105–120.

Hellman, Z. and D. Samet (2012): "How common are common priors?" *Games and Economic Behavior*, 75, 207–218.

Hellwig, M. F. (2013): "From posteriors to priors via cycles: An addendum," *Economics Letters*, 118, 455–458.

Herings, P. J.-J., D. Karos, and T. T. Kerman (2024): "Belief inducibility and informativeness," *Theory and Decision*, 96, 517–553.

Hörner, J., M. Morelli, and F. Squintani (2015): "Mediation and Peace," *The Review of Economic Studies*, 82, 1483–1501, publisher: [Oxford University Press, The Review of Economic Studies, Ltd.].

Ikan, L., D. Lagziel, and O. Raveh (2025): "Resource windfalls, connectivity, and political polarization," *Journal of Environmental Economics and Management*, 132, 103164.

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

Kellerer, H. G. (1961): "Funktionen auf Produkträumen mit vorgegebenen Marginal-Funktionen," *Mathematische Annalen*, 144, 323–344.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017): "Persuasion of a Privately Informed Receiver," *Econometrica*, 85, 1949–1964.

Kosenko, A. (2020): "Mediated Persuasion," ArXiv:2012.00098 [econ].

Lagziel, D., E. Lehrer, and T. Wang (2025): "Comparison of Oracles: Part I," ArXiv:2505.15955 [econ].

Lehrer, E. and D. Samet (2014): "Belief Consistency and Trade Consistency," *Games and Economic Behavior*, 83, 165–177.

Levy, R. (2021): "Social Media, News Consumption, and Polarization: Evidence from a Field Experiment," *American Economic Review*, 111, 831–870.

Matvejka, F. and G. Tabellini (2021): "Electoral Competition with Rationally Inattentive Voters," *Journal of the European Economic Association*, 19, 1899–1935.

Monderer, D. and L. S. Shapley (1996): "Potential games," *Games and Economic Behavior*, 14, 124–143.

Morris, S. (1994): "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, 62, 1327–1347.

——— (2020): "No Trade and Feasible Joint Posterior Beliefs," .

Nimark, K. P. and S. Sundaresan (2019): "Inattention and belief polarization," *Journal of Economic Theory*, 180, 203–228.

Perego, J. and S. Yuksel (2022): "Media Competition and Social Disagreement," *Econometrica*, 90, 223–265, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA16417.

Rockafellar, R. T. (1970): *Convex Analysis*, Princeton, NJ: Princeton University Press.

Rodrigues-Neto, J. A. (2009): "From posteriors to priors via cycles," *Journal of Economic Theory*, 144, 876–883.

Strassen, V. (1965): "The Existence of Probability Measures with Given Marginals," *The Annals of Mathematical Statistics*, 36, 423–439, publisher: Institute of Mathematical Statistics.

Yuksel, S. (2022): "Specialized Learning and Political Polarization," *International Economic Review*, 63, 457–474, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12555.

Ziegler, G. (2020): "Adversarial bilateral information design," .

Özyurt, S. and Y. Zeng (2025): "Peaceful Dispute Resolution with an Imperfectly Informed Mediator," .

A Proofs

A.1 Proof of Lemma 1

For every pair (ω_1, ω_{m+1}) such that $\omega_1 \twoheadrightarrow \omega_{m+1}$, define

$$\varphi(\omega_1, \omega_{m+1}) := \prod_{i=1}^{m} \varphi(\omega_i, \omega_{i+1}),$$

provided that the sequence $\big((\omega_i, \omega_{i+1})\big)_{i=1}^{m}$ lies in E. To show that this is well defined, suppose there are two sequences connecting ω_1 and ω_{m+1}:

$$\big((\omega_i, \omega_{i+1})\big)_{i=1}^{m} \quad \text{and} \quad \big((\xi_j, \xi_{j+1})\big)_{j=1}^{k},$$

with $\xi_1 = \omega_1$ and $\xi_{k+1} = \omega_{m+1}$. Concatenating the two sequences, with the first in its original order and the second reversed, produces an F-cycle.

By **[INC]** and Eq. (5), we have

$$\prod_{i=1}^{m} \varphi(\omega_i, \omega_{i+1}) \cdot \prod_{j=1}^{k} \varphi(\xi_{j+1}, \xi_j) = \prod_{i=1}^{m} \varphi(\omega_i, \omega_{i+1}) \cdot \prod_{j=1}^{k} \frac{1}{\varphi(\xi_j, \xi_{j+1})} = 1.$$

Therefore,

$$\prod_{i=1}^{m} \varphi(\omega_i, \omega_{i+1}) = \prod_{j=1}^{k} \varphi(\xi_j, \xi_{j+1}),$$

shows that the definition of $\varphi(\omega_1, \omega_{m+1})$ does not depend on the particular sequence and is therefore well defined. $\square$

A.2 The proof of Proposition 1.

Part 1: Necessity. Suppose that there is an F-measurable positive function f that satisfies Eq. (11). To show **[INC]**, let $\big((\omega_i, \omega_{i+1})\big)_{i=1}^{m}$ be an F-cycle. Then,

$$\prod_{i=1}^{m} \varphi(\omega_i, \omega_{i+1}) = \prod_{i=1}^{m} \frac{f(\omega_i)}{f(\omega_{i+1})} = \frac{f(\omega_1)}{f(\omega_{m+1})} = 1.$$

The first equality is due to Eq. 11, and the last follows from the assumption that f is F-measurable and that $\omega_{m+1} \in F(\omega_1)$.

To show **[EXC]**, let $\big((\omega_i,\bar{\omega}_i)\big)_{i=1}^m$ be an F-loop. By Eq. (8) and Eq. (11), we have

$$\prod_{i=1}^m \varphi(\omega_i,\bar{\omega}_i) = \prod_{i=1}^m \frac{f(\omega_i)}{f(\bar{\omega}_i)} = \frac{f(\omega_1)}{f(\bar{\omega}_m)} = 1.$$

The first equality follows from Eq. (11), the second equality follows from the fact that f is F-measurable and the second property of F-loops, and the third equality follows from the same property, since $\bar{\omega}_m \in F(\omega_1)$.

Part 2: Sufficiency. Assume **[INC]** and **[EXC]** with the extension provided by Lemma 1. We claim first that for every connected set of V, say C, there is an F-measurable positive function f_C, defined over C, that satisfies Eq. (11). Let ω_0 be an arbitrary state in C. Set $f(\omega_0) = 1$. We proceed by induction on the distance[20] to ω_0.

Suppose that f has been defined on all states in a connected component C whose distance from the state ω_0 is k. Let ω' be a state in C at distance $k+1$ from ω_0. Then, there exists a state ω'' at distance k such that $(\omega'',\omega') \in E$. Define

$$f_C(\omega') := f_C(\omega'') \cdot \varphi(\omega',\omega'').$$

Using **[INC]** and the same argument as in Lemma 1, one can show that f is well defined. To show that f_C is F-measurable, consider $\omega \in F(\omega') \cap C$. Then, there is an F-cycle $\omega_1 = \omega, ..., \omega_{m+1} = \omega'$. By **[INC]** and by the definition of f_C,

$$1 = \prod_{i=1}^m \varphi(\omega_i,\omega_{i+1}) = \prod_{i=1}^m \frac{f_C(\omega_i)}{f_C(\omega_{i+1})} = \frac{f_C(\omega_1)}{f_C(\omega_{m+1})} = \frac{f_C(\omega)}{f_C(\omega')}.$$

Thus, the values that f_C takes on ω and ω' coincide. So far, we have constructed f_C for each connected component C of V. We now show that there exists a single F-measurable function f satisfying Eq. (11).

To this end, we define a new graph, $\bar{G} = (\bar{V},\bar{E})$ as follows. The set of vertices, $\bar{V}$, consists of the connected components of V. Two components $C, C' \in \bar{E}$ are connected by an edge if there exist states $\omega \in C$ and $\omega' \in C'$ such that $\omega \in F(\omega')$. That is, the edge set, $\bar{E}$, consists of all pairs (C, C') for which there exist $\omega \in C$ and $\omega' \in C'$ such that ω and ω' belong to the same information set of F.

Fix a vertex C_0, and set $f := f_{C_0}$ on C_0. We proceed by induction on the distance of a vertex C from C_0. Suppose that an F-measurable function $f > 0$, satisfying Eq. (11), has

[20] The distance between two states in C is defined as the minimal number of edges in a path connecting them.

been defined on all vertices whose distance from C_0 is less than or equal to k. Let C be a vertex at distance $k+1$ from C_0. Then there exists a path of edges in $\bar{E}$,

$$(C_0, C_1), (C_1, C_2), \ldots, (C_k, C),$$

and in particular, there exist $\bar{\omega}_k \in C_k$ and $\omega \in C$ such that ω_k and ω belong to the same information set of F. Define

$$f := f_C \cdot \frac{f(\bar{\omega}_k)}{f_C(\omega)} \tag{17}$$

on C. Clearly, if f is well defined in C, then it satisfies Eq. (11), since it differs from f_C by a multiplicative factor of a positive constant. Moreover, $f(\omega) = f(\bar{\omega}_k)$.

We first verify that f is well defined and then show that it is F-measurable. To establish that f is well defined, suppose that there exists another path

$$(C_0, C'_1), (C'_1, C'_2), \ldots, (C'_\ell, C),$$

connecting C_0 and C, with $\omega'_\ell \in C'_\ell$ and $\omega' \in C$ such that ω'_ℓ and ω' belong to the same information set of F.[21] We have to show that

$$\frac{f(\bar{\omega}_k)}{f_C(\omega)} = \frac{f(\omega'_\ell)}{f_C(\omega')}. \tag{18}$$

Note that the path

$$(C_0, C_1), (C_1, C_2), \ldots, (C_k, C), (C, C'_\ell), \ldots, (C'_2, C'_1), (C'_1, C_0)$$

is a cycle in the graph $\bar{G}$, that corresponds to an F-loop:

$$(\omega_0, \bar{\omega}_0), (\omega_1, \bar{\omega}_1), \ldots, (\omega_{k-1}, \bar{\omega}_{k-1}), (\omega_k, \bar{\omega}_k), (\omega, \omega'), (\omega'_\ell, \bar{\omega}'_\ell), \ldots, (\omega'_2, \bar{\omega}'_2), (\omega'_1, \bar{\omega}'_1).$$

That is, any two states in a pair belong to the same vertex in $\bar{G}$, e.g., ω_0 and $\bar{\omega}_0$ belong to C_0, ω_1 and $\bar{\omega}_1$ belong to C_1; ω and ω' belong to C, and $\omega'_j, \bar{\omega}'_j$ belong to C'_j. Also, in two adjacent pairs, the second coordinate of the left pair belongs to the same information set of F as the first coordinate of the right pair. For instance, $\bar{\omega}_i \in F(\omega_{i+1})$, $\bar{\omega}_k \in F(\omega)$, $\omega' \in F(\omega'_\ell)$ and $\omega'_{j+1} \in F(\bar{\omega}'_j)$. Finally, $\bar{\omega}'_1 \in F(\omega_0)$.

[21] The asymmetry in the notation between $\bar{\omega}_k$ (with a bar) and ω'_ℓ (without) will become clear immediately.

By [**EXC**], and specifically by Eq. (10), the corresponding product equals 1. That is,

$$\left[\prod_{i=0}^{k}\varphi(\omega_i,\bar{\omega}_i)\right]\cdot\varphi(\omega,\omega')\cdot\left[\prod_{j=\ell}^{1}\varphi(\omega'_j,\bar{\omega}'_j)\right]=1. \tag{19}$$

Due to F-measurability and to Eq. (11), we get $\prod_{i=0}^{k}\varphi(\omega_i,\bar{\omega}_i)=\frac{f(\omega_0)}{f(\bar{\omega}_k)}$, $\varphi(\omega,\omega')=\frac{f_C(\omega)}{f_C(\omega')}$, and $\prod_{j=\ell}^{1}\varphi(\omega'_j,\bar{\omega}'_j)=\frac{f(\omega'_k)}{f(\bar{\omega}'_1)}$. Thus, by Eq. (19) we obtain the following:

$$\frac{f(\omega_0)}{f(\bar{\omega}_k)}\cdot\frac{f_C(\omega)}{f_C(\omega')}\cdot\frac{f(\omega'_\ell)}{f(\bar{\omega}'_1)}=1.$$

Due to F-measurability $f(\omega_0)=f(\bar{\omega}'_1)$ and we conclude that

$$\frac{f(\omega'_\ell)}{f(\bar{\omega}_k)}\cdot\frac{f_C(\omega)}{f_C(\omega')}=1.$$

which confirms Eq. (18).

The final step in the proof is to show that f is F-measurable. We consider two cases:

Case I. Let $\omega\in C$ and $\omega'\in C'$, where the distance from C' to C_0 is less than or equal to k, and suppose that $\omega\in F(\omega')$. We need to show that $f(\omega)=f(\omega')$. By assumption, $(C,C')\in\bar{E}$ and $(C,C_k)\in\bar{E}$. Moreover, C' and C_k are connected by paths of length at most k. This forms a cycle in the graph $\bar{G}$, which in turn corresponds to an F-loop. Applying the same technique used earlier, we conclude that $f(\omega)=f(\omega')$.

Case II. Let $\omega\in C$ and $\omega'\in C'$, where the distance from C' to C_0 is exactly $k+1$, so that both C and C' are introduced in the induction process at the same step, and assume that $\omega\in F(\omega')$. Then $(C,C')\in\bar{E}$, and both C and C' are connected to C_0 by paths of length $k+1$. This forms a cycle in $\bar{G}$ passing through C and C', which corresponds to an F-loop. Using [**EXC**], we conclude that $f(\omega)=f(\omega')$, verifying that f is F-measurable. □

A.3 The proof of Theorem 1

Proof. Assume that τ is a signaling function and that s is a signal generated by τ with positive probability such that $\boldsymbol{\mu}_{\tau,s}=\text{JB}$. Then, clearly, the set of states ω such that $\tau(s\mid\omega)>0$ is measurable w.r.t. the mediator's partition F, and matches the set Ω_+, previously defined according to JB.

Similarly to Eq. (4), the equality $\boldsymbol{\mu}_{\tau,s}=\text{JB}$ implies that $\frac{\tau(s|\omega)\mu(\omega)}{\tau(s|\omega')\mu(\omega')}=\frac{\text{JB}(\omega,i)(\omega)}{\text{JB}(\omega',i)(\omega')}$. Applying

Eq. (13), we get

$$\varphi_{\mathrm{JB}}(\omega,\omega') = \frac{\mathrm{JB}(\omega,i)(\omega)}{\mathrm{JB}(\omega',i)(\omega')}\cdot\frac{\mu(\omega')}{\mu(\omega)} = \frac{\tau(s|\omega)}{\tau(s|\omega')},$$

and Proposition 1 guarantees that φ_{JB} satisfies conditions **[INC]** and **[EXC]** in G_+. This proves the necessary direction.

For the sufficiency direction, assume that JB satisfies conditions (i) and (ii). Since Ω_+ is measurable w.r.t. F, we can define $\tau(s \mid \omega) = 0$ for every $\omega \notin \Omega_+$, and otherwise, $\tau(s \mid \omega)$ is positive.

On Ω_+, since φ_{JB} satisfies **[INC]** and **[EXC]** in G_+, Proposition 1 ensures the existence of a function $f > 0$ satisfying Eq. (11) and we get

$$\frac{\mathrm{JB}(\omega,i)(\omega)}{\mathrm{JB}(\omega',i)(\omega')}\cdot\frac{\mu(\omega')}{\mu(\omega)} = \frac{f(\omega)}{f(\omega')},$$

By multiplying f by a constant if needed, we may assume without loss of generality that $\sup_\omega f(\omega) < 1$. We then define $\tau(s \mid \omega) = f(\omega)$ for all $\omega \in \Omega_+$, so that the previous equation can be written as

$$\frac{\tau(s \mid \omega)\mu(\omega)}{\tau(s \mid \omega')\mu(\omega')} = \frac{\mathrm{JB}(\omega,i)(\omega)}{\mathrm{JB}(\omega',i)(\omega')},$$

which matches Eq. (4). This guarantees that $\tau(s \mid \cdot)$ implements JB, as needed. To ensure that τ is a well-defined kernel, we can extend it using an arbitrary signal s_0, so that $\tau(s_0|\omega) = 1 - f(\omega)$ for every $\omega \in \Omega_+$. □

A.4 The proof of Theorem 2

Proof. We start by proving that no informative F-loops implies pseudo-independence. Define the bipartite graph $G = (\mathcal{P}_1, F, E)$ where $(p, f) \in E$ if there exists $\omega \in p \cap f$. So two vertices p and f are connected if and only if they share a state $\omega \in \Omega$. Notice that a cycle exists in G if and only if an informative loops exists. To see this, consider a cycle whose length is even and at least 4 (because this is a bipartite graph). Take the states that yield the edges along this cycle. This sequence of states constitutes a loop across at least 4 different CKCs, two from each side of the graph, thus an informative loop. The same applies in the other direction. We therefore conclude that each connected component in G is a tree.

For every $p \in \mathcal{P}_1$ and $f \in F$ define $m_{p,f} = \mu(p \cap f)$. Take a tree in G, denote a root by $p_0 \in \mathcal{P}_1$ and define $a_{p_0} = 1$. For every f_0 connected to p_0 in G define $b_{f_0} = \mu(p_0 \cap f_0)$. Fix an f_0 and for every p_1 connected to it (in G; not including p_0) define $a_{p_1} = \frac{\mu(p_1 \cap f_0)}{b_{f_0}}$. Continue inductively and similarly for every tree in G, until all parameters a_p and b_f are defined, to conclude that $\mathcal{P}_1$ and F are pseudo-independent.

Next, we prove that the stated pseudo-independence implies a convex feasible set. Consider a feasible posterior belief $\mu_{\tau,1}(\cdot \mid p, s)$ given a signal s and a partition element $p \in \mathcal{P}_1$. For every $\omega \in p \cap \tau^{-1}(s)$, we get

$$\begin{aligned}
\mu_{\tau,1}(\omega \mid p, s) &= \frac{\tau(s|F(\omega))\mu(\omega)}{\sum_{f' \in F} \tau(s|f')\mu(p \cap f')} = \frac{\tau(s|F(\omega)) \sum_{f \in F} \mu(\omega|p \cap f)\mu(p \cap f)}{\sum_{f' \in F} \tau(s|f')\mu(p \cap f')} \\
&= \sum_{f \in F} \left[\frac{\tau(s|f)\mu(p \cap f)}{\sum_{f' \in F} \tau(s|f')\mu(p \cap f')} \right] \mu(\omega|p \cap f) \\
&= \sum_{f \in F} \left[\frac{\tau(s|f)a_p b_f}{\sum_{f' \in F} \tau(s|f')a_p b_{f'}} \right] \mu(\omega|p \cap f) \\
&= \sum_{f \in F} \left[\frac{\tau(s|f)b_f}{\sum_{f' \in F} \tau(s|f')b_{f'}} \right] \mu(\omega|p \cap f),
\end{aligned}$$

where the second line follows the fact that $\mu(\omega|p \cap f) = 0$ if $\omega \notin f$ and the third line follows from pseudo-independence. Define $\lambda_{s,f} = \frac{\tau(s|f)b_f}{\sum_{f' \in F} \tau(s|f')b_{f'}}$ and note that these weights are independent of any partition element $p \in \mathcal{P}_1$. Thus, for every set of partition elements/CKCs $p_1 \subseteq \mathcal{P}_1$, every feasible posterior belief generated by signal s is represented by a probability vector $\lambda_s \in \Delta(F)$. Thus, the map $\lambda_s \mapsto \left(\sum_f \lambda_{s,f}\mu(\cdot|p \cap f) \right)_{p \in p_1 \subseteq \mathcal{P}_1}$ is affine and defined over a convex set $\Delta(F)$, so the feasible set is indeed convex.

For the last part, assume by contradiction that there exists an informative loop L. Reduce it to a minimal size, so it intersects any CKC at most once (one can consider the equivalent cycle in G and contracting it). Assume, w.l.o.g., that $\overline{\omega}_1 \notin F(\omega_1)$. For every $\omega \notin L$, define $\mu(\omega) = 0$ (we can later assign a sufficiently small probability for any such state to maintain a strictly positive prior), and for every $\omega \in L \setminus \{\overline{\omega}_1\}$ set $\mu(\overline{\omega}_1) = 2\mu(\omega)$. This establishes a prior μ.

Define the signaling function τ as follows: $\tau(s_1|F(\overline{\omega}_1)) = 1/3 = 1 - \tau(s_2|F(\overline{\omega}_1))$ and $\tau(s_1|F \setminus F(\overline{\omega}_1)) = 2/3 = 1 - \tau(s_2|F \setminus F(\overline{\omega}_1))$. This strategy induces two feasible posterior beliefs, one for each signal, defined over the loop L. We can take the average of two posterior beliefs and check whether it is also feasible, specifically using the external consistency condition across the loop.

Notice that in every partition element $p \neq \mathcal{P}_1(\omega_1), \mathcal{P}_1(\omega_2)$, we have a uniform prior and non-informative signals, so the posteriors match the priors and the likelihood ratios for external consistency, as in Eq. 13, equal 1. This ratio holds for the partition element $\mathcal{P}_1(\omega_2)$

as well, because the prior is uniform, so

$$\frac{\frac{1}{2}\mu_{\tau,1}(\omega_2 \mid \mathcal{P}_1(\omega_2), s_1) + \frac{1}{2}\mu_{\tau,1}(\omega_2 \mid \mathcal{P}_1(\omega_2), s_2)}{\frac{1}{2}\mu_{\tau,1}(\overline{\omega}_2 \mid \mathcal{P}_1(\overline{\omega}_2), s_1) + \frac{1}{2}\mu_{\tau,1}(\overline{\omega}_2 \mid \mathcal{P}_1(\overline{\omega}_2), s_2)} \cdot \frac{\mu(\overline{\omega}_2)}{\mu(\omega_2)} = \frac{\frac{1}{3} + \frac{2}{3}}{\frac{2}{3} + \frac{1}{3}} \cdot 1 = 1.$$

Hence, the external consistency condition reduces to the likelihood ratio induced in $\mathcal{P}_1(\omega_1)$, however

$$\frac{\frac{1}{2}\mu_{\tau,1}(\omega_1 \mid \mathcal{P}_1(\omega_1), s_1) + \frac{1}{2}\mu_{\tau,1}(\omega_1 \mid \mathcal{P}_1(\omega_1), s_2)}{\frac{1}{2}\mu_{\tau,1}(\overline{\omega}_1 \mid \mathcal{P}_1(\overline{\omega}_1), s_1) + \frac{1}{2}\mu_{\tau,1}(\overline{\omega}_1 \mid \mathcal{P}_1(\overline{\omega}_1), s_2)} \cdot \frac{\mu(\overline{\omega}_1)}{\mu(\omega_1)} = \frac{\frac{1}{2} + \frac{1}{5}}{\frac{1}{2} + \frac{4}{5}} \cdot \frac{2\mu(\omega_1)}{\mu(\omega_1)} = \frac{14}{13} \neq 1.$$

We thus conclude that external consistency fails and the average of the two posterior beliefs is not feasible, implying that the feasible set in not convex. □

A.5 The Proof of Proposition 2

Proof. Fix an informative loop $\big((\omega_i, \bar{\omega}_i)\big)_{i=1}^{m}$. Using the construction of the bipartite graph, as in the proof of Theorem 2, there exists an equivalent cycle, which contains a simple cycle of even length, at least 4. Thus, w.l.o.g., we can assume that $F(\omega_i) \neq F(\bar{\omega}_i)$ for every i and the loop intersect any CKC at most once. For every $1 \leq i \leq m$, denote $f_i = F(\bar{\omega}_i)$ and $p_i = \mathcal{P}_1(\omega_i)$, so that all f_is and p_is are distinct.

Fix two public signals s_1, s_2 and define an F-measurable kernel τ by $\tau(s_1 \mid f_1) = \frac{1}{3}$, $\tau(s_1 \mid f) = \frac{2}{3}$ for every $f \in F \setminus \{f_1\}$, and $\tau(s_2 \mid f) = 1 - \tau(s_1 \mid f)$ for every $f \in F$. Since μ has full support, both s_1 and s_2 occur with positive probability at every state. Let $\mathrm{JB}^{(1)}$ and $\mathrm{JB}^{(2)}$ denote the joint beliefs induced by conditioning on s_1 and s_2, respectively. Both are feasible by construction and have matching support (given any state), so their average $\mathrm{JB}^{(\text{avg})} := \frac{1}{2}\,\mathrm{JB}^{(1)} + \frac{1}{2}\,\mathrm{JB}^{(2)}$ is a well-defined joint belief.

To maintain convexity, [**EXC**] must hold given $\varphi_{\mathrm{JB}^{(\text{avg})}}$. Because s_1 and s_2 are completely uninformative in every $p_i \neq p_1, p_2$ along the loop, $\varphi_{\mathrm{JB}^{(\text{avg})}}(\omega_i, \bar{\omega}_i) = 1$ for every $i = 3, 4, \ldots, m$. Thus, [**EXC**] reduces to $\prod_{i=1}^{2} \varphi_{\mathrm{JB}^{(\text{avg})}}(\omega_i, \bar{\omega}_i) = 1$.

For any CKC $p \in \{p_1, p_2\}$ define the (prior) mass of the mediator partition element f_1 inside p by $x_p := \frac{\mu(p \cap f_1)}{\mu(p)} \in (0,1)$, and let $D_p^{(k)} = \sum_{f \in F} \tau(s_k \mid f)\mu(p \cap f)$ be the (unnormalized) probability that signal s_k occurs given p. We thus get

$$D_p^{(1)} = \frac{1}{3}\,\mu(p \cap f_1) + \frac{2}{3}\big(\mu(p) - \mu(p \cap f_1)\big) = \frac{\mu(p)}{3}(2 - x_p),$$

$$D_p^{(2)} = \frac{2}{3}\,\mu(p \cap f_1) + \frac{1}{3}\big(\mu(p) - \mu(p \cap f_1)\big) = \frac{\mu(p)}{3}(1 + x_p).$$

Because $\mathrm{JB}^{(\mathrm{avg})}$ is an average of two joint beliefs, for any $\omega \in p \cap f$, we get

$$\mathrm{JB}^{(\mathrm{avg})}(\omega, 1)(\omega) = \mu(\omega)\left(\frac{1}{2} \cdot \frac{\tau(s_1 \mid f)}{D_p^{(1)}} + \frac{1}{2} \cdot \frac{\tau(s_2 \mid f)}{D_p^{(2)}}\right).$$

We can now substitute these values in the simplified **[EXC]** condition above, and get

$$\begin{aligned}
1 &= \frac{\mu(\omega_1) \cdot \left(\frac{1}{2} \cdot \frac{2/3}{D_{p_1}^{(1)}} + \frac{1}{2} \cdot \frac{1/3}{D_{p_1}^{(2)}}\right)}{\mu(\bar{\omega}_1) \cdot \left(\frac{1}{2} \cdot \frac{1/3}{D_{p_1}^{(1)}} + \frac{1}{2} \cdot \frac{2/3}{D_{p_1}^{(2)}}\right)} \cdot \frac{\mu(\bar{\omega}_1)}{\mu(\omega_1)} \cdot \frac{\mu(\omega_2) \cdot \left(\frac{1}{2} \cdot \frac{1/3}{D_{p_2}^{(1)}} + \frac{1}{2} \cdot \frac{2/3}{D_{p_2}^{(2)}}\right)}{\mu(\bar{\omega}_2) \cdot \left(\frac{1}{2} \cdot \frac{2/3}{D_{p_2}^{(1)}} + \frac{1}{2} \cdot \frac{1/3}{D_{p_2}^{(2)}}\right)} \cdot \frac{\mu(\bar{\omega}_2)}{\mu(\omega_2)} \\
&= \frac{\left(\frac{2}{D_{p_1}^{(1)}} + \frac{1}{D_{p_1}^{(2)}}\right)}{\left(\frac{1}{D_{p_1}^{(1)}} + \frac{2}{D_{p_1}^{(2)}}\right)} \cdot \frac{\left(\frac{1}{D_{p_2}^{(1)}} + \frac{2}{D_{p_2}^{(2)}}\right)}{\left(\frac{2}{D_{p_2}^{(1)}} + \frac{1}{D_{p_2}^{(2)}}\right)} = \frac{\left(\frac{2}{2-x_{p_1}} + \frac{1}{1+x_{p_1}}\right)}{\left(\frac{1}{2-x_{p_1}} + \frac{2}{1+x_{p_1}}\right)} \cdot \frac{\left(\frac{1}{2-x_{p_2}} + \frac{2}{1+x_{p_2}}\right)}{\left(\frac{2}{2-x_{p_2}} + \frac{1}{1+x_{p_2}}\right)} \\
&= \frac{4+x_{p_1}}{5-x_{p_1}} \cdot \frac{5-x_{p_2}}{4+x_{p_2}},
\end{aligned}$$

where the third equality follows from substituting the different values of $D_p^{(i)}$.

Define the function $R(x) = \frac{4+x}{5-x}$ and note it is strictly increasing on $(0, 1)$. Thus, the last equality is equivalent to

$$R(x_{p_1}) = R(x_{p_2}) \;\Leftrightarrow\; x_{p_1} = x_{p_2} \;\Leftrightarrow \mu(p_1 \cap f_1)\mu(p_2) - \mu(p_2 \cap f_1)\mu(p_1) = 0. \tag{20}$$

Let $\mathcal{N} \subset \Delta(\Omega)$ be the set of priors satisfying Equation (20). The left-hand side of Equation (20) is a nonzero polynomial in the coordinates of μ, hence its zero set $\mathcal{N}$ has Lebesgue measure zero in the interior simplex. Therefore, for every full-support prior, excluding a set of Lebesgue measure zero, the feasible set is not convex. □

A.6 The proof of Theorem 3

Proof. **Part (i)** We start by proving sufficiency. Suppose there exists an F-measurable signaling function τ as described in the theorem. The induced posteriors thus form a martingale. That is, if the posterior μ_i is realized with probability q_i, where $\sum_i q_i = 1$, then

$$\sum_i q_i \mu_i = \mu.$$

Let u be an option. So,

$$\mathbb{E}_{\mu}(u) = \sum_{i} q_i \mathbb{E}_{\mu_i}(u).$$

If each $\mathbb{E}_{\mu_i}(u) \geq 0$, then $\mathbb{E}_{\mu}(u) \geq 0$ as well. Therefore, the family $\{\mu_1, \ldots, \mu_m\}$ PP w.r.t. μ.

Moving on to prove necessity, suppose that no F-measurable signaling function τ exists such that all the posteriors it generates are in $\{\mu_1, \ldots, \mu_n\}$. Note however, that every μ_i can be generated through a specific F-measurable strategy and appropriate signal, and also note that convex combinations of F-measurable kernels are F-measurable as well. Thus, by standard results on Blackwell experiments, if $\mu \in \text{conv}\{\nu_1, \ldots, \nu_k\}$, namely if μ can be expressed as a convex combination of the ν_i's, then there is an experiment that produces signals whose posteriors are precisely $\{\nu_1, \ldots, \nu_m\}$.

Our assumption implies that $\mu \notin \text{conv}\{\mu_1, \ldots, \mu_m\}$. Since this convex hull is closed, the separating hyperplane theorem guarantees the existence of a nonzero vector $u \in \mathbb{R}^{\Omega}$ and a constant c, such that

$$\mathbb{E}_{\mu_i}(u) = \langle \mu_i, u \rangle > c \quad \text{for all } i,$$

while

$$\mathbb{E}_{\mu}(u) = \langle \mu, u \rangle < c.$$

By subtracting c in the two inequalities above and replacing u with $u - (c, \ldots, c)$, noting that the separation theorem is applied here to probability distributions, we conclude that the family $\{\mu_1, \ldots, \mu_m\}$ does not PP w.r.t. μ. This is a contradiction.

Part (ii) We prove sufficiency first. Suppose there exists an F-measurable signaling function τ as described in Part (ii) of the theorem. Let u be an option such that $\mathbb{E}_{\mu}(u) = 0$ and $\mathbb{E}_{\mu_i}(u) \geq 0$ for every i. We show that $\mathbb{E}_{\mu_i}(u) = 0$ for every i. As before, if the posterior μ_i is realized with probability $q_i > 0$, where $\sum_i q_i = 1$, then

$$\sum_{i} q_i \mu_i = \mu.$$

Let u be an option. So,

$$\mathbb{E}_{\mu}(u) = \sum_{i=1}^{m} q_i \mathbb{E}_{\mu_i}(u).$$

If each $\mathbb{E}_{\mu_i}(u) \geq 0$ and $\mathbb{E}_{\mu}(u) = 0$, then $\mathbb{E}_{\mu_i}(u) = 0$ for every i . Therefore, the family $\{\mu_1, \ldots, \mu_m\}$ SPP w.r.t. μ.

To establish necessity and similarly to the proof of Part (i), we show that μ can be expressed as a convex combination of $\mu_1, \ldots, \mu_m$, each with a positive weight. For this purpose, we prove the two following lemmas.

Lemma 2. SPP *implies that for every j there is a convex combination $\mu = \sum_i q_i \mu_i$, where $q_j > 0$.*

Proof. Fix an index j. Suppose, to the contrary, that there is no convex combination $\mu = \sum_i q_i \mu_i$ with $q_j > 0$. This implies that the set

$$D := \Big\{ q_j \mu_j + \sum_{i \neq j} q_i \mu_i \; ; \; q_j > 0, \; q_i \geq 0 \text{ for } i \neq j, \text{ and } q_j + \sum_{i \neq j} q_i = 1 \Big\}$$

does not contain μ. Note that D is convex and has a nonempty interior relative to the simplex of all distributions over Ω; in particular, the point $\frac{1}{n} \sum_i \mu_i$ is an interior point of D.

Consider now the closure of D, denoted $\bar{D}$. This set is also convex and contains all the μ_i's (including μ_j). The point μ may lie outside $\bar{D}$ or on its boundary. In either case, there exists a nonzero hyperplane u (referred to here as an option) such that $\mathbb{E}_\mu(u) = 0$ and $\mathbb{E}_\nu(u) \geq 0$ for every $\nu \in \bar{D}$.[22] In particular, this implies that $\mathbb{E}_{\mu_i}(u) \geq 0$ for every i.

We now apply SPP to the option u, which yields $\mathbb{E}_{\mu_i}[u] = 0$ for all i. Consider, however, the interior point $\nu = \frac{1}{m} \sum_i \mu_i$, which belongs to $\bar{D}$. Being an interior point, we have on the one hand $\mathbb{E}_\nu[u] > 0$ (see Theorem 11.3 in Rockafellar, 1970), while on the other hand,

$$\mathbb{E}_\nu[u] = \mathbb{E}_{\frac{1}{m} \sum_i \mu_i}[u] = \frac{1}{m} \sum_i \mathbb{E}_{\mu_i}[u] = 0.$$

This is a contradiction. Hence, there must exist a convex combination $\mu = \sum_i q_i \mu_i$ with $q_j > 0$, as required. □

Given Lemma 2, we now turn to the next lemma, which shows that μ can be expressed as a convex combination of $\mu_1, \ldots, \mu_m$, each with a positive weight.

Lemma 3. *If for every j there is a convex combination of $\mu_1, \ldots, \mu_m$ such that $\mu = \sum_i q_i \mu_i$, where $q_j > 0$, then there exists a combination where all q_i are strictly positive.*

Proof. Fix j and suppose that $\mu = \sum_i q_i^j \mu_i$ is a convex combination, where $q_j^j > 0$. Take $\mu = \frac{1}{n} \sum_{j=1}^m \sum_i q_i^j \mu_i$, and note that in this combination, the weight of every μ_i is positive. □

Combining the two lemmas, we derive from SPP that μ is a convex combination (with strictly positive weights) of $\mu_1, \ldots, \mu_m$, as needed. □

[22]Note that every option u derived from the separating hyperplane theorem can be adjusted with a positive constant to ensure the stated conditions.