

Constrained Mediation: Bayesian Implementability of Joint Posteriors*

David Lagziel[†]

Ehud Lehrer[‡]

Ben-Gurion University

Durham University

October 23, 2025

Abstract

We examine information structures in settings with privately informed agents and an informationally constrained mediator who supplies additional public signals. Our focus is on characterizing the set of posteriors that the mediator can induce. To this end, we employ a graph-theoretic framework: states are represented as vertices, information sets correspond to edges, and a likelihood ratio function on edges encodes the posterior beliefs. Within this framework, we derive necessary and sufficient conditions, internal and external consistency, for the rationalization of posteriors. Finally, we identify conditions under which a single mediator can implement multiple posteriors, effectively serving as a generator of Blackwell experiments.

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

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

*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-Sheva 8410501, Israel. E-mail: David-lag@bgu.ac.il.

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

1 Introduction

What is the role of an objective mediator in reaching an agreement, such as a peace treaty, between two opposing sides? The answer seems straightforward: to persuade both parties to sign the deal. Yet this form of persuasion is particularly intriguing, 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.

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. A fundamental question, therefore, is whether a given collection of posterior beliefs, interpreted as arising from Bayesian updating, can actually be implemented by some signal structure consistent with an external agent’s knowledge. In this paper, we address this question by analyzing the conditions under which posterior beliefs are *implementable* via a signal generated by an external-information provider, a mediator, endowed with partial knowledge of the state.

Our framework assumes that the mediator’s information is described by a partition over a finite state space, and that it can emit public signals governed by stochastic rules measurable with respect to its own information. Players, upon observing a signal, update their beliefs via Bayes’ rule. We ask: given a profile of posterior beliefs or likelihoods across possible states, can we find a mediator-compatible signal that induces these beliefs?

This question is central in economics and game theory, especially in the design and comparison of information structures, and in understanding how much control an information provider has over the beliefs and behavior of rational agents. While it is always 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.

To address this question, we introduce a compact representation of Bayesian updating through a *posterior likelihood function*, which specifies the ratio of posteriors across *adjacent states*, i.e., states that lie in the same information set of some player. We then ask under what conditions such a function can be rationalized by a mediator-generated signal. 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 positive function φ that captures comparative likelihoods between adjacent states. 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. Thus, our central question translates to: under what conditions can such a function be rationalized as the output of a signal generated by a mediator, constrained by the limits of its own knowledge?

To address this, we develop conditions under which φ admits a representation in terms of an F -measurable function (where F denotes the mediator’s partition), allowing φ to be expressed as the ratio of values of a positive function defined over states. These conditions, named *internal and external consistency*, 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.

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 recently 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. The realized payoffs depend on the state: $x = 2$ if $\omega \in \{\omega_1, \omega_2\}$, and $x = 3$ if $\omega \in \{\omega_3, \omega_4\}$. 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	C
Player 1	A	$(2, -5)$	$(x, -4)$
	C	$(-1, -1)$	$(0, 0)$

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

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

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

Figure 1: Two payoff matrices with Player 1 (row player) and Player 2 (column player). If ω_1 or ω_2 are realized, then $x = 2$, and if ω_3 or ω_4 are realized, then $x = 3$.


The interpretation of the game is straightforward. Each player can either attack the other or compromise by signing 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 $(\frac{1}{2}, \frac{1}{2}, 0, 0)$. 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 $(0, 0, \frac{1}{2}, \frac{1}{2})$, 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 $F = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. The mediator's task is to design a signaling mechanism that may persuade 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 .


To illustrate the mediator's role, Figure 2 depicts the players' and the mediator's information structures. If the mediator fully discloses this information, both players become perfectly informed about the state, and the peaceful outcome (C, C) is never sustained in equilibrium. By contrast, if the mediator could hypothetically induce a posterior of $(\frac{1}{3}, \frac{2}{3}, 0, 0)$ for player 1

¹A 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 (green).

while player 2 remains fully informed, then player 1 is indifferent between A and C since action A yields an expected payoff of $\frac{1}{3} \cdot 2 + \frac{2}{3} \cdot (-5) = -\frac{8}{3}$, while 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 a posterior of $(0, 0, \frac{3}{4}, \frac{1}{4})$ for player 2 while player 1 remains fully informed, then player 2 is indifferent between A and C , because action A yields an expected payoff of $\frac{3}{4} \cdot (-5) + \frac{1}{4} \cdot 3 = -3$, while action C yields $\frac{3}{4} \cdot (-4) + \frac{1}{4} \cdot 0 = -3$. Note that the previous proportions of $\frac{1}{3}$ and $\frac{2}{3}$ change to $\frac{1}{4}$ and $\frac{3}{4}$ because x varies between states. 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 $(\frac{1}{3}, \frac{2}{3}, 0, 0)$ for player 1 when the realized state lies in $\{\omega_1, \omega_2\}$, while at the same time generating a posterior $(0, 0, \frac{3}{4}, \frac{1}{4})$ 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.

So why does the mediator, in this specific example, fails to generate a joint posterior of the stated form? The obstacle arises from measurability (information) constraints. Any signaling function based on the partition F must send the same signal in both ω_1 and ω_3 , as well as 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 when a mediator

can 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 and loops.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, the tail of the loop, ω_3 , belongs to the same mediator information set as the head, ω_1 . Thus, the sequence is a chain of states alternating between belonging to a player's CKC and to a mediator information set.

The sequence $(\omega_1, \omega_2, \omega_4, \omega_3)$ is an example of a *loop*, formally defined in this paper (see Section 4.2) and also employed in Lagziel et al. (2025a,b), in which the mediator's knowledge 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 $(\frac{1}{3}, \frac{2}{3}, 0, 0)$ 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 $(0, 0, \frac{3}{4}, \frac{1}{4})$ the likelihood ratio of ω_4 to ω_3 is $\varphi(\omega_4, \omega_3) = \frac{1/4}{3/4} = \frac{1}{3}$. 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*, 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 is a more advanced problem: the preliminary step is to identify which posteriors are feasible. In other words, before addressing the question of optimality, one must

²Notably, 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.

³Formally, 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.

first determine the feasible actions, and this is the focus of the current research.

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, 1974) 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 additional 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 (1953), 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 Kellner (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 recently 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.⁴

Independently of Arieli et al. (2021), the study of Ziegler (2020) deals with 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 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.

Finally, our results contribute to the broader literature on higher-order beliefs (see the review

⁴See 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.

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 introduces the model: a finite state space, players' information partitions $\mathcal{P}_i$, a common prior μ , and an external mediator whose information is represented by a partition F . Section 3 discusses joint beliefs and joint posteriors. It formulates the central research question: given a joint belief, does there exist an F -measurable signal that generates it as a joint posterior?

Section 4 presents the graph of information, where nodes correspond to states and edges reflect indistinguishability under some player's partition. The section defines the posterior likelihood function on edges and uses it to reformulate the implementability question. Section 4.2 introduces two key conditions, *internal consistency* and *external consistency*, which play a central role in resolving the implementability problem. It then establishes a technical graph-theoretic result (Theorem 1) that serves as the backbone of our main theorem.

Section 5 reformulates the problem in terms of joint posteriors by defining a posterior likelihood function that incorporates both the joint posterior and the prior. Theorem 2, the main result of the paper, characterizes when a joint posterior is implementable by a mediator in terms of the two consistency conditions applied to this function.

Section 6 generalizes the framework to accommodate for multiple signals. It introduces the notion of *positivity preservation* and shows that a family of posterior likelihood functions can be simultaneously implemented by a mediator if and only if the family preserves positivity with respect to the prior.

Finally, Section 7 offers additional interpretations and extensions. Section 7.1 extends the analysis to settings where different players or groups of players observe different signals. It shows that the consistency conditions apply not only globally but also within any subgroup. Section 7.2 reformulates the multiplicative consistency conditions in logarithmic terms and relates them to potential games (see Monderer and Shapley, 1996) in Corollary 2.

2 The Model

Let $N = \{1, 2, \dots, n\}$ with $n \geq 2$, 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, 1974). 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, he uses a public signaling function $\tau : \Omega \rightarrow \Delta(S)$, where S is a finite set of signals and $\Delta(S)$ is a distribution over S . Let $\tau(s | \omega)$ 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 .⁵ Formally, τ is an F -measurable stochastic kernel:

- (i) For each state $\omega \in \Omega$, $\tau(\cdot | \omega)$ is a probability distribution over a finite set of signals S .
- (ii) For all $\omega' \in F(\omega)$ and $s \in S$, we have $\tau(s | \omega') = \tau(s | \omega)$.

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.

3 Bayesian Updating

Recall that μ is the common prior and assigns positive probability to every state. Otherwise, we may restrict attention to the support of μ . Let τ denote the mediator's signaling function.

3.1 Joint posterior beliefs

When the realized state is ω , player i is informed of $\mathcal{P}_i(\omega)$. Assuming that $\tau(s | \mathcal{P}_i(\omega)) > 0$ and upon observing s , player i updates their belief. His posterior probability of $\omega' \in \mathcal{P}_i(\omega)$, given $\mathcal{P}_i(\omega)$ and s , is

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

Thus, $\mu_{\tau,i}(\cdot | \mathcal{P}_i(\omega), s)$ is a probability distribution, conditioned on both $\mathcal{P}_i(\omega)$ and s , over Ω whose support is a subset of $\mathcal{P}_i(\omega)$. For notational convenience, we define $\mu_{\tau,i}(\omega' | \mathcal{P}_i(\omega), s) = 0$ for every ω' , whenever $\tau(s | \mathcal{P}_i(\omega)) = 0$.

⁵Measurability 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 $\tau(s | \omega') = \tau(s | \omega)$.

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 that can generate s :

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

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.

3.2 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 generate.

Definition 1. A joint belief is a stochastic map $\text{JB} : \Omega \times N \rightarrow \Delta(\Omega) \cup \{\mathbf{1}_\emptyset\}$, where $\mathbf{1}_\emptyset$ is the zero vector, such that the following conditions hold: (i) if $\text{JB}(\omega, i)(\omega) > 0$ for some i , then $\text{JB}(\omega, j)(\omega) > 0$ for every j ; and (ii) $\frac{\text{JB}(\omega, i)(\omega)}{\text{JB}(\omega', i)(\omega')} = \frac{\text{JB}(\omega, j)(\omega)}{\text{JB}(\omega', j)(\omega')}$ for every i, j, ω and $\omega' \in \mathcal{P}_i(\omega) \cap \mathcal{P}_j(\omega)$, assuming the denominator is nonzero.

When well-defined, $\text{JB}(\omega, i)$ 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 $\text{JB}(\omega, i)(\omega')$.

A few remarks are in order. First, the two conditions in Definition 1 are necessary for JB to be induced as a posterior by a public signal. Indeed, they follow from the basic properties of Bayesian updating given the existence of a common prior and a mediator's *public* signal.⁶ Second, a general joint posterior specifies a profile of posterior beliefs for every feasible state. In cases without a common prior, for instance, condition (ii) in the definition need not be satisfied. Third, we include the indicator $\mathbf{1}_\emptyset$ in the definition to capture situations in which the mediator provides additional information that assigns zero probability to some state in $\mathcal{P}_i(\omega)$. Specifically, prior to receiving the mediator's signal, each player is subjectively informed about an information set $\mathcal{P}_i(\omega)$. The mediator may then send a signal that has zero likelihood in certain states within this set. In such cases, the conditional posterior is not well-defined, and we therefore set $\text{JB}(\omega, i) = \mathbf{1}_\emptyset$ for all such states.

Given a JB and a common prior μ , a natural question is whether there exists a signaling function τ of the mediator 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,

⁶A broader discussion of this point is provided in connection with Eq. (3) below.

although posterior updating is typically player-dependent (as the denominator reflects player i 's information partition), the following likelihood ratios are not:

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

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.

3.3 Posterior likelihood and the leading question

To formally define the problem, fix a JB and a common prior μ , and let Ω_+ denote the set of states such that $\text{JB}(\omega, i)(\omega) > 0$ for some player i , and hence for all players. In particular, we restrict attention to the states that are assigned positive probability. Normalizing μ to this set, we obtain the distribution $\mu(\cdot \mid \Omega_+)$.⁷

For a pair $\omega, \omega' \in \Omega_+$ such that $\omega' \in \mathcal{P}_i(\omega)$ for some i , consider the ratio $\frac{\text{JB}(\omega, i)(\omega)}{\text{JB}(\omega', i)(\omega')}$. 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

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

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.

⁷We discuss the case where some states are assigned zero probability in Section 5.




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.

	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, $\frac{1/2}{1/3}$ for player 1 and $\frac{1/3}{1/2}$ 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 stochastic signal, say s , sent by the mediator according to the following conditional probabilities: $\mathbb{P}(s \mid \omega_i) = 1/5$ for $i = 1, 4$, and $\mathbb{P}(s \mid \omega_i) = 2/5$ for $i = 2, 3, 5$. Note that the conditional probabilities given ω_1 and ω_4 are equal, and the same applies to ω_2 , ω_3 , and ω_5 . Therefore, the signal s is F -measurable.

Suppose now that the common prior is uniform, and that in the Table 2, the posteriors of player 1 in ω_4 and ω_5 were $(0, 0, 0, \frac{1}{5}, \frac{4}{5})$ rather than those originally stated. The table then forms a joint belief (see Definition 1). Is there still a signal that the mediator could make public that would induce this joint belief? The answer is **no**.

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

In a strategic setting, the entire profile of beliefs, as described in the tables above, plays a fundamental role. In equilibrium, players choose actions based not only on their own beliefs, but also on their beliefs about the beliefs of others, and further on higher-order beliefs, i.e., what players believe about others' beliefs about their beliefs, and so forth. This recursive structure, commonly referred to as a *hierarchy of beliefs*, is fully captured by the specification of each player's belief at every possible state. The richness of this hierarchy ensures that players' strategic reasoning accurately reflects the underlying information structure of the game.

Consider, for instance, Table 1 and the belief of player 1 when state ω_4 is realized. At this state, player 1 assigns equal probabilities to ω_4 and ω_5 , and thus assigns equal probabilities to the events that player 2 knows the true state is ω_4 or that it is ω_5 . In other words, player 1 is uncertain not only about the true state, but also about player 2's information, illustrating the higher-order beliefs embedded in the structure.

The importance of generating appropriate priors in each environment lies in the need to initialize the entire hierarchy of beliefs correctly. Priors determine how beliefs virtually propagate through the system, shape expectations, and ultimately influence equilibrium behavior. Without a well-specified prior structure, the analysis of strategic interaction under incomplete information remains fundamentally incomplete. This consideration forms the foundation of the present study.

This example also highlights the central role of likelihood ratios in characterizing which posterior beliefs can be implemented by a given mediator. Specifically, likelihood ratios capture the relative plausibility of states within the same information set from the perspective of the players. Any feasible public signal must preserve these ratios, in accordance with the underlying measurability constraints imposed by the information structure. As a result, analyzing likelihood ratios offers crucial insight into both the design and the limitations of the signals that a mediator can publicly disclose.

4 The Graph of Information and Consistency Conditions

4.1 The Graph and the Posterior Likelihood Function

In this section, we introduce the *graph of information*, which will serve as the key tool for addressing the question of whether a given joint posterior can be implemented by the mediator.

We follow Rodrigues-Neto (2009) and define a graph $G = (V, E)$, where the set of vertices V coincides with the state space Ω ; that is, each vertex represents a state. The set of directed edges E is defined as follows: for any pair of states $\omega, \omega' \in \Omega$, we have $(\omega, \omega') \in E$, and denoted $(\omega \rightarrow \omega')$ if there exists a player $i \in N$ such that $\omega' \in \mathcal{P}_i(\omega)$. Note that whenever $(\omega, \omega') \in E$, it also holds that $(\omega', \omega) \in E$. We explicitly include both directions because it will be essential in what follows.

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:

Using the graph of information, we can define a *posterior likelihood* (PL) function, which will be used to translate a joint posterior to likelihood ratios.

Definition 2. A positive function φ defined over E is a *posterior likelihood function* if it satisfies the condition⁸

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

⁸We abuse notation and use $\varphi(\omega, \omega')$ instead of $\varphi((\omega, \omega'))$.




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

Example 1, continued. Recall Table 2. Note that the likelihood ratio of the probabilities of ω_1 and ω_2 is $1/2$. 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:

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

while preserving the relation in Eq. (5).

4.2 The mediator-induced PL function: Internal and External consistency

The first question we consider concerns the information structure $(\mathcal{P}_i)_i$ and the mediator's partition F through the notion of a PL function. Specifically, fix a PL function φ : does there exist an F -measurable, strictly positive function $f : \Omega \rightarrow \mathbb{R}_{++}$ that yields

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

for every edge $(\omega, \omega') \in E$?

Given the ability and tools to answer this question, we can translate any JB to its respective PL function (as in Theorem 2 below) and check whether the mediator can induce a function f that replicates the ratios given in the last equation, as well as in Eq. (4) above. If the mediator can indeed generate such a function, then it will correspond to the probabilities of a feasible signal that generates the needed JB (see Remark 1 below).

To address this question, we introduce the notions of an F -cycle and a *loop*, and build upon them two conditions: internal and external consistency. As we prove in Theorem 1 below, internal and external consistency are necessary and sufficient conditions to answer our question.

Formally, an F -cycle is a sequence of edges $((\omega_i, \omega_{i+1}))_{i=1}^n$ such that $(\omega_i, \omega_{i+1}) \in E$ for all $i = 1, \dots, n$, and $\omega_{n+1} \in F(\omega_1)$. The next property, referred to as *Internal consistency*, parallels

Definition 2 in Rodrigues-Neto (2009), where cycles are defined solely in terms of the players' information, independently of the mediator.⁹

Definition 3. Let $(\omega_i, \omega_{i+1})_{i=1}^n$ be an F -cycle. Then, Internal consistency [INC] holds if

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

Hellman and Samet (2012) refer to the left-hand side of Eq. (7), in the case where 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, due to the condition given in Eq. (5), a positive function φ defined over E is a PL if and only if Eq. (7) is satisfied for every such F -cycle. In particular, a function φ that satisfies [INC] is a PL.

We can now use [INC] to extend every φ over any connected pair of states within a CKC. To do so, denote by $\twoheadrightarrow$ the transitive closure of $\rightarrow$. That is, $\omega \twoheadrightarrow \omega'$ if there is a sequence of edges $((\omega_i, \omega_{i+1}))_{i=1}^n$ such that $\omega = \omega_1$ and $\omega_{n+1} = \omega'$. It is well known that a set of states connected, w.r.t. $\twoheadrightarrow$, are in the same CKC (see Rodrigues-Neto, 2009). The following Lemma 1 uses [INC] to extend φ in a consistent manner. (The proof is relegated to the appendix; see Section A.1.)

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

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

In Example 1, the sequence $((\omega_1, \omega_2), (\omega_2, \omega_3), (\omega_3, \omega_1))$ 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 shows that if φ satisfies [INC], then it can also induce a probability distribution, denoted by $\mu_\varphi(\cdot \mid C)$, over any connected component C of the graph G . Specifically, for a connected component C ,

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

where ω is an arbitrary state in C .

⁹In 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.

Next, we use the extended φ to define a loop and the external consistency property. Suppose that φ is already defined for every pair connected by $\rightarrow$. An *F-loop* is a sequence of pairs $((\omega_i, \bar{\omega}_i))_{i=1}^n$ such that, for each $i = 1, \dots, n$, the following conditions hold:

- (i) $\omega_i \rightarrow \bar{\omega}_i$,
- (ii) $\bar{\omega}_i \not\rightarrow \omega_{i+1}$,
- (iii) $\bar{\omega}_i \in F(\omega_{i+1})$, with the convention that $\omega_{n+1} := \omega_1$.

Condition (i) requires that each pair of states $(\omega_i, \bar{\omega}_i)$ belongs to the same connected component, that is, they are transitively connected and lie within the same CKC. In contrast, condition (ii) stipulates that successive pairs are separated across CKCs: the states $\bar{\omega}_i$ and ω_{i+1} are not transitively connected and thus belong to different connected components. Nevertheless, as specified in condition (iii), every two such states are indistinguishable to the mediator: they lie within the same cell of the mediator's partition. This represents a situation in which the mediator is less informed than the players: it cannot distinguish between states that all players can. 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 concept of a loop plays a pivotal role in Lagziel et al. (2025a,b) to provide conditions such that one mediator dominates another, in an extension of Blackwell's work on the comparison of experiments (see Blackwell, 1951, 1953).

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

Definition 4. Let $((\omega_i, \bar{\omega}_i))_{i=1}^n$ be an *F-loop*. Then, External consistency [EXC] holds if

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

The [EXC] property has substantive content only in situations where 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 of this is an all-knowing mediator, who knows the exact identity of the realized state, i.e., $F(\omega) = \{\omega\}$ for every $\omega \in \Omega$. In this case, [INC] reduces to simple cycles, as in Rodrigues-Neto (2009) and Hellman and Samet (2012).

Before presenting our first main result and characterization, 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.

Our first main result, stated in Theorem 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.)

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

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

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

Once Theorem 1 establishes a characterization of a given PL function, we can employ it in Section 5 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 any connected components C . The next section is devoted to implementation. In particular, it shows, based on Theorem 1, that when φ satisfies also **[EXC]**, a mediator may generate by a proper signal, say s , the combination

$$\mu_\varphi = \sum_C \mu(C \mid s) \mu_\varphi(\cdot \mid C), \quad (12)$$

where the summation is 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, μ_φ is the posterior distribution over Ω generated after observing the message s , sent by the mediator.

5 Application to Bayesian Updating

In this section, we apply Theorem 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 posterior be induced by a signaling function τ of the mediator and a signal s ?

We employ Theorem 1 to answer this question. To see this, let us focus first on the subset $\Omega_+ \subseteq \Omega$ of states ω for which¹⁰ $\text{JB}(\omega, i)(\omega) > 0$ 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:

$$\varphi_{\text{JB}}(\omega, \omega') = \frac{\text{JB}(\omega, i)(\omega)}{\text{JB}(\omega', i)(\omega')} \cdot \frac{\mu(\omega')}{\mu(\omega)}, \quad \forall (\omega, \omega') \in E_+. \quad (13)$$

Since both μ and $\text{JB}(\omega, i)(\omega)$ are strictly positive on Ω_+ , φ_{JB} is strictly positive as well, and 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 Theorem 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 $\boldsymbol{\mu}_{\tau, s} = \text{JB}$. (The proof is relegated to the appendix; see Section A.3.)

Theorem 2. *Fix a joint belief JB. Then there exists an F -measurable signaling function τ and a signal s , such that $\boldsymbol{\mu}_{\tau, s} = \text{JB}$ if and only if*

- (i) Ω_+ is measurable w.r.t. F ; and
- (ii) φ_{JB} satisfies conditions [INC] and [EXC] in G_+ .¹¹

The economic implication of Theorem 2 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 2 identifies conditions under which there exists a signaling mechanism, implemented by the mediator, that induces these posterior beliefs, as described in Eq. (1).

¹⁰Recall that $\text{JB}(\omega, i)$ is a distribution while $\text{JB}(\omega, i)(\omega)$ is a probability.

¹¹[EXC] is satisfied using the extension provided by Lemma 1, ensured by [INC].

6 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, $\text{JB}_1, \dots, \text{JB}_n$. 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 2. This means that for each i , there exist 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

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

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 $s_1, \dots, s_n$ and

$$\{\mu_\tau(\cdot \mid s_1), \dots, \mu_\tau(\cdot \mid s_n)\} = \{\mu_1, \dots, \mu_n\}? \quad (15)$$

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 $u : \Omega \rightarrow \mathbb{R}$ is called an *option* (or a *state-contingent claim*), as it specifies a monetary payoff for each possible state.

Definition 5. (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 $\mu_1, \dots, \mu_n$ be the posteriors induced by observed signals. If the posteriors $\mu_1, \dots, \mu_n$ *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 $\varphi_{\text{JB}_1}, \dots, \varphi_{\text{JB}_n}$ (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 $\{\mu_1, \dots, \mu_n\}$. The second part of the theorem uses SPP to characterize when there is a signaling function whose set of posterior coincides with $\{\mu_1, \dots, \mu_n\}$. (The proof is given in Section A.4 in the appendix.)

Theorem 3. *Let μ be a common prior, and let $\varphi_{\text{JB}_1}, \dots, \varphi_{\text{JB}_n}$ be PL functions that satisfy [INC] and [EXC]. For every JB_i , assume that Ω_+^i , defined as in in Theorem 2, 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 $\{\mu_1, \dots, \mu_n\}$ if and only if this family PP w.r.t. μ ;*
- (ii) *the set of posteriors it generates with positive probability coincides with $\{\mu_1, \dots, \mu_n\}$ 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 2. 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 JB.

7 Extensions, Comments and Discussion

7.1 From Public to Private Signals: Differential information to different groups

The preceding discussion has focused on the case of public signals, where all players observe the same realization from the information structure. In this section, we turn to a more general setting in which different groups of players receive different signals. Such a structure allows for asymmetric information across groups, introducing a richer set of strategic considerations.

Formally, let $I \subseteq N$ be a subset of players. A larger group of players generally induces a coarser partition of the state space into CKCs. That is, the set of CKCs for I is finer than the set of CKCs generated by the full set of players N . In other words, every CKC of the players in I is a subset of some CKC of the players in N .

This reduction in the group of players can alter both the internal and external properties that depend on the structure of CKCs. To illustrate, consider the model presented in Example 1, where the full set of players is $N = \{1, 2\}$. Now, take the subset $I = \{1\}$, consisting only of player 1. In this case, a CKC corresponds to an information set of player 1, and the collection of CKCs is simply the partition $\mathcal{P}_1$, player 1's information. That is, with only player 1 in the group, common knowledge coincides with individual knowledge.

Denote by G_N the graph induced by the entire group and by G_I its restriction that corresponds to the information structure induced by I . That is, the set of vertices of G_I is the set of states, and (ω, ω') is an edge if there is a player $i \in I$ such that $\omega' \in \mathcal{P}_i(\omega)$. In this case, and for a given joint posterior JB, the PL function $\varphi_{\text{JB}}(\omega, \omega')$ is the same in the two graphs.

Let I -[INC] and I -[EXC] denote the internal and external consistency conditions corresponding to group I in the subgraph G_I . Given a joint posterior profile for group I , one may ask, analogously to the earlier case, whether this profile can be generated by a public signal emitted by a mediator and received by all members of group I . Theorem 1 addresses this question by providing necessary and sufficient conditions for such a representation, expressed in terms of I -[INC] and I -[EXC].

It is evident that 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, Theorem 1 implies that if the corresponding PL function φ_{JB} satisfies N -[INC] and N -[EXC], it must also satisfy I -[INC] and I -[EXC]. In the following Proposition 1, we demonstrate that this implication can be established directly, without appealing to Theorem 1. (For the proof, see Section A.5 in the appendix.)

Proposition 1. *Fix a JB. If the PL function φ_{JB} satisfies conditions N -[INC] and N -[EXC], then it also satisfies I -[INC] and I -[EXC] for any subgroup $I \subseteq N$.*

7.2 An Additive Interpretation and Implications to Potential Games

The concepts, structure, and results derived from the graph of information extend well beyond Bayesian updating and apply to a variety of other frameworks, one of which is that of potential games, as described below. In particular, this section employs an additive version of the internal consistency property to derive a characterization of potential games.

To this end, consider a PL function φ and a function f such that $\varphi(\omega, \omega') = \frac{f(\omega)}{f(\omega')}$. Define $g := \log(f)$ and take the logarithmic transformation of the previous equality. The resulting quantity

$$\rho(\omega, \omega') := \log(\varphi(\omega, \omega')) = g(\omega) - g(\omega'),$$

can be interpreted as capturing the difference (or “gap”) between connected vertices, where g represents a form of “height” or potential assigned to each vertex.

Formally, let (V, E) denote a finite, directed, and connected graph, and let $\rho : E \rightarrow \mathbb{R}$ be a real-valued function defined on its edges. A discrete analogue of the [INC] constraint arises when F is discrete. Specifically, assume that F is discrete, meaning that every atom of F consists of a single state. In this setting, a *basic cycle* is a sequence of edges $((\omega_i, \omega_{i+1}))_{i=1}^n$ such that $(\omega_i, \omega_{i+1}) \in E$ for all $i = 1, \dots, n$, and $\omega_{n+1} = \omega_1$. An additive analogue of the [INC] condition, abstracting from any mediator or probabilistic interpretation, can then be stated as follows:

Definition 6. *Let $(\omega_i, \omega_{i+1})_{i=1}^n$ be a basic cycle where $\omega_{n+1} = \omega_1$. Then, Additive Internal Consistency [INC-ADD] holds if*

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

This condition reflects additive path-independence and echoes a multiplicative version, for

example, as in Rodrigues-Neto (2009) and Hellman and Samet (2012).

We now connect this idea to an immediate implication of Theorem 1. The following corollary states that **[INC-ADD]** is a necessary and sufficient condition for the existence of a function g , whose differences match the values of ρ .

Corollary 1. *Let (V, E) be a finite, directed, connected graph, and let $\rho : E \rightarrow \mathbb{R}$ be a real-valued function. Then there exists a function $g : V \rightarrow \mathbb{R}$ such that*

$$g(\omega) - g(\omega') = \rho(\omega, \omega') \text{ for every } (\omega, \omega') \in E,$$

*if and only if ρ satisfies **[INC-ADD]**.*

This observation naturally leads us to the notion of potential games. A *potential game* is a strategic game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$, where N is the set of players, with $|N| = n$; A_i is the finite set of actions available to player i ; and $u_i : A \rightarrow \mathbb{R}$ is the payoff function of player i , with $A := \times_{i \in N} A_i$ being the set of action profiles.

We say that Γ is a *potential game* (see, Monderer and Shapley, 1996) if there exists a function $g : A \rightarrow \mathbb{R}$ such that for every player $i \in N$, every action profile $a = (a_1, \dots, a_n) \in A$, and every action $a'_i \in A_i$, it holds that

$$g(a) - g(a_1, \dots, a'_i, \dots, a_n) = u_i(a) - u_i(a_1, \dots, a'_i, \dots, a_n).$$

To apply Corollary 1, define a graph associated with the game Γ , denoted $\text{Graph}(\Gamma)$. Its vertex set is A , and its edge set consists of all pairs of the form

$$((a_1, \dots, a_n), (a_1, \dots, a'_i, \dots, a_n))$$

for some $i \in N$ and $a'_i \in A_i$. We then define a function $\rho_\Gamma : E \rightarrow \mathbb{R}$ on this edge set by:

$$\rho_\Gamma((a_1, \dots, a_n), (a_1, \dots, a'_i, \dots, a_n)) := u_i(a) - u_i(a_1, \dots, a'_i, \dots, a_n).$$

Thus, an immediate consequence of Corollary 1 is the following characterization:

Corollary 2. *Let Γ be a strategic game. Then Γ is a potential game if and only if ρ_Γ satisfies **[INC-ADD]** over $\text{Graph}(\Gamma)$.*

Corollary 2 resembles Proposition 2.8 in Monderer and Shapley (1996) that provides a necessary and sufficient condition for a finite game to be an exact potential game in terms of the cross-differences of players' payoffs.

A related, though somewhat tangential, question concerns the class of *ordinal potential games*. We say that a game is an ordinal potential game if there exists a function $g : A \rightarrow \mathbb{R}$ such that for every player $i \in N$, every action profile $a = (a_1, \dots, a_n) \in A$, and every deviation $a'_i \in A_i$, it holds that

$$g(a) - g((a_1, \dots, a'_i, \dots, a_n)) \geq 0 \Leftrightarrow u_i(a) - u_i((a_1, \dots, a'_i, \dots, a_n)) \geq 0.$$

In other words, the direction of improvement in each player's utility is perfectly aligned with that of the potential function. An intriguing question, which remains beyond the current analysis, is what characterizes ordinal potential games in terms of underlying structural or consistency properties. We leave this question for future exploration.

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. (1974): "Subjectivity and correlation in randomized strategies," *Journal of Mathematical Economics*, 1, 67–96.
- 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.
- (1953): "Equivalent comparisons of experiments," *The Annals of Mathematical Statistics*, 24, 265–272.
- 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.
- 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.
- FEINBERG, Y. (2000): "Characterizing Common Priors in the Form of Posteriors," *Journal of Economic Theory*, 91, 127–179.

- GEANAKOPOLOS, J. (1994): “Chapter 40 Common knowledge,” in *Handbook of Game Theory with Economic Applications*, Elsevier, vol. 2, 1437–1496.
- 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.
- 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.].
- 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.
- LAGZIEL, D., E. LEHRER, AND T. WANG (2025a): “Comparison of Oracles: Part I,” .
- (2025b): “Comparison of Oracles: Part II,” .
- LEHRER, E. AND D. SAMET (2014): “Belief Consistency and Trade Consistency,” *Games and Economic Behavior*, 83, 165–177.
- 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,” .
- 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.
- 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_{n+1}) such that $\omega_1 \rightarrow \omega_{n+1}$, define

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

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

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

with $\xi_1 = \omega_1$ and $\xi_{m+1} = \omega_{n+1}$. 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

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

Therefore,

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

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

A.2 The proof of Theorem 1.

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

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

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

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

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

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 $\bar{\omega}_n \in F(\omega_1)$.

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 $f(\omega_0) = 1$. Now, by induction on the distance¹² 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.

¹²The distance between two states in C is defined as the minimal number of edges in a path connecting them.

To show that f_C is F -measurable, consider $\omega \in F(\omega') \cap C$. Then, there is an F -cycle $\omega_1 = \omega, \dots, \omega_{n+1} = \omega'$. By [INC] and by the definition of f_C ,

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

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 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 $\bar{E}$,

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

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

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

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

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

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.¹³ We have to show that

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

Note that the path

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

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

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

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, $\bar{\omega}'_1 \in F(\omega_0)$.

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

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

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 $\prod_{j=\ell}^1 \varphi(\omega'_j, \bar{\omega}'_j) = \frac{f(\omega'_\ell)}{f(\bar{\omega}'_1)}$. 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:

¹³The asymmetry in the notation between $\bar{\omega}_k$ (with a bar) and ω'_ℓ (without) will become clear immediately.

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. $\square$

A.3 The proof of Theorem 2

Proof. Assume that τ is a signaling function and that s is a signal generated by τ with positive probability such that $\mu_{\tau,s} = \text{JB}$. Then, clearly, the set of states ω such that $\tau(s | \omega) > 0$ 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 $\mu_{\tau,s} = \text{JB}$ 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

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

and Theorem 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 $\tau(s | \omega) = 0$ for every $\omega \notin \Omega_+$, and otherwise, $\tau(s | \omega)$ is positive.

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

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

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

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

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 $\tau(s_0 \mid \omega) = 1 - f(\omega)$ for every $\omega \in \Omega_+$. $\square$

A.4 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 $\{\mu_1, \dots, \mu_n\}$ 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 $\{\mu_1, \dots, \mu_n\}$. 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 $\mu \in \text{conv}\{\nu_1, \dots, \nu_k\}$, 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 $\{\nu_1, \dots, \nu_k\}$.

Our assumption implies that $\mu \notin \text{conv}\{\mu_1, \dots, \mu_n\}$. 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 $u - (c, \dots, c)$, noting that the separation theorem is applied here to probability distributions, we conclude that the family $\{\mu_1, \dots, \mu_n\}$ 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,

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

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 $\{\mu_1, \dots, \mu_n\}$ 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 $\mu_1, \dots, \mu_n$, 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

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

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 $\nu \in \bar{D}$.¹⁴ 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 $\nu = \frac{1}{n} \sum_i \mu_i$, 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,

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

¹⁴Note that every option u derived from the separating hyperplane theorem can be adjusted with a positive constant to ensure the stated conditions.

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

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

Lemma 3. *If for every j there is a convex combination of $\mu_1, \dots, \mu_n$ 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 $\mu = \frac{1}{n} \sum_{j=1}^n \sum_i q_i^j \mu_i$, and note that in this combination, the weight of every μ_i is positive. $\square$

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

A.5 The proof of Proposition 1

Proof. To verify that N -[INC] implies I -[INC], observe that any cycle in G_I is also a cycle in G_N , since every connected component of G_I is contained within a connected component of G_N . Therefore, Eq. (7) holds in G_I due to N -[INC].

To verify I -[EXC], let the sequence of pairs $((\omega_i, \bar{\omega}_i))_{i=1}^n$ be an F-loop in G_I . Since the partition into CKCs in G_I is finer than that in G_N , we can group the pairs in this sequence according to their membership in the CKCs of G_N . Specifically, suppose there are k CKCs in G_N , denoted $C_1, \dots, C_k$, such that for each $\ell = 1, \dots, k$, we have

$$(\omega_j, \bar{\omega}_j) \in C_\ell \quad \text{for } j = i_\ell, \dots, i_{\ell+1} - 1.$$

By Lemma 1, the expression $\varphi_{\text{JB}}(\omega_{i_\ell}, \bar{\omega}_{i_{\ell+1}-1})$ is well defined, because they belong to the same CKC of G_N . Moreover, since for each i , the states $\bar{\omega}_i$ and ω_{i+1} belong to the same information set of the mediator, we have for every $\ell = 1, \dots, k$:

$$\prod_{j=i_\ell}^{i_{\ell+1}-1} \varphi_{\text{JB}}(\omega_j, \bar{\omega}_j) = \varphi_{\text{JB}}(\omega_{i_\ell}, \bar{\omega}_{i_{\ell+1}-1}). \quad (20)$$

It follows that:

$$\prod_{i=1}^n \varphi_{\text{JB}}(\omega_i, \bar{\omega}_i) = \prod_{\ell=1}^k \varphi_{\text{JB}}(\omega_{i_\ell}, \bar{\omega}_{i_{\ell+1}-1}).$$

Now consider two cases:

Case 1: $k \geq 2$. In this case, the sequence $(\omega_{i_\ell}, \bar{\omega}_{i_{\ell+1}-1})$ for $\ell = 1, \dots, k$ forms a loop in G_N .

By condition N -[**EXC**], it follows that:

$$\prod_{\ell=1}^k \varphi_{\text{JB}}(\omega_{i_\ell}, \bar{\omega}_{i_{\ell+1}-1}) = 1,$$

and hence $\prod_{i=1}^n \varphi_{\text{JB}}(\omega_i, \bar{\omega}_i) = 1$.

Case 2: $k = 1$. Then, by Eq. (20), $\prod_{i=1}^n \varphi_{\text{JB}}(\omega_i, \bar{\omega}_i) = \varphi_{\text{JB}}(\omega_1, \bar{\omega}_n)$, which equals 1. The reason is that, due to Eqs. (4) and (13), we have $\varphi_{\text{JB}}(\omega, \omega') = 1$ whenever ω and ω' lie in the same information set of the mediator. By the assumption that the sequence forms a loop, we deduce that ω_1 and $\bar{\omega}_n$ belong to the same information set of F .

Thus, in both cases, we conclude that φ_{JB} satisfies I -[**EXC**]. □