

Comparison of Oracles: Part I*

David Lagziel[†]

Ehud Lehrer[‡]

Tao Wang[§]

Ben-Gurion University

Durham University

CUEB

October 23, 2025

Abstract

We analyze incomplete-information games where an oracle publicly shares information with players. One oracle dominates another if, in every game, it can match the set of equilibrium outcomes induced by the latter. Distinct characterizations are provided for deterministic and stochastic signaling functions, based on simultaneous posterior matching, partition refinements, and common knowledge components. This study extends the work of Blackwell (1951) to games, and expands the study of Aumann (1976) on common knowledge, along with the companion Part II, which develops a theory of information loops.

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

Keywords: oracle; information dominance; signaling function; common knowledge component, information loops.

*For their valuable comments, the authors wish to thank participants of the Durham University Economics Seminar, the Adam Smith Business School Micro theory seminar of Glasgow University, INSEAD EPS seminar, the Tel-Aviv University Game Theory Seminar, the Rationality Center Game Theory Seminar, the Technion Game Theory Seminar, the Bar-Ilan University Theoretical Economics Seminar, the Bar-Ilan University Management Seminar, and the BGU Economics 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, and the Israel Science Foundation, Grant #591/21. Wang acknowledges the support of National Natural Science Foundation of China #72303161.

[†]Department of Economics, Ben-Gurion University of the Negev, Beer-Sheba 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.

[§]International School of Economics and Management, Capital University of Economics and Business, Beijing 100070, China. E-mail: tao.wang.nau@hotmail.com.

1 Introduction

In scenarios with incomplete information, players often have limited insight into the factors influencing outcomes. For this reason, an information provider, referred to as an *oracle*, can play a pivotal role in shaping players' strategies by revealing partial information about the underlying conditions. This partial revelation is akin to the information provided by various forecasters (ranging from weather and sports to geopolitics), news media organizations, rating agencies, and even prediction markets. In all these cases, external observers convey partial information to players engaged in strategic interactions.

This paper examines incomplete-information games where players are partially informed, both privately and publicly, about the realized state. The private information is provided to every player by his specific partition, and the public information is disclosed by an external source (namely, an oracle).

The oracle is endowed with a partition of the state space and communicates partial information through a *signaling function*, constrained by its ability to distinguish between different states. The oracle may not be aware of what is commonly known among the players, and typically possesses information that differs from that of at least some individuals. The combination of the original game—defined by the players' subjective information, action sets, and payoff functions—together with the additional information provided by the oracle, constitutes what we refer to as a *guided game*.

Any signaling function constitutes a Blackwell experiment (see Blackwell, 1951). In this sense, an oracle can be viewed as a generator of Blackwell experiments: it may produce multiple such experiments, each defining a distinct guided game. Each guided game, in turn, admits its own set of equilibria, which typically differs from that of the original incomplete-information game.

Our primary objective is to compare two oracles in terms of their ability to induce equilibria across all games. We say that one oracle *dominates* another if, for every game G and every signaling function of the latter, there exists a signaling function of the former such that the sets of equilibrium distributions over state-action profile pairs in the corresponding guided games coincide. In terms of players' payoffs, this means that, in any game, the dominating oracle can

replicate the entire set of equilibrium payoffs achievable by the dominated one.

A few remarks are in order. This study focuses on comparing different oracles while keeping the players' information fixed: both the players' private information and the oracles' informational capabilities are held constant. This stands in contrast to alternative notions that must hold *for every* possible configuration of players' private information (see, e.g., Section A.1 and Proposition 1 in Brooks et al., 2024).

The second remark concerns the definition of dominance among oracles. One could consider an alternative notion, where an oracle is said to dominate another if it can induce a *larger* set of equilibrium distributions, rather than exactly the same set, as required by our current definition. We adopt the stricter definition for the following reason: once partial information is provided, the oracle has no control over which specific equilibrium will emerge. A signaling function that induces a guided game with a large set of equilibria might allow for socially undesirable outcomes, such as those that are Pareto-dominated. By requiring that the dominant oracle replicate *precisely* the same set of equilibrium distributions as the dominated one in the stated game, we ensure that no outcome arises that could not have been generated by the latter. At the same time, any outcome achievable under the dominated oracle remains attainable under the dominating one.

The third remark concerns the objectives that oracles might have. One could define dominance by requiring that one oracle can induce any equilibrium *preferable* outcome that the other can. However, this definition implicitly assumes that the oracle can select which equilibrium will emerge, an assumption that typically does not hold.

An alternative approach is to model the oracle as a player with its own objectives. Under this interpretation, an oracle is said to dominate another if it can secure a payoff that is at least as high as that of the other. The drawback of this approach is that, as in any standard equilibrium analysis, all elements of the game, including the oracle's partition and payoff function, must be common knowledge. This assumption may be unrealistic in settings where oracles act as external information providers without known preferences.

Following Blackwell (1951), who compares signaling structures in the context of single-agent decision problems by focusing on the induced equilibrium distributions, and thus on the players'

achievable payoffs, we isolate the informational power of the oracles, abstracting away from any objectives they may possess.

Our analysis distinguishes between two types of oracles based on their signaling capabilities: deterministic and stochastic. In the deterministic case, when the oracle’s signaling function publicly announces a signal without any probabilistic element, we show that one oracle dominates another if and only if it can replicate the joint posterior beliefs induced by the other oracle (i.e., of *all players simultaneously*), while adjusting for redundancies arising from the players’ private information (see Theorem 1 in Section 4). We refer to this condition as *Individually More Informative* (IMI). In other words, the comparison takes into account that different players may interpret the same public signal differently, depending on their private information. The informational contribution of the oracle’s announcement must therefore be evaluated relative to what each player already knows.

Although the IMI condition may appear intuitive, it departs fundamentally from the refinement criterion implied by Blackwell’s notion of dominance, as becomes evident in the stochastic setting. Moreover, in contrast to Blackwell’s framework, we show that if two oracles dominate each other under the IMI condition, then they must be identical (see Theorem 2 in Section 4). We establish these results in the deterministic case before extending the analysis to the stochastic setting.

The conditions for dominance in the stochastic setting differ significantly from those in the deterministic case. When oracles are allowed to employ stochastic signaling functions, the resulting posterior beliefs become more intricate. As a result, establishing dominance requires additional criteria, which hinge on two key elements derived from the players’ information structures.

The first element is the *common knowledge component* (CKC)—the minimal (inclusion-wise) set that all players commonly agree upon (see Aumann, 1976). Building on the structure of CKCs, we introduce a second essential concept: the *information loop*. To formally define information loops, we first partition the state space into disjoint CKCs. An information loop is then described as a closed path through the state space that links different CKCs via elements of an oracle’s partition.

This paper is the first part out of two (the other being Lagziel et al., 2025).¹ The current paper provides the complete analysis for deterministic signaling functions, while establishing the baseline case of a single CKC given stochastic signaling functions. Part II expands the framework to incorporate information loops and explores their central role in determining when one oracle dominates (or is equivalent to) another when there are multiple CKCs.

Specifically, Theorem 4 in Section 5.2 establishes that in the case of a single CKC, one oracle dominates another if and only if the former refines the latter. While the “if” direction is straightforward, the “only if” direction is more subtle. To prove it, one must construct a counterexample: when Oracle 1’s partition does *not* refine that of Oracle 2, it is necessary to exhibit a game in which Oracle 2 can induce an equilibrium outcome that Oracle 1 cannot.

1.1 Relation to literature

The current research aims to extend the classical framework established by Blackwell (1951, 1953), which focuses on comparing experiments in decision problems. In Blackwell’s framework, one experiment (or information structure) dominates another if it is more informative, enhancing the decision maker’s expected utility across all decision problems. In the context of games, dominance implies that the information structure of one oracle enables it to replicate the equilibrium distribution over outcomes induced by the other oracle.

Another connection to Blackwell’s comparison lies in the fact that, in our study, an oracle can transmit any information through a signaling function, provided it is measurable with respect to the information it possesses. In this sense, an oracle in our framework functions as *a generator of experiments*, rather than a fixed entity as in Blackwell’s framework. However, unlike Blackwell’s comparison of experts (see Blackwell, 1951), our approach does not focus on optimizing the decision maker’s outcome. Instead, we analyze the role of oracles in inducing various equilibria.

Blackwell’s model was recently extended by Brooks et al. (2024), who compare two information sources (signals) that are robust to any external information source and decision problem. They introduce the notion of *strong Blackwell dominance* and characterize when one signal

¹The original (and unified) paper was split due to its length.

dominates another under this criterion: a signal strongly Blackwell dominates another if and only if every realization of the more informative signal either reveals the state or refines the realization of the less informative one.

There are several key differences between their framework and ours. First, while their analysis focuses on a single decision maker, we study multi-player environments. Second, they allow for arbitrary private information structures and decision problems; in fact, their characterization is entirely independent of the decision maker’s information. In contrast, our model assumes fixed private information structures for the players and allows variation only in the payoff functions of the underlying game. As a result, our analysis is specific to each configuration of the players’ information structures: every distinct configuration must be analyzed separately. A third major difference lies in the role of the oracle. In their model, the oracle is a fixed Blackwell experiment. In contrast, in our setting, the oracle can generate any experiment that is measurable with respect to its partition, effectively acting as a generator of Blackwell experiments.

Beyond Blackwell’s work, this project runs parallel to and is inspired by two additional lines of research. The first concerns the topic of Bayesian persuasion. Originating from the classic model of Kamenica and Gentzkow (2011), the literature on Bayesian persuasion explores how an informed sender should communicate with an uninformed receiver to influence the receiver’s choices. The central question revolves around how much information—and in some contexts, when—should the sender disclose to maximize their payoff.²

The second strand of literature explores the role of an external mediator in games with incomplete information. The mediator provides players with differential information to coordinate their actions, resulting in outcomes that correspond to various forms of correlated equilibria, as introduced by Forges (1993). Importantly, in some of these studies, the mediator does not supply additional information about the realized state but focuses solely on coordinating the players’ actions. Gossner (2000) examines games with complete information, comparing mediating structures that induce correlated equilibria. The mediator’s role is exclusively to co-

²See, for example, Hörner and Skrzypacz (2016); Renault et al. (2013); Ganglmair and Tarantino (2014); Hörner and Skrzypacz (2016); Renault et al. (2017); Ely (2017); Ely and Szydlowski (2020); Che and Hörner (2018); Bizzotto et al. (2021); Mezzetti et al. (2022). For a survey of this field, see Kamenica (2019).

ordinate the players' actions. One mediator is considered "richer" than another if the set of correlated equilibria it induces is a superset of those induced by the other. The characterization is based on the concept of compatible interpretation, which aligns with the spirit of Blackwell's notion of garbling.

Other studies, closely aligned with the current project's goals, investigate information structures in incomplete-information games and establish partial orderings among them. Peski (2008) analyzed zero-sum games, offering an analogous result to Blackwell's by characterizing when one information structure is more advantageous for the maximizer. Lehrer et al. (2010) examines a common-interest game, comparing two experiments that generate private signals for players, which may be correlated. The results depend on the type of Blackwell's notion of garbling used, which varies with the solution concept applied. In a follow-up study, Lehrer et al. (2013) extended Blackwell's garbling to characterize the equivalence of information structures in incomplete-information games, specifically by determining when they induce the same equilibria. Likewise, Bergemann and Morris (2016) characterize dominance among two information structures through the concept of individual sufficiency—an extension of Blackwell's notion of garbling to n -player games. They show that an information structure S is individually sufficient for S' if and only if under every game, S yields a smaller set of Bayesian correlated equilibria.

In this study, we fix the players' initial information structures and compare oracles that provide additional information, which in turn influences the players' beliefs. The key distinction of our study lies in two main aspects: (a) the information provided by the oracles is public, and therefore does not serve as a coordinator between the players' actions, as in various versions of correlated equilibrium; (b) since an oracle functions as a generator of experiments, we allow the externally provided information to vary. Additionally, we do not impose any restrictions on the type of game, whether it involves a common objective, a zero-sum structure, or any other form. While previous results align with Blackwell's garbling, our findings differ significantly from any version of it.

This approach presents a unique challenge compared to the problem of comparing two fixed information structures, as explored in previous literature. The distinction becomes evident in the example in Section 2, where the oracles are evaluated based on the full range of signaling

functions they can generate. From an applied perspective, in many real-life scenarios, information providers have multiple ways to share information with the public, making it crucial to compare them as generators of information.

1.2 The structure of the paper

The paper is organized as follows. In Section 2, we provide a simple example to illustrate the key concepts of the paper. Section 3 presents the model and key definitions. Section 4 analyzes deterministic oracles, including a characterization of dominance and a proof that two-sided dominance implies the oracles are identical (given a unique CKC). In Section 5, we examine stochastic oracles in several stages. First, we introduce a two-stage game, referred to as a "game of beliefs," which serves as a foundational tool for our characterization within each CKC. Then, in Section 5.2, we characterize dominance in the case of a unique CKC.

2 A simple example: the rock-concert standoff

Consider a simple competition between two rock bands.³ Assume two bands, 1 and 2, arrive in the same city during their tours and must decide whether to perform on the same day or on different days. The issue arises because the stadiums in that city are partially open, making bad weather a significant factor that adversely affects crowd attendance.

Assume there are 200,000 fans eager to see these bands, with ticket prices fixed at \$20 each. The production cost for each concert is \$500,000, but this cost doubles if attendance exceeds 75,000 people. Further, assume that each fan attends at most one concert.

On a sunny day, all fans attend the concerts, splitting evenly if both bands perform on the same day. However, under stormy conditions, attendance drops to 20,000 fans, who again split evenly if both bands perform simultaneously. If the bands choose to perform on different days, attendance splits such that only 10% of the fans attend the concert on the stormy day, with the remaining fans attending the other concert.

As it turns out, weather conditions are problematic because a storm is coming either today

³We thank Alon Eizenberg from the Hebrew University and two 1990s rock bands who inspired this example.

or tomorrow. More formally, there are four equally likely states: in states n_1 and n_2 , the storm arrives today, while in states s_1 and s_2 , the storm arrives tomorrow. Each band has a unique partition over this state space. Band 1's partition is $\Pi_1 = \{\{n_1, s_2\}, \{n_2, s_1\}\}$, while Band 2's partition is $\Pi_2 = \{\{n_1\}, \{s_1\}, \{n_2, s_2\}\}$. In simple terms, Band 2 cannot differentiate between n_2 and s_2 , while Band 1 cannot distinguish between n_i and s_{-i} for each $i = 1, 2$. Additionally, there are two weather forecasters with the following partitions: $F_1 = \{\{n_1, s_1\}, \{n_2\}, \{s_2\}\}$ and $F_2 = \{\{n_1, n_2\}, \{s_1, s_2\}\}$. These information structures are illustrated in Figure 1.




Figure 1: On the left, Figure (a) illustrates the information structures: $\Pi_1 = \{\{n_1, s_2\}, \{n_2, s_1\}\}$ for Band 1 (green) and $\Pi_2 = \{\{n_1\}, \{s_1\}, \{n_2, s_2\}\}$ for Band 2 (orange). On the right, Figure (b) depicts the information structures $F_1 = \{\{n_1, s_1\}, \{n_2\}, \{s_2\}\}$ for Forecaster 1 (red) and $F_2 = \{\{n_1, n_2\}, \{s_1, s_2\}\}$ for Forecaster 2 (blue). These figures illustrate a unique CKC where neither of the Forecasters' partitions refines the other. Nevertheless, Forecaster 1 is individually more informative (IMI) than Forecaster 2, whereas the converse does not hold. This is because Forecaster 2 cannot replicate the partition $F'_1 = \{\{n_1, s_1, s_2\}, \{n_2\}\}$.

Based on the realized state, the bands engage in the game depicted in Figure 2. Each band decides whether to perform today, an action denoted by D , or tomorrow, denoted by M . The payoffs in the matrices are given in hundreds of thousands of dollars, and the bands' actions have opposing impacts depending on the state of nature.

Conditional on the state, it is evident that each band has a strictly dominant action: to perform on the day with good weather. Consequently, the analysis is straightforward. If both bands know the exact payoff matrix, there is a unique Nash equilibrium. However, this equilibrium is not necessarily optimal in terms of overall profit, which could be maximized if

the bands coordinated and split the performance dates.

		Band 2	
		D	M
Band 1	D	-3, -3	-1, 26
	M	26, -1	10, 10

Payoffs in states n_1 and n_2 (stormy today)

		Band 2	
		D	M
Band 1	D	10, 10	26, -1
	M	-1, 26	-3, -3

Payoffs in states s_1 and s_2 (stormy tomorrow)

Figure 2: Payoff matrices for sunny and stormy conditions.

Assume that the payoff matrix is not common knowledge. If Band 1 knows the exact payoff matrix while Band 2 believes the two matrices are equally likely (and assuming this is common knowledge), an equilibrium exists in which Band 2 randomizes equally between M and D due to symmetry, and Band 1 selects M under $\{n_1, n_2\}$ and D given remaining states. This equilibrium yields, on aggregate, higher expected payoffs of \$1.8 million for Band 1 and \$450,000 for Band 2.

Now, we examine how the two different forecasters can influence the outcome of this game. For simplicity, assume that forecasters are restricted to deterministic strategies, meaning they provide deterministic public signals based on their information. Forecaster 2 has only two options: either provide no information at all (which, in some cases, leads both bands to perform in stormy conditions) or fully reveal all relevant information, which results in an expected payoff of \$1 million for each band. Forecaster 1 also has these two options, as fully revealing his private information makes the realized state common knowledge between the two bands. In such cases, we say that Forecaster 1 is individually more informative than Forecaster 2.

Yet, Forecaster 1 can achieve more than simply matching the beliefs induced by Forecaster 2. Specifically, he can signal the partition $\{\{n_1, s_1\}, \{n_2, s_2\}\}$, ensuring that Band 1 is fully informed about the state and the corresponding payoff matrix, while Band 2 receives no additional information and remains unable to distinguish between n_2 and s_2 . Under these conditions and given either of the states n_2 and s_2 , the previously described equilibrium, in which the expected payoffs are \$1.8 million and \$450,000 for Bands 1 and 2 respectively, still exists. Thus, Forecaster 1 can support a broader set of equilibria while also matching the set of equilibria induced by Forecaster 2. This exemplifies the partial order of dominance characterized in this study.

This simple example offers several additional insights. First, the state space comprises a unique CKC, given the bands' information. In other words, the smallest set (in terms of inclusion) that the bands can agree upon is the entire space. However, the forecasters' partitions do not refine one another, even within this unique CKC, meaning that the IMI condition does not imply refinement. Moreover, when stochastic signals are allowed, we later show that neither forecaster dominates the other.⁴

Second, if this was a decision problem (as in Blackwell, 1951 and Brooks et al., 2024) rather than a game, both forecasters would be equally beneficial to both parties. In decision problems, superior information can only improve the expected outcome, and both forecasters could fully reveal the true state to each party. This highlights a key distinction: the classification in games is fundamentally different from that in decision problems and does not follow from it.

Third, the ability to induce a broader set of outcomes is distinct from coordination in the sense of correlated equilibrium (as in Forges, 1993). The process here relies critically on the forecasters' private information and how it is disclosed to the players.

3 The model

A *guided game* comprises a Bayesian game and an *oracle*. The oracle's role is to provide information that enables a different, and preferably broader, range of equilibria. It does so through signaling, and our analysis seeks to characterize the extent to which oracles can expand the set of equilibrium payoffs.

We begin by defining the underlying Bayesian game. Let $N = \{1, 2, \dots, n\}$ be a finite set of $n \geq 2$ players, and let Ω denote a non-empty, finite state space. Each player $i \in N$ has a non-empty, finite set of actions⁵ A_i and a partition Π_i over Ω , representing the information available to player i . Denote the set of action profiles by $A = \times_{i \in N} A_i$. The utility function for each player $i \in N$ is $u_i : \Omega \times A \rightarrow \mathbb{R}$, which maps states and action profiles to real-valued payoffs.

⁴Notably, given a unique CKC, we prove that two-sided IMI implies that the two partitions coincide. See Section 4.2.1.

⁵In this setting, A_i is independent of the player's information; however, the current framework can also accommodate scenarios where it is not.

To extend the basic game into a guided game, we introduce an oracle who provides public information before players choose their actions. The oracle is endowed with a partition F of the state space Ω , and a countable set S of possible signals.

A *signaling strategy* of the oracle is an F -measurable function $\tau : \Omega \rightarrow \Delta(S)$, where $\Delta(S)$ is the set of probability distributions on S with finite support. This function specifies the distribution over signals conditional on the realized state and must be measurable with respect to F . We denote by $\tau(s|\omega)$ the probability $\tau(\omega)(s)$ that the oracle sends signal s when the realized state is ω .

A *deterministic signaling strategy* is a special case in which τ is a function from F to S , assigning a single signal to each element of the partition. Note that any deterministic signaling strategy is effectively equivalent to a partition, and we will refer to it as such when appropriate.

The guided game evolves as follows. First, the oracle publicly announces a strategy τ . Then, a state $\omega \in \Omega$ is drawn according to a common prior $\mu \in \Delta(\Omega)$. Each player i is privately informed of $\Pi_i(\omega)$, which is a set of states containing ω and also an atom of player i 's private partition. Finally, the signal $\tau(\omega) \in S$ is publicly announced in case τ is deterministic, or $s \in S$ is drawn according to $\tau(\omega)$ and is publicly announced in case τ is stochastic.

Let the join⁶ $\Pi_i \vee F'$ denote the updated information (i.e., partition) of player i given Π_i and some partition F' . In case τ is a deterministic function, let $\mu_{\tau|\omega}^i = \mu(\cdot | [\Pi_i \vee \tau](\omega)) \in \Delta(\Omega)$ denote player i 's posterior belief after observing $\Pi_i(\omega)$ and $\tau(\omega)$. In case τ is stochastic, let $\mu_{\tau|\omega,s}^i = \mu(\cdot | \Pi_i(\omega), \tau, s) \in \Delta(\Omega)$ denote player i 's posterior belief after observing $\Pi_i(\omega)$ and a realized signal s according to $\tau(\omega)$, and let $\mu_{\tau,s} = \{(\mu_{\tau|\omega,s}^i)_{i \in N} : \omega \in \Omega \text{ s.t. } \tau(s|\omega) > 0\}$ be the set of *joint posteriors* associated with τ and a signal s , across all relevant states. The joint posteriors capture each player's belief about the realized state and their beliefs about others' beliefs, as well as higher-order beliefs. We use μ_τ to denote the distribution over all joint posteriors induced by τ across all signals, and use $\text{Post}(\tau) = \text{Supp}(\mu_\tau)$ to denote its support. Thus, every strategy τ yields an incomplete-information game $G(\tau) = (N, (A_i)_{i \in N}, \mu_\tau, (u_i)_{i \in N})$. Since the state space and the action sets are finite, the equilibria of the game exist. When there is no risk of ambiguity, we denote the incomplete-information game without τ by G .

⁶Coarsest common refinement of Π_i and F' ; following the definition of Aumann (1976).

Example 1. *Deterministic and stochastic strategies.*

To illustrate the difference between deterministic and stochastic strategies, consider an information structure where $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}\}$, μ is the uniform distribution on Ω , and Oracle 1 has complete information, $F_1 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}\}$. Under deterministic strategies, the feasible posteriors are generated by either Π_1 (oracle provides no additional information) or F_1 (complete information). On the other hand, the set of feasible posteriors under stochastic strategies includes distributions of the form $(p, 1 - p, 0)$ for every $p \in [0, 1]$.

3.1 Partial ordering of oracles

To discuss the role of the oracle in the current framework, one needs a relevant solution concept. Thus, let us define the following notion of a Guided equilibrium, which incorporates the oracle's strategy. Formally, let $\sigma_i : \Pi_i \times S \rightarrow \Delta(A_i)$ be a strategy of player i . A tuple $(\tau, \sigma_1, \dots, \sigma_n)$ is a *Guided equilibrium* if $(\sigma_1, \dots, \sigma_n)$ is a Nash equilibrium in the incomplete-information game $G(\tau)$.

The notion of a Guided equilibrium defines a partial ordering of oracles, i.e., a partial relation over their partitions according to the sets of equilibria. To define this relation, let $\text{NED}(G(\tau)) \subseteq \Delta(\Omega \times A)$ be the set of distributions over $\Omega \times A$ induced by Nash equilibria given G and τ .⁷ Now consider two oracles, Oracle 1 and Oracle 2, and denote the generic partition and strategy of Oracle j by F_j and τ_j , respectively. Using these notations we define a partial ordering of oracles as follows.

Definition 1 (Partial ordering of oracles). *Fix the players' information structures. We say that Oracle 1 dominates Oracle 2, denoted $F_1 \succeq_{\text{NE}} F_2$, if for every τ_2 and game G , there exists τ_1 such that $\text{NED}(G(\tau_1)) = \text{NED}(G(\tau_2))$.*

In simple terms, dominance implies that one oracle can mimic the signaling structure of the other to induce the same equilibria. Note that a direct comparison of the games' equilibria is problematic because the players' strategies depend on the oracles' signaling functions.

⁷Note that a Nash equilibrium $(\sigma_i^*, \dots, \sigma_n^*)$ induces a probability distribution over $\Omega \times A$. Specifically, fix ω and an action profile a , the probability of (ω, a) under the equilibrium strategy $(\sigma_i^*, \dots, \sigma_n^*)$ and the signaling function τ is given by $\mu(\omega) \sum_{s \in S} \tau(s|\omega) \prod_{i=1}^n \sigma_i^*(a_i|\Pi_i(\omega), s)$. Since multiple equilibria can exist, $\text{NED}(G(\tau))$ is a subset of $\Delta(\Omega \times A)$.

Two points are worth noting here. First, if the players' information structures were unknown, one might consider defining the dominance order between oracles in a more flexible way, allowing for a variety of possible partitions. In that case, the characterization problem would likely become easier. The challenge in our framework arises from the fact that the partitions are predetermined.

The second point highlights that Definition 1 compares the equilibria induced by the oracles. An alternative, weaker condition could involve, for example, an inclusion criterion based on the set of equilibria or the players' expected payoffs. We relate to these possibilities in Section 3.2 below. Nevertheless, we use the more general definition to address potential issues that may arise from different equilibrium-selection processes. Since we do not restrict ourselves to a specific selection process (which may diverge from the Pareto frontier), a broader set of equilibria might not always benefit the players. This approach also addresses complications that could emerge in a parallel setup, if oracles were to maximize some goal function.

Definition 1 also allows us to define equivalent oracles. Formally, we say that Oracle 1 is *equivalent to* Oracle 2, denoted $F_1 \sim F_2$, if each oracle dominates the other. The characterization of equivalent oracles is one of the main results of Lagziel et al. (2025).

3.2 Alternative definitions of dominance

One could consider other notions of dominance, which might involve different types of comparisons between outcomes, such as combinations of (state, action-profiles), or comparisons based on equilibrium payoffs.

An alternative definition of dominance could be based on an inclusion criterion concerning the distribution over outcomes. Specifically, Oracle 1 dominates Oracle 2 in the inclusive sense, if and only if, for every τ_2 and game G , it holds that

$$\text{NED}(G(\tau_2)) \subseteq \bigcup_{\tau_1} \text{NED}(G(\tau_1)).$$

This is a weaker condition than the one currently used. It implies that Oracle 1 dominates Oracle 2 if any equilibrium distribution of outcomes induced by τ_2 can be generated by some

τ_1 . Unlike the condition in Definition 1, this alternative allows for different distributions over outcomes induced by τ_2 to be generated by different τ_1 strategies.

Another approach to the issue of dominance could involve comparisons between equilibrium payoffs. Specifically, for any game G and a signaling function τ , let $\text{NEP}(G(\tau))$ denote the set of Nash-equilibrium expected-payoffs profiles induced by τ . Oracle 1 is said to dominate Oracle 2 in the payoff sense if, for every τ_2 and game G , there exists a τ_1 such that

$$\text{NEP}(G(\tau_1)) = \text{NEP}(G(\tau_2)).$$

Alternatively, Oracle 1 dominates Oracle 2 in the inclusive-payoff sense if, for every τ_2 and game G , it holds that

$$\text{NEP}(G(\tau_2)) \subseteq \bigcup_{\tau_1} \text{NEP}(G(\tau_1)).$$

The concepts related to equilibrium outcome distributions imply their corresponding payoff-related notions. Definitions based on equilibrium outcome distributions are better suited for policy designers that prioritize outcomes, such as individuals' actions and their aggregate effects, over individual payoffs. Conversely, definitions grounded in equilibrium payoffs are more appropriate for contexts where the primary focus is on individual payoffs.

An interesting direction for future research would be to identify the precise settings, if any, where the various definitions diverge. We leave this question open for further investigation. In the following, we adopt Definition 1.

3.3 The oracles as players

Another way to compare oracles is to treat them as players. In the spirit of sender–receiver games, the oracle takes the role of the sender, responsible for providing information, while the other players act as receivers, making decisions based on both their private information and the signals they receive. In this framework, the oracle's objective is to maximize its equilibrium payoff in the resulting game of incomplete information. One could then compare two oracles by saying that one is more informative than the other if, in every such game, the former always secures a (weakly) higher equilibrium payoff than the latter.

However, this approach has several drawbacks relative to ours. First, such games typically admit multiple equilibria, making it unclear which equilibrium payoff should be the basis for comparison. Second, equilibrium analysis generally presumes that players’ information partitions are common knowledge. In particular, it assumes that the oracles know the private information structures of the players. In contrast, our approach imposes significantly weaker assumptions: one oracle can often imitate another without requiring full knowledge of players’ information structures. In fact, even identifying the components that are common knowledge is sometimes unnecessary. While our comparison focuses exclusively on the equilibrium outcomes of the game played by the players, we assume that the private information structures are common knowledge among the players themselves—but not necessarily known to the oracle.

The third advantage of our approach is that, by focusing on the equilibrium outcomes of the game played by the agents, we can analyze the information structures of the oracles independently of any objectives they might have. This enables us to concentrate on informational aspects and to introduce new concepts into the model, such as informational loops and clusters (in Lagziel et al., 2025).


3.4 The case of one decision maker

3.4.1 The oracle contributes to DM’s private information

To illustrate a key contribution of this paper and connect it to the current body of knowledge, consider a decision problem with one decision-maker (DM) and two oracles. When Oracle i employs a signaling strategy τ_i , the DM also gains access to his own partition Π . The combination of the signaling strategy τ_i and the partition Π induces a Blackwell experiment $M_i(\tau_i, \Pi)$.


Example 2. One decision maker and two oracles.

Consider the uniformly distributed state space $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, with a single DM whose private information is represented by the partition $\Pi = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$, while the oracles’ partitions are given by $F_1 = \{\{\omega_1, \omega_4\}, \{\omega_2\}, \{\omega_3\}\}$, and $F_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. This information structure is illustrated in Figure 3.



(a)

The DM's information



(b)

The oracles' information

Figure 3: On the left, Figure (a) illustrates the information structure of the DM (blue). On the right, Figure (b) portrays the information structure of Oracle 1 (orange) and Oracle 2 (green).

Now, consider the stochastic strategy τ_2 given in Figure 4. Combined with Π , this signaling strategy τ_2 is equivalent to the following Blackwell experiment e , given in Figure 5.

$\tau_2(s \omega)$	s_1	s_2	s_3
ω_1	0	1/2	1/2
ω_2	1/4	3/4	0
ω_3	0	1/2	1/2
ω_4	1/4	3/4	0

Figure 4: A stochastic F_2 -measurable signaling strategy of Oracle 2.

$e(s \omega)$	s_1, L	s_1, R	s_2, L	s_2, R	s_3, L	s_3, R
ω_1	0	0	1/2	0	1/2	0
ω_2	1/4	0	3/4	0	0	0
ω_3	0	0	0	1/2	0	1/2
ω_4	0	1/4	0	3/4	0	0

Figure 5: $M_2(\tau_2, \Pi)$ - the matrix consisting of the probabilities.

Blackwell's Theorem states that, given a signaling strategy τ_2 employed by Oracle 2, the DM can achieve at least as much as he could by obtaining information from Oracle 1 with signaling strategy τ_1 if and only if there exists a stochastic matrix G (the garbling) such that:

$$M_1(\tau_1, \Pi)G = M_2(\tau_2, \Pi).$$

This fact immediately implies the following extension of Blackwell's Theorem:

Observation 1. *Suppose there is a single DM with a partition Π and two oracles with partitions F_1 and F_2 , respectively. Then, $F_1 \succeq_{\text{NE}} F_2$ if and only if, for every signaling strategy τ_2 of Oracle 2, there exists a signaling strategy τ_1 of Oracle 1 such that $M_1(\tau_1, \Pi)G = M_2(\tau_2, \Pi)$, for some garbling matrix G .*

Note that in the case of a single decision maker, equilibrium implies that the equilibrium payoff is the best achievable. In addition, the statement that for every signaling strategy τ_2 of Oracle 2, there exists a signaling strategy τ_1 of Oracle 1 such that $M_1(\tau_1, \Pi)G = M_2(\tau_2, \Pi)$, for some garbling matrix G is equivalent to $F_1 \succeq_{\text{NE}} F_2$.

The stochastic matrix $M_i(\tau_i, \Pi)$ is the combination of two separate stochastic matrices, τ_i and the one corresponding to Π . For Blackwell dominance, we considered $M_1(\tau_1, \Pi)$ and $M_2(\tau_2, \Pi)$. Another possibility is to consider the Blackwell dominance between τ_1 and τ_2 first. If τ_1 Blackwell dominates τ_2 and both τ_1 and τ_2 are independent of Π , then $M_1(\tau_1, \Pi)$ Blackwell dominates $M_2(\tau_2, \Pi)$ (see Theorem 12.3.1 of Blackwell and Girshick (1954)).⁸ Nevertheless, the reverse does not hold. Consider, for instance, that Π is fully informative, then $M_1(\tau_1, \Pi)$ Blackwell dominates $M_2(\tau_2, \Pi)$, but it does not imply that τ_1 dominates τ_2 . Hence, dominance in terms of $M_1(\tau_1, \Pi)$ and $M_2(\tau_2, \Pi)$ is weaker than the dominance in terms of signaling functions τ_1 and τ_2 .

This characterization of dominance is expressed in terms of stochastic matrices. Specifically, the question of whether $M_2(\tau_2, \Pi)$ can be obtained from $M_1(\tau_1, \Pi)$ by taking its product with a garbling matrix reduces to a problem about transforming one set of stochastic matrices into another. However, this characterization is not directly expressed in terms of the model's primitives, namely the information partitions.

In this paper, we focus on comparing information structures rather than analyzing the algebraic properties of the corresponding sets of matrices. Our primary objective is to examine the relationship between two oracles based on the model's primitives, specifically their partitions. Referring to Example 2, we later demonstrate that Oracle 2 cannot imitate Oracle 1. This naturally raises the question: why? What is the underlying reason? We aim to shed light on

⁸Note that for this result to hold, Π is fixed and it is independent of τ_1 and τ_2 .

this issue while also pursuing the second objective of the paper—extending Blackwell’s model to a setting with multiple players.

3.5 Common objectives

The game-theoretic setting closest to a one-agent decision problem is one in which all players share a common objective.⁹ A natural conjecture is that one oracle induces at least as high a payoff as another in any common-objective game if and only if its partition refines that of the other. It turns out that this is not the case.

Example 3.

In this example, there are four states and two players. The following Figure 6 illustrates the knowledge structures of the players as well as those of the two oracles. It is clear that the partition of Oracle 2 refines that of Oracle 1.




Figure 6: 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 Oracle 1 (orange) and Oracle 2 (green).

Now consider a game where both player have two actions: D and M , and the payoffs are given by the matrices in Figure 7. The best common payoff is attained when both players know

⁹As this section serves primarily as a comment, we do not undertake a detailed discussion of the definition of a common objective. For our purposes, we assume that all players' payoff functions are identical.

		Player 2	
		D	M
Player 1	D	1,1	0,0
	M	0,0	0,0
		ω_1	

		Player 2	
		D	M
	D	0,0	0,0
	M	0,0	1,1
		ω_2	

		Player 2	
		D	M
	D	0,0	0,0
	M	1,1	0,0
		ω_3	

		Player 2	
		D	M
	D	0,0	1,1
	M	0,0	0,0
		ω_4	

Figure 7: Payoff matrices for each ω

the realized state. Oracle 2, who is fully informed, can simply reveal the true state. Oracle 1, who cannot distinguish between ω_1 and ω_3 , can nonetheless reveal his information; combined with the players' private knowledge, this is sufficient to fully disclose the state.

While our focus is not on comparing oracles based on the highest equilibrium payoffs they can induce, the following proposition provides an affirmative answer to a question naturally motivated by this example.

Proposition 1. *In any common-objective game, Oracle 1 can induce an equilibrium expected payoff at least as high as any induced by Oracle 2 if and only if, for every player i , the combined information of F_1 and Π_i refines that of F_2 and Π_i .*

The proof is deferred to the Appendix and relies on terminology introduced later in the paper.

4 Partial ordering of deterministic oracles

Our first main result characterizes the notion of dominance among oracles, assuming they are restricted to deterministic strategies. That is, throughout this section, we only consider oracles that use deterministic functions, namely $\tau_i : F_i \rightarrow S$ for every oracle i , and we can relate to every such strategy as a partition (as previously noted).

The characterization is based on the ability of one oracle to *match* the players' joint posterior beliefs, for any given strategy of the other oracle. More formally, we say that Oracle 1 is *individually more informative* (IMI) than Oracle 2, if for every strategy τ_2 , there exists a strategy τ_1 that simultaneously matches the posterior partition of every player i .

Definition 2. Oracle 1 is individually more informative than Oracle 2, denoted $F_1 \succeq_{(\mu^i)_i} F_2$, if for every deterministic τ_2 , there exists a deterministic τ_1 such that $\Pi_i \vee \tau_1 = \Pi_i \vee \tau_2$ for every player i .

In other words, one oracle is more informative than another if it can always ensure that every player has the same information as provided by the other oracle, taking into account the player's private information, namely the redundancies given the players' private information, as well as the publicly available signal (restricted to deterministic signaling functions). A different way of defining the same relation is through partitions' refinements, as given in the following observation.

Observation 2. Oracle 1 is individually more informative than Oracle 2 if and only if for every $F'_2 \subseteq F_2$,¹⁰ there exists $F'_1 \subseteq F_1$ such that $\Pi_i \vee F'_1 = \Pi_i \vee F'_2$, for every player i .

Note that Observation 2 follows directly from Definition 2 because every F_i -measurable deterministic strategy τ_i induced a sub-partition F'_i of F_i and vice versa. Nevertheless, what should be clear is that the notion of IMI differs from the notion of refinement, as the following example illustrates.

Example 4. *Individually More Informative versus refinement.*

The partial ordering generated by the notion of “individually more informative than” need not coincide with the notion of “finer than”. Consider, for example, the three partitions $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$, $F_1 = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}\}$ and $F_2 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$. Note that F_2 strictly refines F_1 and Π_1 , but Oracle 1 remains individually more informative than Oracle 2. This is illustrated in Figure 8. Nevertheless, in Section 4.2.1 we prove that if Oracle 1 is IMI than Oracle 2 and vice versa, then their partitions partially coincide.

One can also bridge the gap between the notions of IMI and refinement by considering the possibility that the players' partitions are not fixed.¹¹ In other words, we can also consider the possibility that Oracle 1 is IMI than Oracle 2 for *any* set of the players' partitions. Once we

¹⁰A partition F'_2 is a subset of partition F_2 if the σ -field generated by F'_2 is a subset of the σ -field generated by F_2 .

¹¹This resembles the condition of strong Blackwell dominance, in the context of decision problems, in Brooks et al. (2024).




Figure 8: The notion of “individually more informative than” does not imply “finer than”, though the latter does imply the former. In this figure, F_2 (red) strictly refines F_1 (green) and Π_1 (blue), but for every deterministic τ_2 , there exists a deterministic τ_1 such that $\Pi_1 \vee \tau_1 = \Pi_1 \vee \tau_2$, so F_1 is individually more informative than F_2 .

account for all possible partitions, we must also account for the trivial partition, so that Oracle 1 must match any deterministic strategy of Oracle 2. This implies that F_1 refines F_2 , at least weakly.

4.1 First characterization result - deterministic oracles

Our first main result, given in Theorem 1 below, presents an equivalence between oracle dominance and the notion of individually more informative. Specifically, we prove that one oracle dominates another if and only if it is individually more informative. The proof is constructive. We assume that Oracle 1 is not more informative than Oracle 2, and depict a game such that the players’ expected payoffs given a deterministic strategy τ_2 differ from their expected payoffs for every deterministic strategy τ_1 . The game is constructed such that a strictly more informative τ_1 , in the sense that $\Pi_i \vee \tau_1$ refines $\Pi_i \vee \tau_2$ for some player i , yields a strictly higher expected payoff for the players, whereas a (strictly) less informative τ_1 yields a strictly lower expected payoff. (Unless stated otherwise, all proofs are deferred to the Appendix.)

Theorem 1. *Assume that oracles are deterministic. Then, Oracle 1 dominates Oracle 2 if and only if Oracle 1 is individually more informative than Oracle 2.*

Though the proof of Theorem 1 is deferred to the appendix, let us provide some intuition for it. The first derivation is straightforward—if Oracle 1 can simultaneously match the information

available to every player given τ_2 , then the sets of equilibria coincide. We emphasize that Oracle 1 actually *matches* the information conveyed by Oracle 2, so the set of equilibria can be preserved by Oracle 1, even if, for instance, there exists a specific equilibrium selection process that influences the players' expected payoffs in one way or another.

Proving the reverse statement is a bit more difficult. To gain some intuition for this result, consider a single-player decision problem. If Oracle 1 is not individually more informative than Oracle 2, then there exists a strategy τ_2 such that for every τ_1 there are two possibilities: either $\Pi_1 \vee \tau_1$ strictly refines $\Pi_1 \vee \tau_2$, or there exists an element of $\Pi_1 \vee \tau_1$ that intersects two elements of $\Pi_1 \vee \tau_2$.

For this purpose, we design a game based on the partition elements of $\Pi_1 \vee \tau_2$. Namely, for every element B in $\Pi_1 \vee \tau_2$, take all permutations $p : B \rightarrow \{1, 2, \dots, |B|\}$. The player's action set is the set of all such permutations. Once a state ω is realized and an action p is chosen, the player receives a payoff that depends on $p(\omega)$ in case p is supported on the realized state, or a very low negative payoff otherwise. Figure 9 below depicts a specific example for this payoff function given a uniform distribution on four possible states and two partition elements in $\Pi_1 \vee \tau_2$. Thus, if $\Pi_1 \vee \tau_1$ strictly refines $\Pi_1 \vee \tau_2$, the player can secure a strictly higher expected payoff, and if an element of $\Pi_1 \vee \tau_1$ intersects two disjoint elements of $\Pi_1 \vee \tau_2$, the player receives a very low expected payoff. Either way, expected payoffs are either higher or lower given τ_1 , relative to τ_2 , and the result follows.

An example with 4 states and two partition elements in $\Pi_1 \vee \tau_2$

	ω_1	ω_2	ω_3	ω_4
a_1	1	2	3	-2^{42}
a_2	1	3	2	-2^{42}
a_3	2	1	3	-2^{42}
a_4	2	3	1	-2^{42}
a_5	3	1	2	-2^{42}
a_6	3	2	1	-2^{42}
a_7	-2^{42}	-2^{42}	-2^{42}	1

Figure 9: Assume that $\Omega = \Pi_1 = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ and μ is the uniform distribution. Further assume that $\Pi_1 \vee \tau_2$ consists of two elements $B_1 = \{\omega_1, \omega_2, \omega_3\}$ and $B_2 = \{\omega_4\}$. So, there are 6 permutations/actions for B_1 and a single one for B_2 . If $\tau_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$, then the player can secure a strictly higher expected payoff, and if $\tau_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ the player would get -2^{42} with positive probability, thus generating a strictly lower expected payoff.

Remark 1. *We repeatedly use the fact that if the players' expected payoffs in equilibrium differ when following τ_1 instead of τ_2 , then $\text{NED}(G(\tau_1)) \neq \text{NED}(G(\tau_2))$ for the specified game G . This holds because μ is fixed, meaning that every element in $\Delta(\Omega \times A)$ determines the players' expected payoffs in the corresponding equilibrium. The reverse deduction, however, is not necessarily true, as different such distributions may, in fact, yield the same expected payoffs.*

Remark 2. *In situations where the information available to the players is unknown, a reasonable definition of dominance is that one oracle dominates another if Definition 1 holds, regardless of the players' knowledge. Considering the case where the players have no private information, Theorem 1 implies that this notion of dominance is equivalent to refinement.*

Remark 3. *Though following a different comparison rule, note that Theorem 1 is consistent with Proposition 1 in the setting of common-objective games. The distinction is that Proposition 1 concerns the best (i.e., most preferred) equilibrium outcome, whereas Theorem 1 deals with the entire set of equilibrium outcomes induced by the oracles.*

The proof of Theorem 1 shows that if Oracle 1 is not individually more informative than Oracle 2, then Oracle 1 does not dominate Oracle 2. The constructed game (in the proof of Theorem 1) can be slightly modified by aggregating the players' payoffs into a common objective, yielding a common-objective game in which there exists an equilibrium distribution induced by Oracle 2 that cannot be induced by Oracle 1.

4.2 Common knowledge components

Theorem 1 characterizes dominance (under deterministic signaling functions) using the notion of IMI, and Example 4 shows that if Oracle 1 is IMI than Oracle 2 it does not imply that F_1 refines F_2 . Nevertheless, Example 4 does show that F_1 refines F_2 in every information set of player 1. That is, given an element of player 1's partition, F_1 refines F_2 . This raises the general question of whether the notion of IMI leads, in some way, to a refinement of partitions while taking into account the players' private information.

To study this aspect in the context of games, rather than in decision problems, we first need to define the notion of a "Common Knowledge Component." Following Aumann (1976),

an event $E \subseteq \Omega$ is a *common knowledge component* (CKC) if E is common knowledge (among all players) given some $\omega \in E$, and there is no event $E' \subsetneq E$ which is also common knowledge given some $\omega' \in E'$. Formally, an event E is a CKC of the partitions $\Pi_1, \Pi_2, \dots, \Pi_n$ if it is an element in the meet $\bigwedge_{i=1}^n \Pi_i$, which is the finest common coarsening of all the partitions. For example, Figure 8 depicts two CKCs: $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4\}$.

With respect to players' payoffs, their sole concern is the information available within each CKC. Moreover, all possible posteriors within a given CKC are derived collectively from the players' private and public signals within that CKC. This implies that players' expected payoffs can be decomposed across CKCs. As a result, the impact of each oracle can be analyzed independently within each CKC.

Using this definition, we can now debate the general hypothesis of whether an IMI oracle also has a finer partition in every CKC. The answer for this question is no. The following example shows that even in the case of a unique CKC, the fact that Oracle 1 is IMI than Oracle 2 does not imply that F_1 refines F_2 .

Example 5. *IMI does not imply refinement in every CKC, and refinement in every CKC does not imply IMI.*

To see that IMI does not imply refinement in every CKC, consider the information structure given in Figure 10. It depicts a unique CKC that covers the entire state space, such that $\Pi_1 = \{\{\omega_1, \omega_4\}, \{\omega_2\}, \{\omega_3\}\}$, $\Pi_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3, \omega_4\}\}$, and $\Pi_3 = \{\{\omega_1\}, \{\omega_2, \omega_3\}, \{\omega_4\}\}$. One can see that there exists a unique CKC, Ω , as the finest common coarsening of all players' partitions is Ω . The oracles, however, have the following partitions: $F_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$ and $F_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$.

Both oracles can either withhold all information or fully disclose their information, thereby ensuring that all players become fully informed of the realized state. In addition, Oracle 1 can signal the partition $F'_1 = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}\}$, which provides complete information to players 1 and 2 but provides no information to player 3. Oracle 2 cannot do the same, because any information provided by Oracle 2 (other than the trivial set Ω) gives all players complete information. Thus, Oracle 1 is IMI than Oracle 2 because Oracle 1 can provide full information to all players simultaneously, whereas Oracle 2 is not IMI than Oracle 1. Note that neither of

the two partitions is finer than the other.




Figure 10: On the left, Figure (a) illustrates the information structures: $\Pi_1 = \{\{\omega_1, \omega_4\}, \{\omega_2\}, \{\omega_3\}\}$ of player 1 (blue); $\Pi_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3, \omega_4\}\}$ of player 2 (red); and $\Pi_3 = \{\{\omega_1\}, \{\omega_2, \omega_3\}, \{\omega_4\}\}$ of player 3 (black). On the right, Figure (b) portrays the information structures $F_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$ of Oracle 1 (orange) and $F_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$ of Oracle 2 (green). This illustrates a unique CKC in which neither oracle refines the other. Nevertheless, F_1 is IMI than F_2 whereas the converse is not true, because Oracle 2 cannot replicate the partition $F'_1 = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}\}$.

Another aspect of this example, which resonates with the key insight of the stochastic setting in Section 5, is that there exists a stochastic strategy τ_2 that Oracle 1 cannot imitate. Specifically, consider the stochastic strategy τ_2 given in Figure 11. One can verify that there exists no τ_1 that yields the same vectors of posteriors as the stated strategy τ_2 , and this hinges on the fact that F_1 does not refine F_2 . A broader discussion of this issue is given in Example 6 at the beginning of Section 5.

$\tau_2(s \omega)$	s_1	s_2
ω_1	1/3	2/3
ω_2	2/3	1/3
ω_3	1/3	2/3
ω_4	2/3	1/3

Figure 11: A stochastic F_2 -measurable strategy of Oracle 2.

To demonstrate that refinement in every CKC does not imply IMI, consider the following example with two players whose partitions are $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_4, \omega_5\}, \{\omega_3, \omega_6\}\}$ and $\Pi_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, \{\omega_5, \omega_6\}\}$. In this case, there are two CKCs, $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4, \omega_5, \omega_6\}$. Next, assume the two oracles have the following partitions, $F_1 = \{\{\omega_1, \omega_3, \omega_4\}, \{\omega_2, \omega_5, \omega_6\}\}$,

$F_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, \{\omega_5, \omega_6\}\}$, as illustrated in Figure 12. Observe that in every CKC, F_1 refines F_2 .

Now consider a completely revealing, deterministic strategy τ_2 that maps the three different partition elements of F_2 to three different signals: $\tau_2(s_1|\omega_1) = \tau_2(s_1|\omega_2) = 1$, $\tau_2(s_2|\omega_3) = \tau_2(s_2|\omega_4) = 1$, and $\tau_2(s_3|\omega_5) = \tau_2(s_3|\omega_6) = 1$. Can Oracle 1 produce a signaling function τ_1 such that $\Pi_i \vee \tau_1 = \Pi_i \vee \tau_2$ for every player i ?

Note that under τ_2 , neither player can distinguish ω_1 from ω_2 . Therefore, in order for τ_1 to satisfy $\Pi_i \vee \tau_1 = \Pi_i \vee \tau_2$ for every i , the strategy τ_1 must map all F_1 partition elements to the same signal. Consequently, under τ_1 , Player 1 cannot distinguish ω_4 from ω_5 , which is achievable given τ_2 . We therefore conclude that Oracle 1 is not IMI than Oracle 2, even though F_1 refines F_2 in every CKC. However, in the special case where Ω consists of a single CKC, refinement does imply IMI.




Figure 12: Refinement in every CKC does not imply IMI. Suppose $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_4, \omega_5\}, \{\omega_3, \omega_6\}\}$ and $\Pi_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, \{\omega_5, \omega_6\}\}$. There are two CKCs, $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4, \omega_5, \omega_6\}$. Consider F_1 (orange) and F_2 (teal) depicted in the figure. Despite F_1 refining F_2 in every CKC, F_1 is not individually more informative than F_2 .

4.2.1 Two-sided IMI implies equivalence in every CKC

Though we substantiated that an IMI oracle need not have a finer partition in every CKC, this does hold in case *both* oracles dominate one another, under deterministic signaling strategies. The following theorem provides this equivalence by stating that, given a specific CKC, both oracles dominate each other if and only if their partitions coincide.

Theorem 2. *Fix a unique CKC. Then, Oracle i is IMI than Oracle $-i$ for every i if and only if $F_1 = F_2$.*

By applying the result within each CKC, the theorem asserts that the partitions F_1 and F_2 are equivalent in every CKC if and only if they are mutually IMI within that CKC, given any *fixed* set of players’ partitions. This aligns with our previous observation in Example 4 that IMI with respect to *any* set of partitions implies refinement. As a result, the issue of CKCs arises naturally in the context of deterministic oracles and becomes even more significant when studying stochastic ones, as examined in Section 5.

5 Partial ordering of (stochastic) oracles

In this section we analyze dominance among oracles who can exercise general signaling strategies, not restricted to deterministic ones. The main result characterizes when one oracle dominates another in the case of a single CKC.

To achieve this result, we take the following gradual steps. In Section 5.1 we describe a two-stage game, entitled “a game of beliefs”. Given a profile p of probability distributions, the players’ expected payoffs in this game are maximized if and only if their individual beliefs match p . We use the game of beliefs to show that if an oracle dominates another, it must be able to produce the same joint posteriors as the other oracle. In Section 5.2 we consider a set-up with a unique CKC and show that Oracle 1 dominates Oracle 2 if and only if F_1 refines F_2 .

The next stages of this analysis are provided in Part II of the paper (i.e., in Lagziel et al., 2025), where we introduce the concept of *information loops* between common knowledge components (CKCs). In the absence of such loops, we show—building on the result in Section 5.2—that oracle-dominance is equivalent to partition refinement within each CKC. However, in the general case, the refinement condition alone is not sufficient. To fully characterize dominance, it must be combined with an additional condition expressed in terms of information loops.

Before we proceed with the aforementioned road map, we start with a simple example that illustrates the difference between the deterministic and the stochastic settings. In the following

two-player set-up, we show that even if Oracle 1 is IMI than Oracle 2, it does not mean that Oracle 1 can match the posteriors that Oracle 2 generates under stochastic strategies (whereas this can be achieved under deterministic strategies). This example also resonates with the key issue in Example 5, showing that IMI does not imply refinement in every CKC.

Example 6. *IMI is insufficient under stochastic oracles.*

The ordering generated by the notion of IMI need not hold when we transition to stochastic strategies. Consider, for example, the following uniformly distributed state space $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, with two players whose private information is given by the two partitions $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$ and $\Pi_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3, \omega_4\}\}$. The oracles, to differ, have the following partitions $F_1 = \{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}$ and $F_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. This information structure is illustrated in Figure 13.




Figure 13: 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 Oracle 1 (orange) and Oracle 2 (green).

First, assume that every Oracle i is restricted to a deterministic F_i -measurable strategy. Thus, every oracle can either convey no information, i.e., a constant signaling strategy, or he can reveal his partition element, thus ensuring that all players have complete information. Therefore, we can say that Oracle 1 is IMI than Oracle 2, and vice versa.

Now, consider the stochastic strategy τ_2 given in Figure 14. Given ω_1 and assuming s_2 is

realized, the posteriors of players 1 and 2 are $\mu_{\tau_2|\omega_1,s_2}^1 = (2/5, 3/5, 0, 0)$ and $\mu_{\tau_2|\omega_1,s_2}^2 = e_1 = (1, 0, 0, 0)$, respectively.¹²

$\tau_2(s \omega)$	s_1	s_2	s_3
ω_1	0	1/2	1/2
ω_2	1/4	3/4	0
ω_3	0	1/2	1/2
ω_4	1/4	3/4	0

Figure 14: A stochastic F_2 -measurable strategy of Oracle 2.

To mimic this joint posterior, there must exist a signal s_4 such that $\tau_1(s_4|\omega_1) = \alpha > 0$ and $\tau_1(s_4|\omega_2) = \frac{3}{2}\alpha$. However, τ_1 is F_1 -measurable, so $\tau_1(s_4|\omega_4) = \alpha$ and $\tau_1(s_4|\omega_3) = \frac{3}{2}\alpha$. Hence, given ω_3 and assuming s_4 is realized, we get a joint posterior of $\mu_{\tau_1|\omega_3,s_4}^1 = e_3 = (0, 0, 1, 0)$ and $\mu_{\tau_1|\omega_3,s_4}^2 = (0, 0, 3/5, 2/5)$, which does not exist in the support of τ_2 . So, although Oracle 1 is IMI than Oracle 2 under deterministic strategies, he cannot convey the same information under stochastic ones.

Note that the players' partitions form two CKCs, the first is $\{\omega_1, \omega_2\}$ and the second $\{\omega_3, \omega_4\}$. In every CKC, every oracle refines the other, so each of them can mimic the other, even under stochastic strategies, in that CKC. Yet, the example shows that one cannot extend this result to the entire state space.

This raises the question of the fundamental difference between the deterministic and stochastic settings. This issue should be addressed on two levels: within every CKC and between CKCs. Example 5 suggests that, under stochastic signaling functions, one cannot restrict the discussion to IMI alone but must require that F_1 refines F_2 within every CKC. Example 6 further complicates this problem by demonstrating that even a refinement within every CKC may not be sufficient.

The critical distinction arises from the significance of the joint profile of posteriors. The induced Bayesian game and its equilibria depend not only on the players' marginal posteriors but also on the joint profile of posteriors. In the deterministic setup, there is a *unique* public signal in every state, leading to a *unique* posterior for each player. Consequently, the IMI condition ensures that the profiles of posteriors coincide and the dominant oracle induces the

¹²We use e_i to denote the vector whose i^{th} coordinate is 1, while all other coordinates equal 0.

same Bayesian game as the other oracle. However, this is not necessarily the case in the stochastic setting, where multiple public signals can induce various marginal posteriors in each state. This poses a challenge both within and across CKCs.

The fact that every state has potentially multiple signals allows the oracles to use the same signals, with *different weights*, across various states. The basic structure of the players' partitions is not rich enough to cover all the information that the oracles can convey this way. Namely, one cannot use the players' interim partitions (i.e., given the information conveyed by the oracles), to cover all feasible profiles of posteriors, rather than compare these profiles directly, for every signaling function. Thus, one oracle can dominate another if the former can mimic every signaling function of the latter, and this necessitates refinement within CKCs, as well as a supplementary condition across CKCs.

5.1 A game of beliefs

In this section, we construct a two-stage game for every profile of posteriors p , which we refer to as a *game of beliefs*. The key property of this game is that the sum of equilibrium expected payoffs is maximized if and only if players adhere to the specified profile of beliefs p . Therefore, if one oracle can support that profile of posteriors, the only way for the other to match the players' expected payoffs in equilibrium is to also induce p . We repeatedly use this game in Section 5 to characterize dominance among oracles.

Formally, fix a profile of probability distributions $p = (p^1, \dots, p^n) \in (\Delta(\Omega))^n$, and consider the following game $G(p)$. The actions and utility of every player i are $A_i = \{\omega \in \Omega | p_\omega^i > 0\}$ and

$$u_i(a, \omega | p) = R_i(a_i, \omega | p) - \frac{2}{n-1} \sum_{j \neq i} R_j(a_j, \omega | p) \mathbf{1}_{\{\omega \in A_j\}},$$

respectively, where the function $R_i(a_i, \omega | p)$, for every player i , is defined by

$$R_i(a_i, \omega | p) = \begin{cases} -2, & \text{if } \omega \notin A_i, \\ \frac{1}{p_\omega^i}, & \text{if } a_i = \omega \in A_i, \\ 0, & \text{otherwise.} \end{cases}$$

In simple terms, every player i aims to match the realized state ω , and in any case would suffer a penalty of -2 if the realized state does not have a strictly positive probability according to p . Note that the utility function of every player i also depends on the actions of each player $j \neq i$, but R_j is independent of player i 's actions. The game yields to following result.

Proposition 2. *Consider the game $G(p)$. If p represents the players' actual beliefs, then the expected equilibrium payoff of every player is -1 . However, if there exists a player i with a belief $q^i \neq p^i$, then the aggregate expected payoff (over all players) in equilibrium is strictly below $-n$.*

The result given in Proposition 2 is rather straightforward. If p represents the players' actual beliefs then, in equilibrium, every player i chooses an action $a_i = \omega$ such that $p_\omega^i > 0$. This is the players' best option, given the information conveyed through p . One can easily verify it is indeed an equilibrium that yields an expected payoff of -1 for every player. Any other profile of beliefs would either yield a state with zero-probability according to p thus generating a strictly low payoff, or allow for the player to choose an action that secures an expected payoff above -1 (thus reducing the payoffs of all others).

We use this single-stage game $G(p)$ to construct a two-stage game which enables us to cross-validate the true signal and joint posterior that the players receive. The game is specifically defined given some strategy τ_2 of Oracle 2, to check whether Oracle 1 can indeed mimic the feasible posteriors of τ_2 .

The two-stage game is defined as follows. First, fix a strategy τ_2 of Oracle 2 and consider some signaling function τ . Assume that ω and s^0 are realized according to μ and τ , respectively. Thus, every player i maintains a posterior $\mu_{\tau|\omega, s^0}^i \in \Delta(\Omega)$. Next, every player i privately announces the perceived signal $s^i \in S$ and a posterior $p^i \in \Delta(\Omega)$ from the set of the player's feasible posteriors given the (previously fixed) signaling function τ_2 , private information Π_i and the stated signal s^i . Let $s = (s^1, s^2, \dots, s^n)$ be the profile of declared signals and denote by $p = (p^i)_{i \in N}$ the declared posteriors of all players. If s and p are not feasible profiles according to the information induced by every Π_i and τ_2 (including a mismatch between signals so that $s^i \neq s^j$ for any two players i and j), then all players receive $-M$ for some $M \gg 1$. However, if $s^1 = s^2 = \dots = s^n \in S_{\tau_2}$ and $p = (\mu_{\tau_2|\omega, s^1}^i)_{i \in N} \in \text{Post}(\tau_2)$, then all players proceed to the

second stage in which they play $G(p)$. The two-stage game $\mathbf{G}_{\tau_2}$ is illustrated in Figure 15.




Figure 15: The two-stage game $\mathbf{G}_{\tau_2}$, under any signaling strategy τ .

This two-stage game $\mathbf{G}_{\tau_2}$ is constructed such that players have to match their declared signals and posteriors between themselves because every mismatch leads to a very low expected payoff. Moreover, for the same reason, the players must also ensure that the declared signals and subsequent posteriors match a feasible profile (s, p) given their private information and signaling function τ_2 .

The following claim analyzes the two-stage game $\mathbf{G}_{\tau_2}$ given that the signaling function τ is either τ_2 or τ_1 , and assuming that the set $\text{Post}(\tau_1)$ is not a subset of $\text{Post}(\tau_2)$, i.e., assuming that $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$. It proves that under τ_2 , players can achieve a strictly higher aggregate expected payoff compared to what they can achieve in equilibrium under τ_1 .

Lemma 1. *Consider the two-stage game $\mathbf{G}_{\tau_2}$. If τ_2 is the signaling function, then there exists an equilibrium so that the aggregate expected payoff is $-n$. However, given τ_1 and assuming that $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$, then the aggregate expected payoff in equilibrium is strictly below $-n$.*

An immediate conclusion from Lemma 1 is Proposition 3, which establishes a condition for the existence of a strategy τ_2 such that $\text{NED}(G(\tau_2)) \neq \text{NED}(G(\tau_1))$ for every τ_1 . Proposition 3 states that, given a strategy τ_2 and for every τ_1 such that $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$, there exists a game in which Oracle 1 cannot dominate Oracle 2 due to its inability to match the set of equilibria induced by the latter. The proof is straightforward, given the construction of $\mathbf{G}_{\tau_2}$ and Lemma 1, and is therefore omitted. Yet, as in the proof of Theorem 1, we emphasize that the deduction follows from the fact that once the expected payoffs in equilibrium do not align between $G(\tau_1)$ and $G(\tau_2)$, then the equilibrium distributions over profiles of actions and states cannot match.

Proposition 3. *Fix τ_2 and consider the game $\mathbf{G}_{\tau_2}$. For every τ_1 satisfying $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$, the maximal aggregate expected equilibrium payoff in $\mathbf{G}_{\tau_2}(\tau_2)$ is strictly greater than in $\mathbf{G}'_{\tau_2}(\tau_1)$, which also implies that $\text{NED}(\mathbf{G}_{\tau_2}(\tau_2)) \neq \text{NED}(\mathbf{G}_{\tau_2}(\tau_1))$.*

In other words, given the game $\mathbf{G}_{\tau_2}$, a necessary condition for Oracle 1 to dominate Oracle 2 is that, for every strategy τ_2 , there exists a strategy τ_1 , such that $\text{Post}(\tau_1) \subseteq \text{Post}(\tau_2)$. Henceforth, we relate to this as *the inclusion condition*.

The next proposition proves the reverse inclusion condition, such that a necessary condition for Oracle 1 to dominate Oracle 2 is that for every strategy τ_2 , there exists a strategy τ_1 , such that $\text{Post}(\tau_2) \subseteq \text{Post}(\tau_1)$. This builds on a different game which exploits the Kullback-Leibler divergence (KLD) to elicit a unilateral and truthful revelation of individual posteriors.

Proposition 4. *Fix τ_2 . There exists a game $\mathbf{G}'_{\tau_2}$ such that for every τ_1 satisfying $\text{Post}(\tau_2) \not\subseteq \text{Post}(\tau_1)$, it follows that $\text{NED}(\mathbf{G}'_{\tau_2}(\tau_2)) \neq \text{NED}(\mathbf{G}'_{\tau_2}(\tau_1))$.*

The combination of Propositions 3 and 4 provides a key insight into the dominance of one oracle over another: the dominant oracle can *match* the set of posterior beliefs induced by the other oracle. To formalize this, we define a combined game that integrates the game of beliefs with a KLD-based game. The following Theorem 3 establishes this result.

Theorem 3. *If $F_1 \succeq_{\text{NE}} F_2$, then for every τ_2 , there exists τ_1 , such that $\text{Post}(\tau_1) = \text{Post}(\tau_2)$.*

The intuition for this result follows from the previous propositions such that the players need to align their signals and posteriors with each other, as well as to *truthfully* match them with the feasible outcomes of τ_2 . When players are unable to achieve a truthful alignment, they encounter the issue of mismatched beliefs and misaligned incentives while playing the sub-games $\mathbf{G}_{\tau_2}$ and $\mathbf{G}'_{\tau_2}$. Notice that one can reach the result of Theorem 3 even when using the weaker (previously mentioned) dominance condition which states that Oracle 1 dominates Oracle 2 if and only if for every τ_2 and game G , it follows that $\text{NED}(G(\tau_2)) \subseteq \bigcup_{\tau_1} \text{NED}(G(\tau_1))$. Yet, the general question of whether matching the set of posteriors is not only a necessary condition for dominance, but also a sufficient one, is left for future research.

Remark 4. *Recall the weaker dominance notion in the inclusive sense (see Subsection 3.2). The proof of Theorem 3 also demonstrates that if F_1 dominates F_2 in the inclusive sense, then the conclusion of this theorem holds. Specifically, there exists τ_1 such that $\text{Post}(\tau_1) = \text{Post}(\tau_2)$.*

Beyond Theorem 3, the result given in Proposition 3 also raises an immediate question about the implications of the inclusion condition on the signaling functions τ_1 and τ_2 . Namely, how does the inclusion condition translate to the oracles' strategies, which in turn reflect on the oracles' partitions? We provide an analysis of this condition in Lemma 2 below, focusing on a specific binary signaling function τ_2 . The lemma shows that the distribution of each signal of τ_1 is proportional to the distribution of some signal of τ_2 .

Formally, fix two distinct signals $\{s_1, s_2\}$ and assume that the partition $F_2 = \{A_1, A_2, \dots, A_m\}$ has m elements, as noted. Let $p_1, p_2, \dots, p_m$ be m distinct probabilities such that all ratios of two distinct numbers from the set $\mathbb{A} = \{p_j, 1 - p_j : j = 1, 2, \dots, m\}$ are pairwise different.¹³ Define the signaling function τ_2 such that

$$\tau_2(s_1|A_j) = 1 - \tau_2(s_2|A_j) = p_j, \quad \forall 1 \leq j \leq m. \quad (1)$$

Given this signaling function and assuming that the state space comprises a unique CKC,

¹³To achieve this, one can consider m distinct prime numbers $r_1 < r_2 < \dots < r_m$. Define $\mathbb{T}_0 = \mathbb{Q}$, and for every $j \geq 1$, let $\mathbb{T}_j$ be the extended field of $\mathbb{T}_{j-1}$ with $\sqrt{r_j}$. Take $p_j \in \mathbb{T}_j \setminus \mathbb{T}_{j-1}$.

Lemma 2 states that the inclusion condition implies that τ_1 is partially proportional to τ_2 , restricted to a subset of feasible signals.

Lemma 2. *Fix τ_2 given in Equation (1) and a unique CKC. If $\text{Post}(\tau_1) \subseteq \text{Post}(\tau_2)$, then for every signal $t \in \text{Supp}(\tau_1)$ there exists a signal $s \in \{s_1, s_2\}$ and a constant $c > 0$ such that $\tau_1(t|\omega) = c\tau_2(s|\omega)$ for every $\omega \in \Omega$.*

The result in Lemma 2 pertains to fundamental aspects of Bayesian inference. When the inclusion condition holds, the probability weights for each signal of τ_1 must be proportional to the weights of some signal of τ_2 ; otherwise, the posteriors would not align. The impact of this condition is rather extensive, because it implies (at least in some cases) that the partition of Oracle 1 refines that of Oracle 2. We utilize this result in the characterization of oracle dominance under a unique CKC in the following Section 5.2.

5.2 A unique CKC

In this section, we characterize oracle dominance under the assumption that Ω consists of a unique CKC. Specifically, we prove in Theorem 4 that, given a unique CKC, Oracle 1 dominates Oracle 2 if and only if F_1 refines F_2 . This is also equivalent to the condition that for every strategy τ_2 , there exists a strategy τ_1 such that the inclusion condition holds (by itself *and as an equality*), and it is also equivalent to the condition that the set of distributions over posteriors profiles are identical (namely, that for every strategy τ_2 , there exists a strategy τ_1 such that $\mu_{\tau_1} = \mu_{\tau_2}$). While this result has significant merits on its own, it also serves as a foundational building block for subsequent results in Lagziel et al. (2025) that address the partial ordering of oracles in more general probability spaces.

Theorem 4. *Assume that Ω comprises a unique common knowledge component. Then, the following are equivalent:*

- F_1 refines F_2 ;
- $F_1 \succeq_{\text{NE}} F_2$;
- For every τ_2 , there exists τ_1 , so that $\text{Post}(\tau_1) \subseteq \text{Post}(\tau_2)$;

- For every τ_2 , there exists τ_1 , so that $\text{Post}(\tau_1) = \text{Post}(\tau_2)$;
- For every τ_2 , there exists τ_1 , so that $\mu_{\tau_1} = \mu_{\tau_2}$.

Theorem 4, which builds on Lemma 2, presents an intriguing *equivalence* between partition refinements and the inclusion condition. Notably, this result applies to any information structure with a unique CKC, independent of any specific game. Furthermore, the refinement condition implies that Oracle 1 can effectively mimic any strategy of Oracle 2, allowing Oracle 1 to support the same sets of distributions on $\Omega \times A$ induced by Nash equilibria in incomplete-information games for any given τ_2 .

6 Final comments and a pointer to part II

6.1 The main results in part I

This paper provides a comprehensive analysis of the case where the signaling strategy is deterministic. In contrast, the stochastic signaling case is only partially explored, and even then, only in the specific context where there is a single CKC.

This naturally leads to the following question: what happens when there are multiple CKCs, and the oracles lack access to the players' common knowledge—in particular, the ability to distinguish between states that belong to different CKCs? In such scenarios, the informational limitations of the oracles may lead to significant complications. The implications of this gap are addressed more thoroughly in Part II of the paper.

6.2 Information loops

Revisiting Example 6, we ask: What is the fundamental reason that Oracles 1 and 2 dominate each other—hence are equivalent—when attention is restricted to deterministic signaling functions, yet fail to be equivalent once stochastic signaling strategies are allowed?

The key lies in the structure of the information that oracles can induce. Under deterministic signaling, both oracles generate the same joint posterior beliefs, accounting for redundancies due to private information, and thus satisfy the *Individually More Informative* (IMI) condition

in both directions. As a result, they are equivalent with respect to the set of equilibrium outcomes they can support.

However, when oracles are allowed to use stochastic signaling, the informational structure becomes richer and more nuanced. Lagziel et al. (2025) introduces the concept of *information loops*, which capture the recursive flow of information across distinct components of common knowledge. In general, an F_i -loop is a closed path among distinct CKCs, connected via the information sets of Oracle i .

While Oracles 1 and 2 may be indistinguishable under deterministic signaling (generating the same posterior beliefs and satisfying the IMI condition), they can induce *different* information loops when employing stochastic strategies. These differences affect how information is disseminated and interpreted, and in particular, influence the extent to which each oracle is constrained in its ability to shape the players' beliefs.

As a result, the sets of achievable equilibria may diverge. The failure of equivalence under stochastic signaling thus stems from the oracles' ability to influence the game's informational dynamics in fundamentally different ways—differences that deterministic signaling cannot capture. These structural distinctions are precisely what the notion of information loops is designed to formalize.

References

- Aumann, R. J. (1976). Agreeing to Disagree. *The Annals of Statistics* 4(6), 1236–1239. Publisher: Institute of Mathematical Statistics.
- Bergemann, D. and S. Morris (2016, May). Bayes correlated equilibrium and the comparison of information structures in games: Bayes correlated equilibrium. *Theoretical Economics* 11(2), 487–522.
- Bizzotto, J., J. Rüdiger, and A. Vigier (2021, February). Dynamic Persuasion with Outside Information. *American Economic Journal: Microeconomics* 13(1), 179–194.

- Blackwell, D. (1951). Comparison of Experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* 2, 93–102.
- Blackwell, D. (1953). Equivalent Comparisons of Experiments. *Annals of Mathematical Statistics* 24(2), 265–272. Publisher: Institute of Mathematical Statistics.
- Blackwell, D. and M. A. Girshick (1954). *Theory of games and statistical decisions*. Theory of games and statistical decisions. Oxford, England: Wiley. Pages: xi, 355.
- Brooks, B., A. Frankel, and E. Kamenica (2024, September). Comparisons of Signals. *American Economic Review* 114(9), 2981–3006.
- Che, Y.-K. and J. Hörner (2018, May). Recommender Systems as Mechanisms for Social Learning*. *The Quarterly Journal of Economics* 133(2), 871–925.
- Ely, J. C. (2017, January). Beeps. *American Economic Review* 107(1), 31–53.
- Ely, J. C. and M. Szydlowski (2020, February). Moving the Goalposts. *Journal of Political Economy* 128(2), 468–506. Publisher: The University of Chicago Press.
- Forges, F. (1993, November). Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision* 35(3), 277–310.
- Ganglmair, B. and E. Tarantino (2014). Conversation with secrets. *The RAND Journal of Economics* 45(2), 273–302.
- Gossner, O. (2000, January). Comparison of Information Structures. *Games and Economic Behavior* 30(1), 44–63.
- Hörner, J. and A. Skrzypacz (2016, December). Selling Information. *Journal of Political Economy* 124(6), 1515–1562. Publisher: The University of Chicago Press.
- Kamenica, E. (2019, August). Bayesian Persuasion and Information Design. *Annual Review of Economics* 11(1), 249–272.

- Kamenica, E. and M. Gentzkow (2011, October). Bayesian Persuasion. *American Economic Review* 101(6), 2590–2615.
- Lagziel, D., E. Lehrer, and T. Wang (2025, May). Comparison of Oracles: Part II.
- Lehrer, E., D. Rosenberg, and E. Shmaya (2010, March). Signaling and mediation in games with common interests. *Games and Economic Behavior* 68(2), 670–682.
- Lehrer, E., D. Rosenberg, and E. Shmaya (2013, September). Garbling of signals and outcome equivalence. *Games and Economic Behavior* 81, 179–191.
- Mezzetti, C., L. Renou, T. Tomala, and W. Zhao (2022, January). Contracting Over Persistent Information.
- Peski, M. (2008, March). Comparison of information structures in zero-sum games. *Games and Economic Behavior* 62(2), 732–735.
- Renault, J., E. Solan, and N. Vieille (2013, March). Dynamic sender–receiver games. *Journal of Economic Theory* 148(2), 502–534.
- Renault, J., E. Solan, and N. Vieille (2017, July). Optimal dynamic information provision. *Games and Economic Behavior* 104, 329–349.

A Appendices

A.1 Proof of Proposition 1

Proof. Necessity. Suppose, by way of contradiction, that there exists a player, say player i , such that the combined information of F_1 and Π_i does not refine that of F_2 and Π_i . Then there exists an information set of Π_i on which F_1 does not refine F_2 . By Blackwell (1953), this implies that there is a decision problem defined on this information set in which F_2 induces a higher expected payoff than F_1 .

Now consider a common objective game in which all players except player i are dummies (i.e., have only one available action). Suppose that payoffs are zero outside this information set

and coincide with player i 's payoff within it. In this game, the highest equilibrium expected payoff induced by F_2 is strictly greater than that induced by F_1 , contradicting the assumption.

Sufficiency. Assume that for every player i , the combined information of F_1 and Π_i refines that of F_2 and Π_i . Fix a CKC. We first show that in any common objective game, confined to this CKC, and for every partition F , the highest equilibrium payoff is achieved when F is fully revealed. In fact, we prove a stronger statement.

Claim 1. *Let τ be a signaling function measurable with respect to F . Then the highest equilibrium payoff under τ is at least as high as the highest equilibrium payoff under any garbling of τ ,¹⁴ denoted τM .*

Suppose that the experiment τ uses signals in the set S , while τM uses signals in the set T . Let $(\sigma_i)_{i \in N}$ be the equilibrium profile that maximizes the players' payoff, using signals produced by τM and the private information available to the players. Finally, let $M = (m_{st})$ be the garbling matrix, where $m_{st} \geq 0$ for every $(s, t) \in S \times T$ and $\sum_{t \in T} m_{st} = 1$ for every $s \in S$.

Unlike the case with a single decision-maker, the players cannot use the signal generated by τ in conjunction with M to replicate the signal of τM . The reason is that M is typically stochastic, and if the players were to use M privately, they would generate independent signals, thus lacking coordination.

To prove the assertion, we construct an auxiliary signaling strategy, $\bar{\tau}$, that players can follow and generate the same distribution over pairs of state and action profiles as under τM and $(\sigma_i)_{i \in N}$. The set of signals that $\bar{\tau}$ uses is $S \times T$. Define

$$\bar{\tau}((s, t)|\omega) := m_{st}\tau(s|\omega).$$

Note that for any fixed $s \in S$, all signals of the form $(s, t) \in S \times T$ induce the same posterior—namely, the posterior that s induces under τ . Define the following strategy profile: for each player i , let

$$\bar{\sigma}_i((s, t), \pi_i) := \sigma_i(t, \pi_i),$$

¹⁴Here we refer to τ as a Blackwell experiment.

where π_i denotes the private information of player i , that is, the element of Π_i containing the realized state. In other words, when player i observes the signal (s, t) and the private information π_i , he plays according to $\sigma_i(t, \pi_i)$. The signaling strategy $\bar{\tau}$ serves to coordinate the players regarding the outcome of the garbling.

The profile $(\bar{\sigma}_i)_{i \in N}$, when used in conjunction with the signal generated by $\bar{\tau}$, induces the same distribution over states and action profiles as the original strategy profile $(\sigma_i)_{i \in N}$ under the signal generated by τM . Consequently, it yields the same expected payoffs.

The profile $(\bar{\sigma}_i)_{i \in N}$ may not constitute an equilibrium, however. In that case, a sequence of pure-strategy, payoff-improving deviations by individual players benefits all players and eventually (after finitely many such deviations) leads to an equilibrium induced by $\bar{\tau}$. The resulting payoff is at least as high as the one generated by τM and the profile $(\bar{\sigma}_i)_{i \in N}$.

Since, for a fixed $s \in S$, all signals of the form (s, t) induce the same posterior, we can assume that for every player i and private information π_i , the actions $\bar{\sigma}_i((s, t), \pi_i)$ are identical across all $t \in T$. It follows that the strategies $\bar{\sigma}_i$ can be equivalently defined on the signal set S associated with τ .

We conclude that there exists an equilibrium under τ that yields a payoff at least as high as that generated by the profile $(\bar{\sigma}_i)_{i \in N}$. This completes the proof of Claim 1.

Observe that any F -measurable signaling function is a garble of the full revelation of F . Thus, the highest equilibrium payoff induced by $F_i, i = 1, 2$ is when it is fully revealed. Finally, since for every player i , the join of F_1 and Π_i refines the join of F_2 and Π_i , any equilibrium strategy that is measurable with respect to the latter is also measurable with respect to the former.¹⁵ If these strategies do not constitute an equilibrium under F_1 , then a process of sequential improvement—where players unilaterally deviate one at a time—leads to an equilibrium that yields a higher payoff. This concludes the proof. □

¹⁵We cannot reuse Claim 1 here because there is no common garbling for all players: each has its own garbling matrix.

A.2 Proof of Theorem 1

Proof. One derivation is straightforward. Assume that $F_1 \succeq_{(\mu^i)_i} F_2$. For every τ_2 , take τ_1 such that $\Pi_i \vee \tau_1 = \Pi_i \vee \tau_2$ for every player i . Thus, we get $\text{NED}(G(\tau_1)) = \text{NED}(G(\tau_2))$ for every game G . This holds for every strategy τ_2 , so $F_1 \succeq_{\text{NE}} F_2$ as needed.

To establish the converse derivation of the theorem, we assume that Oracle 1 is not individually more informative than Oracle 2, and prove that Oracle 1 does not dominate Oracle 2. Fix a strategy τ_2 , so that for every τ_1 , there exists a player i such that $\Pi_i \vee \tau_1 \neq \Pi_i \vee \tau_2$. Consider such τ_1 , and with no loss of generality, assume that $\Pi_1 \vee \tau_1 \neq \Pi_1 \vee \tau_2$. Denote $\Pi_1 \vee \tau_2 = \{B_1, \dots, B_k\}$ where $B_j = \{\omega_1^j, \dots, \omega_{|B_j|}^j\} \subseteq \Omega$ for every $1 \leq j \leq k$.

Consider the following decision problem. Define P_{B_j} to be the set of all permutations of B_j , so that every element $p \in P_{B_j}$ is a function $p : B_j \rightarrow \{1, 2, \dots, |B_j|\}$ where $p(\omega_l^j)$ is the location of ω_l^j according to that permutation. Let $A_1 = \bigcup_j P_{B_j}$ be the action set of player 1, so that player 1 chooses a permutation p over a partial set of Ω . Define the following utility function

$$u_1(a, \omega) = u_1(p, \omega_l^j) = \begin{cases} \frac{p(\omega_l^j)}{\mu(\omega_l^j|B_j)|B_j|}, & \text{if } p \in P_{B_j}, \\ -\frac{2^{10|\Omega|}}{\min_{\omega} \mu(\omega)}, & \text{if } p \notin P_{B_j}, \end{cases}$$

where $\mu(\omega_l^j|B_j)$ is the probability of ω_l^j conditional on B_j . In simple terms, player i needs to match his action, i.e., a permutation, to the realized state ω_l^j . If the action of player 1 is not a permutation on the states of the realized element of the partition (generated by his private information and the information that Oracle 2 conveys), he gets an extremely low negative payoff. However, in case the action of player 1 is a permutation on the relevant block, he receives a positive payoff based on the ordinal location of the realized state according to the chosen permutation.

Let us compare the expected payoffs of player 1 given the additional information conveyed separately by the two oracles. Given the partition $\Pi_1 \vee \tau_2$ and after ω is realized, player 1 is

informed of the relevant block B_j of the partition such that $\omega \in B_j$. Thus, for every $p \in P_{B_j}$,

$$\mathbf{E}[u_1(p, \omega)|B_j] = \sum_{\omega_l^j \in B_j} \mu(\omega_l^j|B_j)u_1(p, \omega_l^j) = \sum_{\omega_l^j \in B_j} \mu(\omega_l^j|B_j) \frac{p(\omega_l^j)}{\mu(\omega_l^j|B_j)|B_j|} = \sum_{\omega_l^j \in B_j} \frac{p(\omega_l^j)}{|B_j|} = \frac{|B_j| + 1}{2}.$$

Note that the expected payoff is independent of the chosen permutation p given that $p \in P_{B_j}$.

Hence,

$$\max_p \mathbf{E}[u_1(p, \omega)|\Pi_1 \vee \tau_2] = \sum_{j=1}^k \mu(B_j) \frac{|B_j| + 1}{2}.$$

Now consider the two possible scenarios given that $\Pi_1 \vee \tau_1 \neq \Pi_1 \vee \tau_2$: either $\Pi_1 \vee \tau_1$ is a strict refinement of $\Pi_1 \vee \tau_2$, or there exists at least one block of $\Pi_1 \vee \tau_1$ that intersects two disjoint blocks of $\Pi_1 \vee \tau_2$.

Starting with the former, assume that $\Pi_1 \vee \tau_1$ is a strict refinement of $\Pi_1 \vee \tau_2$, so there exists a block B_j^* that $\Pi_1 \vee \tau_1$ splits into at least two separate blocks. Without loss of generality, assume that B_1 is such a block, and denote the two disjoint sub-blocks by $B_{1,1}$ and $B_{1,2}$, so that $B_1 = B_{1,1} \cup B_{1,2}$. Assume that for every $B_j \neq B_1$, player 1 follows the same strategy as with $\Pi_1 \vee \tau_2$ so that we can focus on the difference in expected payoffs given B_1 . Evidently,

$$\begin{aligned} \mathbf{E}[u_1(p, \omega)|B_{1,1}] &= \sum_{\omega_l^1 \in B_{1,1}} \mu(\omega_l^1|B_{1,1})u_1(p, \omega_l^1) = \sum_{\omega_l^1 \in B_{1,1}} \mu(\omega_l^1|B_{1,1}) \frac{p(\omega_l^1)}{\mu(\omega_l^1|B_1)|B_1|} \\ &= \sum_{\omega_l^1 \in B_{1,1}} \mu(\omega_l^1|B_1) \frac{\mu(B_1)}{\mu(B_{1,1})} \cdot \frac{p(\omega_l^1)}{\mu(\omega_l^1|B_1)|B_1|} \\ &= \frac{\mu(B_1)}{\mu(B_{1,1})|B_1|} \sum_{\omega_l^1 \in B_{1,1}} p(\omega_l^1). \end{aligned}$$

Note that player 1 can choose a permutation on B_1 which maximizes the sum of all states in $B_{1,1}$, i.e.,

$$\max_p \sum_{\omega_l^1 \in B_{1,1}} p(\omega_l^1) = |B_1| + |B_1| - 1 + \dots + |B_1| - |B_{1,1}| + 1 > |B_{1,1}| \frac{|B_1| + 1}{2}.$$

Thus,

$$\max_{p \in P_{B_1}} \mathbf{E}[u_1(p, \omega) | B_{1,1}] > \frac{\mu(B_1)}{\mu(B_{1,1})|B_1|} |B_{1,1}| \frac{|B_1| + 1}{2},$$

and a similar computation holds for $B_{1,2}$. Therefore,

$$\begin{aligned} \max_p \mathbf{E}[u_1(p, \omega) | \Pi_1 \vee \tau_1] &> \sum_{j=2}^k \mu(B_j) \frac{|B_j| + 1}{2} + \mu(B_{1,1}) \frac{\mu(B_1)}{\mu(B_{1,1})|B_1|} |B_{1,1}| \frac{|B_1| + 1}{2} \\ &+ \mu(B_{1,2}) \frac{\mu(B_1)}{\mu(B_{1,2})|B_1|} |B_{1,2}| \frac{|B_1| + 1}{2} \\ &= \sum_{j=2}^k \mu(B_j) \frac{|B_j| + 1}{2} + \left[\frac{|B_{1,1}|}{|B_1|} + \frac{|B_{1,2}|}{|B_1|} \right] \mu(B_1) \frac{|B_1| + 1}{2} \\ &= \sum_{j=1}^k \mu(B_j) \frac{|B_j| + 1}{2} = \max_p \mathbf{E}[u_1(p, \omega) | \Pi_1 \vee \tau_2], \end{aligned}$$

and player 1 can guarantee a strictly higher expected payoff using the information conveyed through Oracle 1 than through Oracle 2.

Next, consider the other possibility that $\Pi_1 \vee \tau_1$ is not a refinement of $\Pi_1 \vee \tau_2$. This implies that there exists at least one block of $\Pi_1 \vee \tau_1$ that intersects two disjoint blocks of $\Pi_1 \vee \tau_2$. Denote this block by B^* . For every state ω_l^j and every permutation $p \in P_{B_j}$, note that $p(\omega_l^j) \leq |B_j|$, so $u_1(p, \omega_l^j) \leq \frac{1}{\mu(\omega_l^j | B_j)}$. Hence, in the optimal case in which player 1 is completely informed of the realized state, his payoff cannot exceed $|\Omega|$. However, in case player 1 wrongfully chooses a permutation that does not match the realized block in $\Pi_1 \vee \tau_2$, his payoff is given by $-\frac{2^{10|\Omega|}}{\min_{\omega} \mu(\omega)}$.

Thus,

$$\begin{aligned} \mathbf{E}[u_1(p, \omega) | B^*] &= \sum_{\omega \in B^*} \mu(\omega | B^*) u_1(p, \omega) \\ &< \sum_{\omega \in B^*} \mu(\omega | B^*) \frac{1}{\mu(\omega | B^*)} + \min_{\omega \in B^*} \mu(\omega | B^*) \left[-\frac{2^{10|\Omega|}}{\min_{\omega} \mu(\omega)} \right] \\ &< |B^*| - \frac{2^{10|\Omega|}}{\mu(B^*)}. \end{aligned}$$

This suggests that the expected payoff of player 1 given $\Pi_1 \vee \tau_1$ is bounded from above by

$$\max_p \mathbf{E}[u_1(p, \omega) | \Pi_1 \vee \tau_1] < |\Omega| - 2^{10|\Omega|} < 0,$$

which is strictly below the expected payoff given the information transmitted through Oracle 2.

To conclude, for every player i , we can define a decision problem such that whenever $\Pi_i \vee \tau_1 \neq \Pi_i \vee \tau_2$, it follows that the expected payoff of player i given τ_2 differs from the player's expected payoff given τ_1 . Hence, there exists τ_2 which yields a unique profile of expected payoffs in equilibrium that cannot be matched by any τ_1 , thus for every τ_1 , we get $\text{NED}(\text{G}(\tau_2)) \neq \text{NED}(\text{G}(\tau_1))$, and this concludes the proof. $\square$

A.3 Proof of Theorem 2

Proof. Fix a unique CKC. One direction is trivial, so assume that F_i is IMI than F_{-i} for every $i = 1, 2$, and let us prove that $F_1 = F_2$. Assume, to the contrary, that $F_1 \neq F_2$. W.l.o.g, there exist $\omega_1 \neq \omega_2$, such that $F_1(\omega_1) = F_1(\omega_2)$ whereas $F_2(\omega_1) \neq F_2(\omega_2)$.

Consider the partition $F'_2 = \{F_2(\omega), (F_2(\omega))^c\}$. By assumption, there exists a partition F'_1 such that $\Pi_i \vee F'_1 = \Pi_i \vee F'_2$, for every player i . Denote $A = F'_1(\omega_1) \cap F_2(\omega_1)$, $B = F'_1(\omega_1) \cap (F_2(\omega_1))^c$, $C = (F'_1(\omega_1))^c \cap F_2(\omega_1)$, and $D = (F'_1(\omega_1))^c \cap (F_2(\omega_1))^c$. See Figure 16.(a).

If there exists a player i such that $\Pi_i(\omega_1) = \Pi_i(\omega_2)$, then $\omega_2 \in (F'_1 \vee \Pi_i)(\omega_1)$, while $\omega_2 \notin (F'_2 \vee \Pi_i)(\omega_2)$, which contradicts the equation $\Pi_i \vee F'_1 = \Pi_i \vee F'_2$. Thus, for every $(\omega, \omega') \in A \times B \cup A \times D \cup B \times C$ and for every player i , we conclude that $\Pi_i(\omega) \neq \Pi_i(\omega')$.

Because this is a unique CKC, every two states ω and ω' have a connected path, in the sense that there exists a finite sequence of states starting with ω and ending with ω' where every two adjacent states are in the same information set of some player. Fix such a path from ω_1 to ω_2 , and denote it by $(\omega_1, \omega_{1.1}, \omega_{1.2}, \dots, \omega_{1.l}, \omega_3, \dots, \omega_2)$ where $\{\omega_{1.t} : 1 \leq t \leq l\} \in C$ and $\omega_3 \in D$. This holds, w.l.o.g., because states in A are directly connected (through a partition element of some player) only to states in $A \cup C$, and the same holds for states in B that are directly connected only to states in $B \cup D$. Note that $\omega_{1.t} \in (F_1(\omega_1))^c$ for every t and

Illustrations of sub-partitions in the proof of Theorem 2




Figure 16: Figures (a) and (b) depict the partition F'_2 and the sub-partition F'_1 that mimics it. Figure (b) also illustrates the path between ω_1 and ω_2 , as well as the possible connections between the different sets. Figure (c) depicts the partitions F'_1 and F'_2 along with the path from ω_1 to ω_2 .

$\omega_3 \in F_2(\omega_1) \cap (F_1(\omega_1))^c$. See Figure 16.(b).

Now consider the partition $F''_1 = \{F_1(\omega_1), (F_1(\omega_1))^c\}$. By assumption, there exists a partition F''_2 such that $\Pi_i \vee F''_1 = \Pi_i \vee F''_2$, for every player i . Denote $A' = F_1(\omega_1) \cap F''_2(\omega_1)$, $B' = (F_1(\omega_1))^c \cap F''_2(\omega_1)$, $C' = (F_1(\omega_1))^c \cap (F''_2(\omega_1))^c$, and $D' = F_1(\omega_1) \cap (F''_2(\omega_1))^c$. See Figure 16.(c).

Similarly to the previous analysis, states in A' are directly connected only to states in $A' \cup C'$, and states in B' are directly connected only to states in $B' \cup D'$. In addition, note that $\omega_1 \in F_1(\omega_1) \cap F_2(\omega_1) \subseteq A'$, $\omega_{1,t} \in (F_1(\omega_1))^c \subseteq B' \cup C'$ for every t , and $\omega_3 \in F_2(\omega_1) \cap (F_1(\omega_1))^c \subseteq B'$. If $\omega_{1,1} \in B'$, we can make a direct connection between A' and B' , which yields a contradiction. Thus, $\omega_{1,1} \in C'$, and the sequence $(\omega_{1,1}, \omega_{1,2}, \dots, \omega_{1,l_1}, \omega_3)$ which starts in C' and ends in B' has at least one direct connection between B' and C' . A contradiction, as well. Thus, for every

$\omega_1 \neq \omega_2$, we conclude that $F_1(\omega_1) = F_1(\omega_2)$ if and only if $F_2(\omega_1) = F_2(\omega_2)$, and the result follows. $\square$

A.4 Proof of Proposition 2

Proof. For every player i , we can focus our analysis on the function R_i . Assuming that player i 's belief is q^i , we get

$$\max_{a_i \in A_i} \mathbb{E}_{q^i}[R_i(a_i, \omega|p)] = \max_{a_i \in A_i} \left[\sum_{\omega \in \Omega} q_\omega^i R_i(a_i, \omega|p) \right] = \max_{a_i \in A_i} \left[\sum_{\omega \in A_i} \frac{q_\omega^i}{p_{a_i}^i} \mathbf{1}_{\{\omega=a_i\}} \right] - 2 \sum_{\omega \notin A_i} q_\omega^i.$$

The second term is independent of a_i , so player i maximizes only the first one. If $p^i = q^i$ for every player i , then

$$\max_{a_i \in A_i} \left[\sum_{\omega \in A_i} \frac{q_\omega^i}{p_{a_i}^i} \mathbf{1}_{\{\omega=a_i\}} \right] - 2 \sum_{\omega \notin A_i} q_\omega^i = \max_{a_i \in A_i} \frac{q_{a_i}^i}{p_{a_i}^i} - 2 \sum_{\omega \notin A_i} 0 = 1,$$

independently of the chosen action $a_i \in A_i$. Therefore,

$$\max_{a_i \in A_i} \mathbb{E}_{q^i}[u_i(a, \omega|p)] = 1 - \frac{2}{n-1} \sum_{j \neq i} 1 = -1,$$

as stated.

Moving on to the second part of the proposition, assume that there exists a player i whose actual belief is $q^i \neq p^i$. The proof is now divided into two parts: either q^i is supported on a subset of $\text{Supp}(p^i)$, namely $\text{Supp}(q^i) \subseteq \text{Supp}(p^i)$, or not.

Starting with the former, assume that $\text{Supp}(q^i) \subseteq \text{Supp}(p^i)$. Evidently,

$$\max_{a_i \in A_i} \mathbb{E}_{q^i}[R_i(a_i, \omega|p)] = \max_{a_i \in A_i} \frac{q_{a_i}^i}{p_{a_i}^i} > 1.$$

Denote $\max_{a_i \in A_i} \mathbb{E}_{q^i}[R_i(a_i, \omega|p)] = 1 + c$. Assuming that the beliefs of all other players align

with p , the expected equilibrium payoffs of player i and of every other player $j \neq i$ are

$$\begin{aligned}\mathbb{E}_{q^i}[u_i(a_i, \omega|p)] &= 1 + c - \frac{2}{n-1}(n-1) = -1 + c, \\ \mathbb{E}_{p^j}[u_j(a_j, \omega|p)] &= 1 - \frac{2}{n-1}(n-1+c) = -1 - \frac{2c}{n-1},\end{aligned}$$

respectively. Thus, the aggregate expected payoff in equilibrium is

$$(-1 + c) + (n-1) \left[-1 - \frac{2c}{n-1} \right] = -n - c < -n,$$

as stated. Note that we get a similar result for every additional player j whose belief is $q^j \neq p^j$.

Next, assume that there exists a player i with a belief q^i such that $\text{Supp}(q^i) \not\subseteq \text{Supp}(p^i)$. If $\text{Supp}(q^i) \cap \text{Supp}(p^i) = \emptyset$, then the player's expected payoff is

$$\mathbb{E}_{q^i}[u_i(a_i, \omega|p)] = -2 - \frac{2}{n-1}(n-1) = -4.$$

For players other than player i , since $\mathbf{1}_{\{\omega \in A_i\}} = 0$, it follows that their expected payoff is

$$\mathbb{E}_{p^j}[u_j(a_j, \omega|p)] = 1 - \frac{2}{n-1}(n-2).$$

The aggregate expected payoff over all players is $-4 + (n-1) \left[1 - \frac{2}{n-1}(n-2) \right] = -n - 1$, as needed.

If $\text{Supp}(q^i) \cap \text{Supp}(p^i) \neq \emptyset$, denote $q_0 = \sum_{\omega \notin A_i} q_\omega^i \in (0, 1)$ and $r_\omega^i = \frac{q_\omega^i}{1-q_0}$, for every $\omega \in A_i$. Thus, $\sum_{\omega \in A_i} r_\omega^i = 1$, and we get

$$\max_{a_i \in A_i} \sum_{\omega \in A_i} \frac{r_\omega^i}{p_{a_i}^i} \mathbf{1}_{\{a_i = \omega\}} \geq 1,$$

which implies that

$$d := \max_{a_i \in A_i} \sum_{\omega \in A_i} \frac{q_\omega^i}{p_{a_i}^i} \mathbf{1}_{\{a_i = \omega\}} = \max_{a_i \in A_i} \sum_{\omega \in A_i} \frac{[1-q_0]r_\omega^i}{p_{a_i}^i} \mathbf{1}_{\{a_i = \omega\}} \geq 1 - q_0.$$

Thus, the expected payoff of player i , assuming that $q^j = p^j$ for every other player $j \neq i$, is

$$\begin{aligned} \max_{a_i \in A_i} \mathbb{E}_{q^i}[u_i(a_i, \omega|p)] &= \max_{a_i \in A_i} \left[\sum_{\omega \in A_i} \frac{q_\omega^i}{p_{a_i}^i} \mathbf{1}_{\{a_i=\omega\}} \right] - 2 \sum_{\omega \notin A_i} q_\omega^i - \frac{2}{n-1} \sum_{j \neq i} 1 \\ &= \max_{a_i \in A_i} \left[\sum_{\omega \in A_i} \frac{q_\omega^i}{p_{a_i}^i} \mathbf{1}_{\{a_i=\omega\}} \right] - 2q_0 - 2 = d - 2q_0 - 2, \end{aligned}$$

and the expected equilibrium payoff of every other player $j \neq i$ is

$$\mathbb{E}_{p^j}[u_j(a_j, \omega|p)] = 1 - \frac{2}{n-1}(n-2+d),$$

Aggregating over all players,

$$\begin{aligned} \sum_j \mathbb{E}_{p^j}[u_j(a_j, \omega|p)] &= d - 2q_0 - 2 + (n-1) \left[1 - \frac{2}{n-1}(n-2+d) \right] \\ &= -n - q_0 + (1 - q_0 - d) \\ &\leq -n - q_0 < -n, \end{aligned}$$

where the two inequalities follow from $d \geq 1 - q_0$ and $q_0 > 0$, as stated above. Again, we get a similar result for every additional player j whose belief is $\text{Supp}(q^j) \not\subseteq \text{Supp}(p^j)$, and the statement holds. $\square$

A.5 Proof of Lemma 1

Proof. We start by analyzing the game given that the signaling function is τ_2 . Consider the profiles $s = (s^1, s^2, \dots, s^n)$ and $p = (p^i)_{i \in N}$, so that all players declare the true public signal $s^i = s^j$ for every two players i and j , and $p^i = \mu_{\tau_2|\omega, s^i}^i$ is the true posterior of every player i . In the second-stage sub-game, as stated in Proposition 2, every player receives a payoff of -1 and the aggregate expected payoff in the two-stage game $\mathbf{G}_{\tau_2}$ is $-n$. Let us prove that this is indeed an equilibrium.

The negative payoff $-M$ ensures that a unilateral deviation to a different signal is sub-optimal, so we need only to consider the case in which some player i deviates to a posterior

$p^i \neq \mu_{\tau_2|\omega, s^i}^i$. Notice that, given an element in Π_i and for every signal $s^i \in S_{\tau_2}$, there exists a unique feasible posterior on Π_i . Thus, there are only two possible deviations concerning p^i : either the updated profile p is no longer feasible and again all players receive a payoff of $-M$, or p is feasible, but p^i is supported on a different partition element whose probability is zero given player i 's actual partition element. Due to the negative expected payoff of $-M$ in the former case, we need only to consider the latter possibility. If player i declares a zero-probability belief (relative to the true posterior), then the proof of Proposition 2 shows that the player's expected payoff in the second stage is -2 . Thus, we conclude that a truthful revelation of all information comprises an equilibrium, and the aggregate expected payoff given this equilibrium is $-n$.

Next, consider the signaling function τ_1 so that $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$, and fix any equilibrium profile. Evidently, the players must coordinate on some feasible combination of s and p according to τ_2 , otherwise they all get $-M$. However, with some positive probability, the declared posterior p^i of some player i mismatches the realized one $\mu_{\tau_1|\omega, s^i}^i$. In that case, Proposition 2 shows that the aggregate expected payoff is strictly below $-n$. So, the expected aggregate payoff in the two-stage game $\mathbf{G}_{\tau_2}$, given the stated strategy τ_1 , is also strictly below $-n$, as needed. $\square$

A.6 Proof of Proposition 4

Proof. Fix τ_2 and let $\text{Post}^i(\tau_2)$ be the set of feasible posterior beliefs of player i under τ_2 . Define the game $\mathbf{G}'_{\tau_2}$ as follows. The set of player i 's actions is $A_i = \text{Post}^i(\tau_2)$. His payoff function is $u_i(p^i, \omega) = \lim_{\epsilon \rightarrow 0^+} \log(p_\omega^i + \epsilon)$. For every player, the game is a single-person decision problem in which the objective of a player is to choose a belief in $\text{Post}^i(\tau_2)$ that maximizes his expected payoff, given his actual belief q^i , which may be different from p^i .

Claim 1. *If the actual belief is $q^i \in \text{Post}^i(\tau_2)$, then the optimal strategy for player i is $p^i = q^i$. Any $p^i \in \text{Post}^i(\tau_2)$ that is different from q^i would yield player i a strictly lower payoff.*

To prove this claim, first observe that it is not optimal to choose a p^i where $\text{Supp}(q^i) \not\subseteq \text{Supp}(p^i)$. Otherwise, there exists $\omega \in \text{Supp}(q^i) \setminus \text{Supp}(p^i)$, such that with a positive probability q_ω^i , player i would receive a payoff that tends to $-\infty$.

Next, we show that among those p^i that share the same support as q^i , the unique optimal

choice is $p^i = q^i$. To see this, note that

$$\sum_{\omega \in \text{Supp}(q^i)} q_\omega^i \log(p_\omega^i) = \sum_{\omega \in \text{Supp}(q^i)} q_\omega^i \log(q_\omega^i) - D_{\text{KL}}(q^i \| p^i),$$

where $D_{\text{KL}}(q^i \| p^i)$ is the Kullback-Leibler divergence. Since $D_{\text{KL}}(q^i \| p^i)$ is uniquely minimized when $p^i = q^i$, it follows that player i 's expected payoff is uniquely maximized when $p^i = q^i$.

Finally, we show that it is not optimal to choose p^i where $\text{Supp}(q^i) \subsetneq \text{Supp}(p^i)$. Consider such a p^i . Since $\sum_{\omega \in \text{Supp}(q^i)} p_\omega^i < 1$, we can allocate the remaining probability mass to states in $\text{Supp}(q^i)$ to obtain a probability distribution $\hat{p}^i$ where $\text{Supp}(\hat{p}^i) = \text{Supp}(q^i)$ and $\hat{p}_\omega^i > p_\omega^i$ for every $\omega \in \text{Supp}(q^i)$. Hence,

$$\sum_{\omega \in \text{Supp}(q^i)} q_\omega^i \log(q_\omega^i) \geq \sum_{\omega \in \text{Supp}(q^i)} q_\omega^i \log(\hat{p}_\omega^i) > \sum_{\omega \in \text{Supp}(q^i)} q_\omega^i \log(p_\omega^i),$$

where the first inequality follows from the fact that q^i is the unique optimal choice among probability distributions that share the same support, and the second inequality follows from $\hat{p}_\omega^i > p_\omega^i$ for every $\omega \in \text{Supp}(q^i)$. This concludes the claim.

It follows from Claim 1 that under τ_2 , the set of posterior belief profiles in $\text{Post}(\tau_2)$ are all chosen with positive probability in the equilibria of the game $\mathbf{G}'_{\tau_2}(\tau_2)$. On the other hand, for every strategy τ_1 satisfying $\text{Post}(\tau_2) \not\subseteq \text{Post}(\tau_1)$, there exists a posterior belief profile $p \in \text{Post}(\tau_2) \setminus \text{Post}(\tau_1)$, that is chosen with zero probability in every equilibrium of the game $\mathbf{G}'_{\tau_2}(\tau_1)$. Thus, for every τ_1 that satisfies $\text{Post}(\tau_2) \not\subseteq \text{Post}(\tau_1)$, we conclude that $\text{NED}(\mathbf{G}'_{\tau_2}(\tau_2)) \neq \text{NED}(\mathbf{G}'_{\tau_2}(\tau_1))$. $\square$

A.7 Proof of Theorem 3

Proof. Fix τ_2 , and consider the games $\mathbf{G}_{\tau_2}$ and $\mathbf{G}'_{\tau_2}$, as defined above, where the sets of actions for each player in these games are disjoint. Define the game $\mathbf{G}$ as the one in which $\mathbf{G}_{\tau_2}$ and $\mathbf{G}'_{\tau_2}$ are played with equal probability, i.e., with probability 1/2 each.

If $\text{Post}(\tau_1) \neq \text{Post}(\tau_2)$, then either there exists a posterior profile $p \in \text{Post}(\tau_1) \setminus \text{Post}(\tau_2)$, or there exists a posterior profile $p \in \text{Post}(\tau_2) \setminus \text{Post}(\tau_1)$. Following Proposition 3 and 4,

in each of the mentioned sub-games, it follows that $\text{NED}(G(\tau_2)) \neq \text{NED}(G(\tau_1))$ where $G \in \{\mathbf{G}_{\tau_2}, \mathbf{G}'_{\tau_2}\}$. Thus, if no τ_1 satisfies $\text{Post}(\tau_1) = \text{Post}(\tau_2)$, there exists a game G and τ_2 , such that $\text{NED}(G(\tau_2)) \neq \text{NED}(G(\tau_1))$ for every τ_1 , which contradicts the dominance assumption. $\square$

A.8 Proof of Lemma 2

Proof. Assume, to the contrary, there exists a signal $t \in \text{Supp}(\tau_1)$ such that for every signal $s_i \in \{s_1, s_2\}$, there exist two states $\omega_1, \omega^* \in \Omega$ such that

$$\frac{\tau_1(t|\omega_1)}{\tau_2(s_i|\omega_1)} \neq \frac{\tau_1(t|\omega^*)}{\tau_2(s_i|\omega^*)}. \quad (2)$$

Note that $\tau_2(s_i|\omega) > 0$ for every s_i and ω , so the fractions are well defined. In addition, it must be that either $\tau_1(t|\omega_1) > 0$ or $\tau_1(t|\omega^*) > 0$, so assume that $\tau_1(t|\omega_1) > 0$. Because ω_1 and ω^* are in the same CKC, there exists a finite sequence $(\omega_1, \omega_2, \omega_3, \dots, \omega^*)$ such that every two adjacent states are in the same partition element for some player. Assume, w.l.o.g., that $\{\omega_1, \omega_2\}$ and $\{\omega_2, \omega_3\}$ are in the same partition elements of players l_1 and l_2 respectively. Using the definition of τ_2 , it follows that in every posterior $(\mu_{\tau_2|\omega, s_i}^l)_{l \in N} \in \text{Post}(\tau_2)$, the coordinates relating to $\Pi_l(\omega)$ are strictly positive (for every player l and every signal s_i). Thus, for every state ω and signal s_i ,

$$\mu_{\tau_2|\omega, s_i}^{l_1}(\omega_1) > 0 \Leftrightarrow \mu_{\tau_2|\omega, s_i}^{l_1}(\omega_2) > 0,$$

and

$$\mu_{\tau_2|\omega, s_i}^{l_2}(\omega_2) > 0 \Leftrightarrow \mu_{\tau_2|\omega, s_i}^{l_2}(\omega_3) > 0.$$

Take a posterior $(\mu_{\tau_1|\omega, t}^l)_{l \in N}$ such that $\mu_{\tau_1|\omega, t}^{l_1}(\omega_1) > 0$. Because $\text{Post}(\tau_1) \subseteq \text{Post}(\tau_2)$, it follows that $\mu_{\tau_1|\omega, t}^{l_1}(\omega_2) > 0$, hence $\tau_1(t|\omega_2) > 0$. The fact that $\tau_1(t|\omega_2) > 0$ implies that $\mu_{\tau_1|\omega_2, t}^{l_2}(\omega_2) > 0$, and so $\mu_{\tau_1|\omega_2, t}^{l_2}(\omega_3) > 0$. We thus conclude that $\tau_1(t|\omega_3) > 0$. Continuing inductively, it follows that $\tau_1(t|\omega) > 0$ for every $\omega \in \{\omega_1, \omega_2, \dots, \omega^*\}$.

According to the definition of τ_2 and using Bayes' rule, for every signal s_i and for every

posterior where $\mu_{\tau_2|\omega'',s_i}^l(\omega) > 0$, which implies that $\omega \in \Pi_l(\omega'')$, we know that

$$\mu_{\tau_2|\omega'',s_i}^l(\omega) = \frac{\mu_{\tau_2}^l(\omega'', s_i|\omega)\mu(\omega)}{\mu_{\tau_2}^l(\omega'', s_i)} = \frac{\tau_2(s_i|\omega)\mu(\omega)}{|\Pi_l(\omega'')|\mu_{\tau_2}^l(\omega'', s_i)}.$$

Thus, for every $\omega' \in \Pi_l(\omega)$, we get

$$\frac{\mu_{\tau_2|\omega'',s_i}^l(\omega)}{\mu(\omega)} = \frac{\tau_2(s_i|\omega)}{\tau_2(s_i|\omega')} \cdot \frac{\mu_{\tau_2|\omega'',s_i}^l(\omega')}{\mu(\omega')}.$$

Note that $\frac{\tau_2(s_i|\omega)}{\tau_2(s_i|\omega')} = 1$ if and only if $F_2(\omega) = F_2(\omega')$, and otherwise, the ratio $\frac{\tau_2(s_i|\omega)}{\tau_2(s_i|\omega')}$ is given by $c \in \{\frac{x}{y} : x, y \in \mathbb{A}\}$. Thus, for every such s_i where $\mu_{\tau_2|\omega'',s_i}^l(\omega) \cdot \mu_{\tau_2|\omega'',s_i}^l(\omega') > 0$, there exists a unique $c \in \{\frac{x}{y} : x, y \in \mathbb{A}\} \cup \{1\}$ such that

$$\frac{\mu_{\tau_2|\omega'',s_i}^l(\omega)}{\mu(\omega)} = c \cdot \frac{\mu_{\tau_2|\omega'',s_i}^l(\omega')}{\mu(\omega')}.$$

In case $c = 1$, then the last equation holds for every signal s_i because $\tau_2(s_i|\omega) = \tau_2(s_i|\omega')$ if and only if $\omega' \in F_2(\omega)$.

By the inclusion criterion, for every posterior $(\mu_{\tau_1|\omega_2,t}^l)_{l \in N}$ generated by τ_1 , there exists a posterior $(\mu_{\tau_2|\omega'',s_i}^l)_{l \in N}$ generated by τ_2 , such that the two are identical. We thus conclude that

$$\frac{\mu_{\tau_1|\omega_2,t}^{l_1}(\omega_1)}{\mu(\omega_1)} = \frac{\mu_{\tau_2|\omega'',s_i}^{l_1}(\omega_1)}{\mu(\omega_1)} = c_1 \cdot \frac{\mu_{\tau_2|\omega'',s_i}^{l_1}(\omega_2)}{\mu(\omega_2)} = c_1 \cdot \frac{\mu_{\tau_1|\omega_2,t}^{l_1}(\omega_2)}{\mu(\omega_2)},$$

and

$$\frac{\mu_{\tau_1|\omega_2,t}^{l_2}(\omega_2)}{\mu(\omega_2)} = \frac{\mu_{\tau_2|\omega'',s_i}^{l_2}(\omega_2)}{\mu(\omega_2)} = c_2 \cdot \frac{\mu_{\tau_2|\omega'',s_i}^{l_2}(\omega_3)}{\mu(\omega_3)} = c_2 \cdot \frac{\mu_{\tau_1|\omega_2,t}^{l_2}(\omega_3)}{\mu(\omega_3)},$$

as well. Using Bayes' rule, the last two equations are equivalent to

$$\begin{aligned} \tau_2(s_i|\omega_1) &= c_1 \cdot \tau_2(s_i|\omega_2) = c_1 \cdot c_2 \cdot \tau_2(s_i|\omega_3), \\ \tau_1(t|\omega_1) &= c_1 \cdot \tau_1(t|\omega_2) = c_1 \cdot c_2 \cdot \tau_1(t|\omega_3). \end{aligned} \tag{3}$$

Note that these equations hold for every s_i in case $c_1 = c_2 = 1$, and otherwise hold for a specific

signal, which could be taken as s_1 without loss of generality.

One can continue inductively along the sequence $(\omega_1, \omega_2, \omega_3, \dots, \omega^*)$ to get

$$\tau_2(s_i|\omega_2) = c_2 \cdot \tau_2(s_i|\omega_3) = c_2 \cdot c_3 \cdot \tau_2(s_i|\omega_4), \quad (4)$$

$$\tau_1(t|\omega_2) = c_2 \cdot \tau_1(t|\omega_3) = c_2 \cdot c_3 \cdot \tau_1(t|\omega_4),$$

and the first equality in Equation (4) coincides with the second equality in Equation (3).

Namely, Equations (3) and (4) either hold for every signal s_i , or hold for the same signal s_1 .

Repeatedly following the same procedure, we get that

$$\tau_2(s_i|\omega_1) = c_1 \cdot \tau_2(s_i|\omega_2) = \dots = [\prod_{k \geq 1} c_k] \cdot \tau_2(s_i|\omega^*), \quad (5)$$

$$\tau_1(t|\omega_1) = c_1 \cdot \tau_1(t|\omega_2) = \dots = [\prod_{k \geq 1} c_k] \cdot \tau_1(t|\omega^*). \quad (6)$$

Dividing Equation (6) by Equation (5), we get $\frac{\tau_1(t|\omega_1)}{\tau_2(s_i|\omega_1)} = \frac{\tau_1(t|\omega^*)}{\tau_2(s_i|\omega^*)}$, which contradicts (2), as needed. $\square$

A.9 Proof of Theorem 4

Proof. Proving that the first condition yields the second which, in turn, yields the third, is immediate. Assume that F_1 refines F_2 . Then, for every τ_2 , there exists τ_1 such that $\tau_1 = \tau_2$. It thus follows that Oracle 1 dominates Oracle 2. Next, assume that there exists τ_2 such that for every τ_1 , it follows that $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$. According to Proposition 3, Oracle 1 does not dominate Oracle 2. Now, let us prove that the third condition yields the first, that is: if F_1 does not refine F_2 , then there exists τ_2 such that for every τ_1 , it follows that $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$.

If F_1 does not refine F_2 , there exists ω_0 and ω^* , so that $F_1(\omega_0) = F_1(\omega^*)$ and $F_2(\omega_0) \neq F_2(\omega^*)$. Consider the signaling function τ_2 defined in (1) and take any strategy τ_1 . Assume, to the contrary that $\text{Post}(\tau_1) \subseteq \text{Post}(\tau_2)$. According to Lemma 2, for every signal $t \in \text{Supp}(\tau_1)$ there exists a signal $s_i \in \text{Supp}(\tau_2)$ and a constant $c > 0$ such that $\tau_1(t|\omega) = c\tau_2(s_i|\omega)$ for every ω . In addition, the measurability condition of τ_1 imply that $\tau_1(t|\omega_0) = \tau_1(t|\omega^*)$ for every signal t . Thus, $\tau_2(s_i|\omega_0) = \tau_2(s_i|\omega^*)$ and this contradicts the definition of τ_2 . This establishes the

equivalence between the first three conditions.

Now, notice that the first (refinement) condition implies the equivalence of distributions over posteriors profiles (fifth condition), because Oracle 1 can exercise any strategy of Oracle 2. The fifth condition in turn implies the forth condition (so that the set of joint posterior profiles match), which implies the third condition, thus concluding the proof. $\square$