

Comparison of Oracles*

David Lagziel[†]

Ehud Lehrer[‡]

Tao Wang[§]

Ben-Gurion University

Durham University

CUEB

May 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 by developing 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 relevant 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 these cases, external observers convey partial information to players engaged in strategic interactions.

This paper examines incomplete-information games in which an external oracle publicly discloses information to players, potentially altering the game’s equilibria. Our primary objective is to explore and characterize the conditions under which one oracle can be said to ‘dominate’ another. To this end, we define a partial order of dominance: One oracle dominates another if, in every game, the information structure of the former can induce at least the same set of equilibrium outcomes as the latter. This framework generalizes the classical results of Blackwell (1951), who focused on comparing signaling structures in decision problems.

The analysis is divided based on the oracles’ signaling capabilities: deterministic and stochastic. When oracles are limited to deterministic signaling functions, we show that an oracle dominates another if and only if it can simultaneously match the players’ posterior beliefs induced by the other oracle, while accounting for potential redundancies due to players’ private information (see Theorem 1 in Section 4). We refer to this condition as ‘Individually More Informative’ (IMI). Although the IMI condition may seem intuitive at first, it is fundamentally different from the refinement condition implied by Blackwell’s criterion for dominance, as evident from the stochastic characterization. Moreover, in our framework and unlike Blackwell’s result, if two oracles dominate each other under the IMI condition, then they must be identical (see Theorem 2 in Section 4). We prove this before extending our analysis to the stochastic setting.

The conditions for dominance in the stochastic setting differ from those in the deterministic one. When oracles are permitted to use stochastic signaling functions, the resulting posteriors become more complex. In this case, dominance requires additional criteria that depend on two

key elements based on the information structures of all players.

The first element is the common knowledge component (CKC), the smallest set, in terms of inclusion, that all players can agree upon (see Aumann, 1976). Using the structure of CKCs, we introduce the concept of an information loop, the second key element in our characterization. To formally define these loops, we first partition the state space into distinct CKCs. An information loop is then defined as a closed path of states that connects different CKCs through elements of an oracle's partition.

For example, assume there are 4 states $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ and two players whose private information is given by the following partitions: $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$ and $\Pi_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3, \omega_4\}\}$. The players' private information induces two CKCs: $C_1 = \{\omega_1, \omega_2\}$ and $C_2 = \{\omega_3, \omega_4\}$. That is, the two players can agree on each of these two events. See the illustration in Figure 1. If the oracle's information is given by the partition $F_1 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$, we say that a loop exists, as the different partition elements of F_1 form a closed path between the two CKCs. Namely, $\omega_1 \in C_1$ and $\omega_3 \in C_2$ are joined by a partition element of F_1 and the same holds for $\omega_2 \in C_1$ and $\omega_4 \in C_2$. This yields a sequence of states, that starts in C_1 , transitions to C_2 and reverts back again to C_1 , through different states that serve as entry and exit points from each CKC.




Figure 1: There are two CKCs $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4\}$. The oracle's partition F_1 generates a loop $(\omega_1, \omega_3, \omega_4, \omega_2)$, which is a closed path connecting the two CKCs using the oracle's partition elements.

Assuming that an oracle does not generate information loops (which includes the case where the entire state space comprises a unique CKC), we prove that it dominates the other oracle if

and only if its partition refines that of the other within every CKC (see Theorem 4 in Section 5.2 and Theorem 5 in Section 5.3). Importantly, this refinement condition does not follow from the IMI criterion used in the deterministic setting.

At this stage, we also prove that the refinement and dominance notions, given a unique CKC, are both equivalent to the *inclusion condition* which states that for every signaling strategy τ_2 of Oracle 2, there exists a signaling function τ_1 of Oracle 1 such that the set of the players' posterior beliefs profiles based on τ_1 is a subset of that based on τ_2 . Again, this holds given a single CKC, which obviously cannot admit information loops.

However, if a loop exists, the characterization becomes more complex. An information loop imposes (measurability) constraints on the information the oracle can convey. In the previous example, notice that every signaling function of the oracle over $\{\omega_1, \omega_2\}$, uniquely defines the signaling over $\{\omega_3, \omega_4\}$. Thus, the oracle is not free to signal any information it wants in one CKC, without restricting its ability to convey different information in the other CKC.

An obvious question that goes to the heart of information loops and our results is why should we care specifically about the signaling structure over the *pairs of states* that form the loop in every CKC? Moreover, why should a loop consist of separate entry and exit points in every CKC? The answer is that, given a CKC, Bayesian updating depends on the ratio of signal-probabilities for the different states. Thus, an effective constraint imposes restrictions over such ratios, thus relating to at least two states in every CKC (while keeping in mind the refinement condition in every CKC).

To tackle this issue, we need to thoroughly study the properties of information loops, and the first property is *non-informativeness*. A loop is called *non-informative* if, in every CKC that it intersects, all the states of the loop are in the same partition element of that oracle. We refer to this as non-informativeness because, conditional on the CKC and loop, the oracle has no-information to convey to the players. For example, in Figure 1, consider an oracle with a trivial partition $F'_1 = \Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$. This partition creates a closed path between the two CKCs, as well as joining all the states of the loop (given a CKC) to a single partition element of F'_1 . Building on this notion and assuming that the partition of Oracle 1 refines that of Oracle 2 in every CKC (as in the previously stated characterization), then non-informative loops do

not pose a problem for dominance and Oracle 1 dominates the other (see Theorem 7 in Section 5.5).

However, once a loop is *informative* (i.e., in at least one CKC that it intersects, there are states in the loop from different partition elements of the oracle; see Figure 2), then we require additional conditions for characterization. More specifically, in case there are only two CKCs, an additional condition is that Oracle 2 also has information loops whose states cover Oracle 1's loop (the notion of a cover is formally defined in Section 5.4). Using this condition we provide a characterization for the case of two CKCs (see Proposition 7 in Section 5.6), while the question of characterization in case of more than two CKCs remains open.

Yet, we should point out that the concept of information loops hints at a significant connection to Aumann's theory of common knowledge, from Aumann (1976). This link appears to be central to understanding how shared and differing information structures impact equilibrium outcomes in incomplete-information games. For this reason we provide an extensive set of result concerning information loops (see Section 5.4).

Another property that proves crucial for our analysis is the notion of irreducibility, which splits to two levels. The first is *irreducible loops*, which implies that there exists no (smaller) loop that is based on a strict subset of states taken from the original loop. The second is referred to as *type-2 irreducible loops*, and it implies that the loop does not contain four states from the same partition element of the oracle (again see Figure 2). On the one hand, type-2 irreducibility is a weaker notion compared to irreducible loops, because it allows for a loop to intersect the same CKC several times, whereas an irreducible loop cannot. On the other hand, a type-2 irreducible loop must be informative, because it does not allow for the entry and exit point in every CKC to be in the same partition element of that oracle. In fact, it is *fully-informative* because this condition holds in every CKC, rather than in a specific CKC.

The somewhat-delicate understanding of the relations between these loops properties allows us to achieve another main result: the characterization of equivalent oracles. Formally, we say that two oracles are equivalent if they simultaneously dominate one another. The characterization of equivalence, given in Theorem 8 in Section 6, is based on: (i) equivalence in every CKC; (ii) equivalence of irreducible-informative loops; and (iii) a cover over loops. To prove




Figure 2: An illustration of a fully-informative and irreducible loop, which intersects three CKCs C_1, C_2 and C_3 with two states in each.

this result, we use type-2 irreducible loops to compare the information of both oracles. Specifically, we consider the sets of type-2 irreducible loops that intersect a joint CKC (i.e., *connected loops*), also taking into account sequential intersections (i.e., the transitive closure) where loop 1 is connected to loop 2 which is then connected to loop 3 and so on. We observe the set of CKCs for each of these groups and refer to these sets as *clusters*. These are used as building blocks in our analysis, and we prove that the information of equivalent oracles must match on these clusters. This, in turn, provides some insight into possible future characterization of general dominance between oracles, as well as provides another level of extending the theory of Aumann (1976) on common knowledge, beyond information loops.

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 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.¹

¹See, 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).

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 coordinate 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) explores common-interest games, characterizing dominance through the concept of individual sufficiency—an extension of Blackwell's notion of garbling to n -player games.

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, Sections 5.2 and 5.3 characterize dominance in the cases of a unique CKC and multiple CKCs without loops, respectively. Section 5.4 outlines necessary and sufficient conditions for dominance. In Section 5.6, we provide a characterization of dominance in the context of two CKCs. Finally, in Section 6 we characterize the equivalence relation between oracles.

2 A simple example: the rock-concert standoff

To understand these concepts, consider a simple example of competition between two rock bands.² 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

²We thank Alon Eizenberg from the Hebrew University and two 1990s rock bands who inspired this example.

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 would prefer to 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 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 1 cannot differentiate between n_2 and s_2 , while Band 2 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 3.

Based on the realized state, the bands engage in the game depicted in Figure 4. 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.

However, 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



(a)

The bands' information



(b)

The weather forecasters' information

Figure 3: 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 $F'_1 = \{\{n_1, s_1, s_2\}, \{n_2\}\}$.

		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 4: Payoff matrices for sunny and stormy conditions.

Band 2 randomizes equally between M and I 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.

However, 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 1800 and 950 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.³

Second, if this were 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.

³Notably, given a unique CKC, we prove that two-sided IMI implies that the two partitions coincide. See Section 4.2.1.

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 actions⁴ 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.

To extend the basic game to a guided game, we introduce an oracle who provides public information before players take their actions. For that purpose, the oracle has a partition F over Ω , and a countable set S of signals. A strategy of the oracle is an F -measurable function $\tau : F \rightarrow \Delta(S)$ used to transmit information to all players N , where $\Delta(S)$ is the set of all distributions on finite subsets of S . We denote by $\tau(s|\omega)$ the probability $\tau(\omega)(s)$ at which τ transmits the signal s when the realized state is ω . Note that any deterministic strategy $\tau : F \rightarrow S$ is essentially 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 the case τ is deterministic, or $s \in S$ is drawn according to $\tau(\omega)$ and is publicly announced in the case where τ is stochastic.

Let the join⁵ $\Pi_i \vee F'$ 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 distribution after observing $\Pi_i(\omega)$ and $\tau(\omega)$. In case τ is stochastic,

⁴In this setting, A_i is independent of the player's information; however, the current framework can also accommodate scenarios where it is not.

⁵Coarsest common refinement of Π_i and F' ; following the definition of Aumann (1976).

let $\mu_{\tau|\omega,s}^i = \mu(\cdot|\Pi_i(\omega), \tau, s) \in \Delta(\Omega)$ denote player i 's posterior distribution after observing $\Pi_i(\omega)$ and a realized signal s according to $\tau(\omega)$. Thus, every strategy τ yields an incomplete-information game $G(\tau) = (N, (A_i)_{i \in N}, (\mu_{\tau}^i)_{i \in N}, (u_i)_{i \in N})$. 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 .

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. 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 .⁶ 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.

⁶Note that a Nash equilibrium $(\sigma_1^*, \dots, \sigma_n^*)$ 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 $(\sigma_1^*, \dots, \sigma_n^*)$ 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 $\Delta(\Omega \times A)$.

Definition 1 (Partial ordering of Oracles). 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.

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. We provide a necessary condition for equivalent oracles in Section 5.

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 oracles—such as the Federal Reserve—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 (see Sections 5.4 and 6).

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 5.




Figure 5: 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 6.

$\tau_2(s \omega)$	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 6: A stochastic F_2 -measurable signaling strategy of Oracle 2.

Combined with Π , this signaling strategy τ_2 is equivalent to the following Blackwell experiment e , given in Figure 7.

$e(s \omega)$	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 7: $M_2(\tau_2, \Pi)$ - 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 $M_2(\tau_2, \Pi)$ Blackwell dominates $M_2(\tau_2, \Pi)$ (see Theorem 12.3.1 of Blackwell and Girshick (1954)).⁷ 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

⁷Note that for this result to hold, Π is fixed and it is independent of τ_1 and τ_2 .

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? Simultaneously, the second objective of this paper is to extend 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.⁸ 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. The following Figure 8 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.


Now consider a game where both player have two actions: D and M , and the payoffs are given by the matrices in Figure 9. The best common payoff is attained when both players know the realized state. Oracle 2, who is fully informed, can simply reveal the true state. Oracle 1,

⁸As 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.



(a)

The players' information



(b)

The oracles' information

Figure 8: 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).

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

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

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

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

Figure 9: Payoff matrices for each ω

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.⁹

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.

⁹This example provides a concrete instance of Theorem 5.

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 simultaneously *match* the information of each player, 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 and the publicly available signal (restricted to deterministic signaling functions). In other words, Oracle 1 only needs the ability to match the information that Oracle 2 transmits *simultaneously* to each player, considering the redundancies given the private information of the players. 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.*

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 .

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 10. Nevertheless, in Section 4.2.1 we prove that if F_1 is IMI then F_2 and vice versa, it implies that the two partitions do partially coincide.

Individually More Informative versus Refinement




Figure 10: 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 .

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 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.

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

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 11 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 11: 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. *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 F_1 is IMI than F_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 decision problems as in Blackwell (1951, 1953) and Example 4 here), 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 10 depicts two CKCs: $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4\}$.

Regarding 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 separately across different 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 12. 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 $F_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$.

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 12: 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 13. 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 13: 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 14. 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 strategy 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.

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, F_i is IMI than F_{-i} for every Oracle i if and only if*




Figure 14: 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 refines F_2 in every CKC, F_1 is not individually more informative than F_2 .

$$F_1 = F_2.$$

In other words, 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 characterize dominance among oracles given they can exercise general signaling strategies, not restricted to deterministic ones. This goal is achieved in several 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, he 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 . In Section 5.3 we introduce the concept of *information loops* between CKCs. In general, an F_i -loop is a closed path among different CKCs, connected

through information sets of Oracle i . In case there are no such loops, we extend the result described in Section 5.2, and prove that oracle-dominance is equivalent to partition refinement in every CKC. In Section 5.5 we provide necessary and sufficient conditions for dominance, in general. In Section 5.6 we connect the stated (necessary and sufficient) conditions in a setting with two CKCs, providing a characterization for this set-up as well.

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 15.

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 16. 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

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

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




Figure 15: 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).

$\tau_2(s \omega)$	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 16: A stochastic F_2 -measurable strategy of Oracle 2.

$\mu_{\tau_1|\omega_3, s_4}^2 = (0, 0, 3/5, 2/5)$, 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 *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 (based on the concept of loops).

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 .

For this purpose, let us define the sets of feasible signals and posteriors. Formally, for every

strategy τ , let $S_\tau = \{s \in S : \exists \omega \in \Omega : \tau(s|\omega) > 0\}$ be the set of feasible signals, and let $\text{Post}(\tau)$ denote the set of feasible posteriors profiles,

$$\text{Post}(\tau) = \left\{ p \in (\Delta(\Omega))^n : \exists (\omega, s) \text{ s.t. } \tau(s|\omega) > 0 \text{ and } p = (\mu_{\tau|\omega, s}^i)_{i \in N} \right\}.$$

Note that for every (τ, ω, s) , where $\tau(s|\omega) > 0$, there exists a unique posterior $p \in (\Delta(\Omega))^n$ where $p = (\mu_{\tau|\omega, s}^i)_{i \in N} \in \text{Post}(\tau)$, so the sets are well-defined. Let $\mu_\tau \in \Delta(\Delta(\Omega)^n)$ be the distribution over posteriors profiles given a strategy τ .

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 17.

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*




Figure 17: The two-stage game $\mathbf{G}_{\tau_2}$, under any signaling strategy τ .

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 the 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 the ratio of every two distinct numbers from the set $\mathbb{A} = \{p_j, 1 - p_j : j = 1, 2, \dots, m\}$ is distinct.¹³ Define the signaling function τ_2 such that

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

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

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 $p_j \in \mathbb{T}_j \setminus \mathbb{T}_{j-1}$.

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 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 .

5.2.1 More than one CKC: two examples

The refinement condition given in Theorem 4 ensures that Oracle 1 can produce the *exact* same strategy as Oracle 2. This however hinges on the existence of a unique CKC. In case there are several CKCs, Oracle 1 may need to follow a different strategy in order to match the distribution on posteriors generated by τ_2 . Namely, τ_1 may require more signals than τ_2 , even if both oracles have the same (complete) information in every CKC. Let us provide a concrete example for this.

Example 7. *More signals are needed.*

Consider a uniformly distributed state space $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, with two players whose private information is $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$ and $\Pi_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3, \omega_4\}\}$. The oracles have the following partitions $F_1 = \{\{\omega_1, \omega_3\}, \{\omega_2\}, \{\omega_4\}\}$ and $F_2 = \{\{\omega_1\}, \{\omega_3\}, \{\omega_2, \omega_4\}\}$. This information structure is illustrated in Figure 18. Notice that there are two CKCs, $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4\}$, and both oracles have complete information in each of these components. That is, F_1 refines F_2 in every CKC, and vice versa.




Figure 18: 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).

Consider the stochastic strategy τ_2 given in Figure 19. Notice it is F_2 -measurable, as $\tau_2(s|\omega_2) = \tau_2(s|\omega_4)$ for every signal s , but not F_1 -measurable.

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

Figure 19: A stochastic F_2 -measurable strategy of Oracle 2.

The set $\text{Post}(\tau_2)$ of τ_2 -posteriors is

$$\text{Post}(\tau_2) = \left\{ \begin{array}{ll} (e_i, e_i), & \forall 1 \leq i \leq 4, \\ ((\frac{3}{7}, \frac{4}{7}, 0, 0), e_j), & j = 1, 2, \\ (e_k, (0, 0, \frac{1}{2}, \frac{1}{2})), & k = 3, 4 \end{array} \right\},$$

and we can now try to mimic τ_2 using an F_1 -measurable strategy. First, this requires at least two signals to distinguish between ω_1 and ω_2 , as well as ω_3 and ω_4 . Second, the posterior $((\frac{3}{7}, \frac{4}{7}, 0, 0), e_1)$ requires another signal s so that $\tau(s|\omega_1) = \alpha > 0$ and $\tau(s|\omega_3) = \frac{4}{3}\alpha > 0$. However, the F_1 -measurability requirement implies that $\tau(s|\omega_3) = \alpha$, and the τ_2 -posterior $(e_3, (0, 0, \frac{1}{2}, \frac{1}{2}))$ necessitates that $\tau(s|\omega_4) = \alpha$ as well. These conditions are jointly given in Table (a) within Figure 20.

$\tau_1(s \omega)$	s_3	s_4	s_5
ω_1	α	β	0
ω_2	$\frac{4}{3}\alpha$	0	γ
ω_3	α	β	0
ω_4	α	0	γ

(a)

$\tau_1(s \omega)$	s_3	s_4	s_5	s_6
ω_1	1/2	1/3	0	1/6
ω_2	2/3	0	1/3	0
ω_3	1/2	1/3	0	1/6
ω_4	1/2	0	1/3	1/6

(b)

Figure 20: A strategy τ_1 , either with 3 signals as given in Table (a), or with 4 signals as in Table (b).

Evidently, it must be that $\alpha, \beta, \gamma > 0$ in order to mimic τ_2 , but the second and fourth rows in Table (a) cannot jointly sum to 1 unless $\alpha = 0$, which eliminates the possibility of a well-defined mimicking strategy. Thus, in order to mimic the stated strategy τ_2 , Oracle 1 requires an additional signal as presented in Table (b), in Figure 20. To conclude, though the oracles' partitions refine one another in every CKC, they cannot always produce the exact same

strategy when trying to mimic each other.

Example 8. *Dominance need not imply refinement with multiple CKCs*

In this example we wish to show that when there are multiple CKCs, Oracle 1 can dominate Oracle 2 although F_1 does not refine F_2 . To see this, we revisit Example 4. Consider the following signaling strategy of Oracle 2 given in Figure 21.

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

Figure 21: A stochastic F_2 -measurable strategy of Oracle 2.

Here, Oracle 2 provides the players with no additional information regarding states ω_1 and ω_2 . Thus, the posterior over these states remains the original one. On the other hand, given the states ω_3 and ω_4 , the strategy τ_2 reveals the true state with a positive probability and induces the posterior $(0, 0, 2/5, 3/5)$ with the remaining probability.

While Oracle 2 can assign different probabilities to a signal conditioned on ω_2 and ω_3 , Oracle 1 cannot. However, there is a signaling strategy for Oracle 1 that produces the same distribution over the posteriors as τ_2 does. The following strategy τ_1 , given in Figure 22, does that.

$\tau_1(s \omega)$	s_1	s_2	s_3
ω_1	1/2	0	1/2
ω_2	1/2	0	1/2
ω_3	1/2	0	1/2
ω_4	0	1/4	3/4

Figure 22: A stochastic F_1 -measurable strategy of Oracle 1.

In this example, it is straightforward to prove that Oracle 1 can mimic every strategy τ_2 of Oracle 2, and we prove this result under more general conditions in Theorem 5 and Proposition 7. Yet, it is clear that F_1 is not a refinement of F_2 in general, but it is a refinement in every CKC.

5.3 Multiple CKCs and no loops

We now turn to the general setting in which the players' information structures induce any (finite) number of CKCs. Assume that $C_1, \dots, C_l$ are mutually exclusive CKCs such that $\Omega = \bigcup_{j=1}^l C_j$. A key aspect of our analysis is the presence of measurability constraints, where different CKCs are connected by atoms of the oracles' partitions. To understand the significance of this, consider a setting where F_1 does not contain any element intersecting multiple CKCs. In this case, Theorem 4 applies separately to each CKC, as Oracle 1 faces no constraints when attempting to mimic some strategy of Oracle 2.

However, when elements of Oracle 1's partition intersect different CKCs, the analysis becomes more complex, because we must account for measurability constraints when attempting to use the same strategy τ_1 across different CKCs. Such intersections impose constraints on τ_1 , preventing us from naively applying Theorem 4.

This issue becomes even more complicated when multiple elements of Oracle 1's partition intersect different CKCs, forming what we call an (information) *loop*.¹⁴

Generally, a loop is an ordered sequence of states from different CKCs such that the partition of an oracle groups together distinct pairs of states from different CKCs, creating a closed path. The main result of this section, presented in Theorem 5 below, states that in the absence of such loops, Oracle 1 dominates Oracle 2 if and only if F_1 refines F_2 in every CKC. The formal definition of a loop is provided in Definition 3.

Definition 3. *An F_i -loop is a sequence $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$, where $m + 1 \equiv 1$ and $m \geq 2$, such that*

- $\omega_j, \bar{\omega}_j \in C_{r_j}$ and $\omega_j \neq \bar{\omega}_j$ for all $j = 1, \dots, m$.¹⁵
- $\omega_{j+1} \in F_i(\bar{\omega}_j)$ for all $j = 1, \dots, m$.
- $C_{r_j} \neq C_{r_{j+1}}$ for all $j = 1, \dots, m$.
- The sets $\{\bar{\omega}_j, \omega_{j+1}\}$ are pairwise disjoint for all $j = 1, \dots, m$.

¹⁴An (information) loop is different from a loop in graph theory. In graph theory, a loop refers to an edge that connects a vertex to itself.

¹⁵Here C_{r_j} refers to the CKC that contains the j -th pair of states $(\omega_j, \bar{\omega}_j)$.

To understand information loops, one can view the CKCs as vertices of a graph. An edge connects two CKCs if there exist ω_{j+1} and $\bar{\omega}_j$ such that they belong to the same F_i -partition element (this corresponds to the second requirement). An information loop then parallels an Eulerian graph, where there is a walk that includes every edge exactly once (the last requirement in the definition) and ends back at the initial vertex (hence the requirement $m + 1 \equiv 1$). As noted at the beginning of Section 5.3, the key aspect of the general analysis is to consider the case when the oracle partition atoms intersect different CKCs, so we require that $C_{r_j} \neq C_{r_{j+1}}$ for all $j = 1, \dots, m$.

An example of an F_1 -loop is provided in Figure 23.(a), which depicts a loop consisting of six states across three CKCs. Note that a loop can intersect the same CKC multiple times, as long as the sets $\{\bar{\omega}_j, \omega_{j+1}\}$ remain pairwise disjoint for each j .

We use the concept of a loop in our first general characterization, presented in Theorem 5. This theorem builds on the assumption that F_1 contains no loops and extends Theorem 4 by showing that one oracle dominates another if the former's partition refines that of the latter in every CKC. It is important to note that the proof is extensive, as it must account for the measurability constraints of τ_1 across all CKCs.

Theorem 5. *Assume there is no F_1 -loop. Then, Oracle 1 dominates Oracle 2 if and only if F_1 refines F_2 in every CKC.*

The proof of Theorem 5 builds on the concept of a *sub-strategy*. A sub-strategy is a signaling function without the requirement that the probabilities sum to 1. This relaxation allows us to study functions that partially mimic a strategy τ_2 , meaning each posterior is drawn from $\text{Post}(\tau_2)$ and is induced with a probability that does not exceed the probability with which τ_2 induces it. We show that the set of sub-strategies is compact, allowing us to consider an optimal sub-strategy for mimicking τ_2 . The proof then proceeds by contradiction: if the optimal sub-strategy is not a complete strategy, we can extend it by constructing an additional sub-strategy to complement the optimal one for posteriors that are not fully supported (relative to the probabilities induced by τ_2). This part is rather extensive as it requires some graph theory and several supporting claims given in the proof in the appendix.




Figure 23: Figure (a) depicts an F_1 -loop with three CKCs and six states overall. Figure (b) illustrates how the F_1 -loop, presented in (a), is non-balanced with respect to F_2 . Namely, F_2 has two elements $A = \{\omega_1, \omega_2, \omega_3\}$, and $B = \{\bar{\omega}_1, \bar{\omega}_2, \bar{\omega}_3\}$ such that the number of transitions from A to B are 3, while the reverse equals 0.

5.4 Information loops

Previous sections have examined the problem of oracle dominance in the absence of loops, considering either a unique CKC or multiple CKCs without loops. However, in order to confront the general question of dominance in the presence of information loops, we need to have a clear understanding of their properties and implications.

Specifically, when an F_1 -loop exists, it may create challenges for Oracle 1 in mimicking Oracle 2, because loops introduce measurability constraints across CKCs. Although Oracle 1 can mimic Oracle 2 within each CKC individually, it may be impossible to do so simultaneously across CKCs if the required combined strategy is not measurable with respect to F_1 . This suggests that any F_1 -loop must satisfy certain conditions to ensure that such a strategy is indeed F_1 -measurable. The first condition that we study, which turns out to be a necessary condition for dominance, is generally referred to as F_2 -balanced.

The idea starts with an F_1 -loop. We examine all states in this loop and determine how they can be covered by two F_2 -measurable sets. In other words, the loop is divided into two disjoint sets, each contained in an F_2 -measurable set, denoted A and B . Next, we count the number of transitions along the loop from A to B , where the entry point into one CKC is through a state in A and the exit is through a state in B . We do the same for transitions from B to A . An F_1 -loop is called F_2 -balanced if the number of transitions between A and B is equal in both directions. The formal definition follows.

Definition 4. An F_i -loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$ is F_{-i} -balanced if for every F_{-i} -measurable partition of the loop's states into two disjoint sets $\{A, B\}$ such that $\cup_j \{\omega_j, \bar{\omega}_j\} \subseteq A \cup B$, it follows that:

$$\#(A \rightarrow B) := |\{j; \omega_j \in A \text{ and } \bar{\omega}_j \in B\}| = |\{j; \omega_j \in B \text{ and } \bar{\omega}_j \in A\}| =: \#(B \rightarrow A). \quad (2)$$

Note that an F_1 -loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$, where $\omega_j \in F_2(\bar{\omega}_j)$ for all $j = 1, \dots, m$, is F_2 -balanced. Figure 23.(b) examines the F_1 -loop from Figure 23.(a). The sets A and B are F_2 -measurable, restricted to the six states under consideration. The partition into A and B renders the loop non- F_2 -balanced, as $\#(A \rightarrow B) = 3$, while $\#(B \rightarrow A) = 0$.

Why are balanced loops crucial? The intuition follows from Lemma 2, which must hold in any CKC, but presents a challenge when a loop is non-balanced. Consider, for example, a non-balanced loop as depicted in Figure 23, and assume that $\tau_2(s|\omega) = \frac{1}{2} - \frac{1}{4}\mathbf{1}_{\{\omega \in A\}}$ for some signal $s \in S$. This imposes a specific 1 : 2 ratio between any two states described in each CKC, so that $\Pi_i \frac{\tau_2(s|\omega_i)}{\tau_2(s|\bar{\omega}_i)} = \frac{1}{8}$. However, since $\bar{\omega}_i$ and ω_{i+1} belong to the same F_1 partition element, the measurability constraints on Oracle 1 along the loop require that $\tau_1(s|\bar{\omega}_i) = \tau_1(s|\omega_{i+1})$, hence $\Pi_i \frac{\tau_1(s|\omega_i)}{\tau_1(s|\bar{\omega}_i)} = 1$ for any s in the support of all states. In other words, Oracle 1 cannot match the ratio dictated by τ_2 , therefore Lemma 2 does not hold in at least one CKC.

If the loop were balanced—say, with $A = \{\bar{\omega}_1, \omega_2\}$ and $B = \{\omega_1, \bar{\omega}_2, \omega_3, \bar{\omega}_3\}$ —then the same strategy τ_2 would yield $\Pi_i \frac{\tau_2(s|\omega_i)}{\tau_2(s|\bar{\omega}_i)} = 1$, as required. In general, when all loops are balanced, this discrepancy is eliminated for any two such sets A and B . The notion of balanced loops is closely related to the following notion of *covered loops*, which implies that an F_1 -loop can be decomposed to loops of F_2 .

Definition 5. An F_i -loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$ is F_{-i} -covered if

- The set $\{1, \dots, m\}$ is partitioned to disjoint sets of indices, $J, I_1, \dots, I_r$, i.e., $\{1, \dots, m\} = J \cup (\cup_{t=1}^r I_t)$.
- For each $t = 1, \dots, r$, $\left((\omega_j, \bar{\omega}_j)\right)_{j \in I_t}$ is an F_{-i} -loop, also referred to as a sub-loop.¹⁶

¹⁶The order of the pairs $(\omega_j, \bar{\omega}_j)$ in the F_{-i} -loop does not have to coincide with their order under the F_i -loop.

- $J = \{j; \omega_j \in F_{-i}(\bar{\omega}_j)\}$.

The cover is order-preserving if every F_{-i} -loop $\left((\omega_j, \bar{\omega}_j)\right)_{j \in I_t}$ in the cover follows the same ordering of pairs as the F_i -loop.

In simple terms, the definition states that, given an F_1 -loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$, we can partition its states to several F_2 -loops and a set of states where $\omega_j \in F_2(\bar{\omega}_j)$. Figure 24 (a) depicts an F_1 -loop consisting of $((\omega_j, \bar{\omega}_j))_{j=1,\dots,4}$, which is covered by two F_2 -loops: $(\omega_1, \bar{\omega}_1, \omega_3, \bar{\omega}_3)$ and $(\omega_2, \bar{\omega}_2, \omega_4, \bar{\omega}_4)$. In this case, the set J (defined in Definition 5) is empty. Figure 24 (b) depicts another case in which the F_1 -loop is covered by F_2 -loops, but $J = \{2, 4\}$. Note that the sub-loops in Figure 24 (a) are order-preserving, whereas those in Figure 24 (b) are not.




Figure 24: Two states connected by a colored line are in the same information set of F_2 . In (a), the sub-loops are order-preserving, i.e., following the ordering of pairs in the original F_1 -loop, whereas those in (b) are not.

The following Proposition 5 proves that an F_1 -loop is F_2 -balanced if and only if it is F_2 -covered. This proposition assists with the proof of Theorem 6 below, which provides a necessary condition for dominance.

Proposition 5. *Let $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$ be an F_1 -loop. The following statements are equivalent:*

For instance, an F_1 -loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \omega_3, \bar{\omega}_3)$ might be covered by the following F_2 -loop $(\omega_1, \bar{\omega}_1, \omega_3, \bar{\omega}_3, \omega_2, \bar{\omega}_2)$.

- i. The loop is F_2 -balanced;
- ii. The loop is F_2 -covered;
- iii. For every F_2 -measurable function $f : \{\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m\} \rightarrow (0, \infty)$,

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

The next two properties that we study are *irreducible* and *informative* loops. Starting with the former, an F_i -loop is irreducible if it does not have a *sub-loop*, namely, there exists no ‘smaller’ F_i -loop that comprises a strictly smaller set of states taken solely from the original loop. Our analysis would use irreducible loops as building blocks to decompose and compare loops generated by the oracles’ partitions.

Definition 6. Let $L_i = (\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$ be an F_i -loop. We say that the loop is irreducible if there exists no strict subset of the set $\{\omega_j, \bar{\omega}_j : j = 1, \dots, m\}$ that forms an F_i -loop.

We use the definition of an irreducible loop in the context of covers as well, stating that a cover is *irreducible* if every loop in the cover is irreducible. Furthermore, the idea of irreducible loops is closely related to the concept of covers, and specifically to the set $J = \{j; \omega_j \in F_{-i}(\bar{\omega}_j)\}$ given in Definition 5 above. Specifically, if there exists an F_i -loop with a pair of states $(\omega_j, \bar{\omega}_j)$ such that $\bar{\omega}_j \in F_i(\omega_j)$, then it cannot be irreducible unless it comprises only 4 states.¹⁷ We typically refer to such cases where $\bar{\omega}_j \in F_i(\omega_j)$ as *non-informative* because Oracle i cannot distinguish between the two states. This condition is essentially equivalent to every F_1 -loop being F_2 -balanced at 0, meaning that for any choice of the specified F_2 -measurable sets A and B , the number of transitions between these sets is zero. The following Definition 7 captures the idea of *informative* loops, which would later be used in Theorem 7 as a sufficient condition for dominance.

Definition 7. An F_i -loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$ is F_k -non-informative if $F_k(\omega_j) = F_k(\bar{\omega}_j)$ for every j . The loop is F_k -fully-informative if $F_k(\omega_j) \neq F_k(\bar{\omega}_j)$ for every j .

¹⁷In general, the smallest possible loop has at least 4 states, so any such loop is, by definition, irreducible.

To understand the motivation behind this definition, consider any F_1 -loop denoted by $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$. If this loop is F_2 -non-informative, it suggests that the ratios $\frac{\tau_2(s|\omega_i)}{\tau_2(s|\bar{\omega}_i)}$ equals 1 for every signal s supported on these states. In simple terms, conditional on any $\{\omega_i, \bar{\omega}_i\}$, Oracle 2 does not provide any additional information, so the constraints that an F_1 -loop imposes on Oracle 1 in every CKC (i.e., that the product of probability ratios along the loop equals 1) are met by the measurability requirements of F_2 .

The following proposition summarizes key properties of informative and irreducible loops. It states that an irreducible loop intersects every CKC at most once and must be fully-informative (unless it has only 4 states). In addition, the proposition shows that an informative loop has a fully-informative sub-loop, as well.

Proposition 6. *Consider an F_i -loop L_i .*

- *If L_i intersects the same CKC more than once, then it is not irreducible.*
- *If L_i is irreducible and consists of at least 6 states, then it is F_i -fully-informative.*
- *If L_i is F_i -informative, then it has an F_i -fully-informative sub-loop.*
- *If L_i is F_i -fully-informative, then it can be decomposed to irreducible F_i -loops.*
- *If L_i is not irreducible, then either it intersects the same CKC more than once, or it has at least 4 states in the same partition element of F_i .*

We use this proposition in the following subsection to provide necessary and sufficient condition for the dominance of one oracle over another.

5.5 Necessary and Sufficient conditions for dominance

In the following section, we address the general case where F_1 has loops, which imposes constraints on Oracle 1 *across* CKCs. Due to the complexity of this problem, we divide our analysis into two parts: a necessary condition for dominance presented in Theorem 6, and a sufficient condition given in Theorem 7. These theorems depend strongly on the properties of information loops, and specifically on the notions of covers, irreducibility and non-informativeness.

Starting with the necessary conditions, the following theorem, which builds on Propositions 5 and 6, states that if Oracle 1 dominates Oracle 2, then besides the refinement condition in every CKC, already established in Theorem 5, it must be that every F_1 -loop is covered by loops of F_2 . In addition, it states that every irreducible F_2 -loop that cover an irreducible F_1 -loop is order-preserving, essentially stating that the two loops coincide.

Theorem 6. *If Oracle 1 dominates Oracle 2, then:*

- F_1 refines F_2 in every CKC;
- Any F_1 -loop has a cover by F_2 -loops; and
- Every irreducible F_2 -loop that covers an irreducible F_1 -loop is order-preserving.

The proof of the first part is immediate, as it follows directly from Theorem 4. The proof of the second part relies on Proposition 5 by assuming that an F_1 -loop is not F_2 -balanced, and constructing a strategy τ_2 that Oracle 1 cannot mimic without violating measurability constraints. The last part relies on Proposition 6, as well as Lemma 2, by depicting a two-signal strategy τ_2 that one cannot mimic without following the same order of pairs throughout the F_2 -loop.

Next, we use the understanding regarding covered and balanced loops to present a sufficient condition for dominance, which indirectly requires that any loop is balanced at 0—meaning that there are no transitions between sets A and B . This leads to the following Theorem 7, which uses the non-informative notion for dominance.

Theorem 7. *If F_1 refines F_2 in every CKC and every F_1 -loop is F_2 -non-informative, then Oracle 1 dominates Oracle 2.*

Though we do not yet provide a full characterization, it becomes rather clear that the requirement that every F_1 -loop is F_2 -balanced should be the main focus, as it is a necessary condition, as well as a sufficient one when the balance is set to zero. In the following section we show that the balance condition is both necessary and sufficient for the case of two CKCs.

5.6 Toward a general characterization: two CKCs

In this section, we assume there are only two CKCs. This assumption simplifies the analysis, as the case of two CKCs can be resolved using our prior results, allowing us to examine all possible loops directly. Formally, Proposition 7 states that, given two CKCs, the necessary condition of an F_2 -balanced loop from Theorem 6 is also a sufficient condition.

To build intuition, consider the scenario with two CKCs depicted in Figure 25, featuring an F_1 -loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2)$ across four states. Fix some τ_2 and assume the loop is F_2 -balanced. There are then only two possibilities: either the loop is F_2 -non-informative, as shown in cases (a) and (b) in Figure 25, or it is also an F_2 -loop, illustrated in case (c) in Figure 25. The first possibility was covered in Theorem 7, while the second allows Oracle 1 to meet the constraints imposed by the F_1 -loop when attempting to mimic τ_2 .

Proposition 7. *Assume there are only two CKCs. Then, Oracle 1 dominates Oracle 2 if and only if F_1 refines F_2 in every CKC and any F_1 -loop is F_2 -balanced.*




Figure 25: Two CKCs with an F_1 -loop described by $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2)$. Graph (a) and (b) depict two F_2 -balanced loops, that are also F_2 -non-informative, and (c) describes an F_2 -loop. Any other structure of F_2 yields a non-balanced loop.

6 Equivalent oracles

In this section we tackle a parallel question to dominance, which is the problem of oracles' equivalence. Specifically, we characterize necessary and sufficient conditions such that both oracles dominate one another simultaneously, as formally given in the following definition:

Definition 8. F_1 is equivalent to F_2 , denoted $F_1 \sim F_2$, if the two oracles dominate one another, that is, if $F_i \succeq_{\text{NE}} F_{-i}$ for every $i = 1, 2$.

Based on the results for the case that loops do not exist and the case of two CKCs, equivalence between oracles obviously requires two-sided refinement within every CKC (i.e., equivalence), and that every F_i -loop is F_{-i} -balanced for every Oracle i . This, however, is insufficient and equivalence also requires that every irreducible F_i -loop with at least 6 states is also an irreducible F_{-i} -loop. This result is given in the following Theorem 8.

Theorem 8. F_1 is equivalent to F_2 if and only if for every Oracle i , the partition F_i refines F_{-i} in every CKC, any F_i -loop has a cover of F_{-i} -loops, and every irreducible F_i -loop with at least 6 states is an irreducible F_{-i} -loop.

The equivalence condition concerning irreducible loops is based on the ability of both oracles to follow similar measurability constraints when signaling to players in every CKC. That is, if one oracle is constrained by an information loop, then we require the other to follow suit. Yet, this still raises the question of why do we need to focus on irreducible loops? To understand this, consider a single partition element of F_i that intersects at least two CKCs where each intersection contains at least two states. This evidently generates a non-informative loop, because all pairs are non-informative. But as long as the other oracle cannot distinguish between the two states in each pair, the ability to separate different pairs in different CKCs is not needed, as each pair is common knowledge among the players themselves within every CKC.

The proof of Theorem 8 also builds on an intermediate irreducibility notion that we refer to as *type-2 irreducible loop*. More formally, an F_i -loop is type-2 irreducible if it does not have four states from the same partition element of F_i . This notion refines that of fully-informative loops (as every type-2 irreducible loop is fully-informative), but also weakens that of irreducible loops, because a type-2 irreducible loop can intersect the same CKC multiple times, and so be decomposed to sub-loops.

The notion of type-2 irreducible loops is crucial for our analysis and results, but also in a more general manner. We use type-2 irreducible loops to generate the basic elements, *building blocks*, upon which two oracles must match one another (in terms of their information). These

building blocks are referred to as *clusters* and they are constructed as follows. First, we take the set of type-2 irreducible loops. Then, we consider such loops that intersect the same CKC and consider them as connected. Next, we take the transitive-closure of this relation, which yield disjoint sets of connected type-2 irreducible loops. Finally, we take every such set (of connected loops) and consider all the CKCs that it intersects - this is a cluster. We prove that the oracles' partitions match one another in each of these clusters. That is, the clusters are the basic structure upon which we derive an equivalence, and later extend it to "simpler" connections between clusters that involve only a single partition element of F_i .

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.
- 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 ,¹⁸ 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

¹⁸Here we refer to τ as a Blackwell experiment.

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

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

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.¹⁹ 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. $\square$

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

¹⁹We cannot reuse Claim 1 here because there is no common garbling for all players: each has its own garbling matrix.

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 26.(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 $\Pi_i(\omega) \neq \Pi_i(\omega')$.

Illustrations of sub-partitions in the proof of Theorem 2




Figure 26: 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 .

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 $\omega_3 \in F_2(\omega_1) \cap (F_1(\omega_1))^c$. See Figure 26.(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

26.(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 $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$, 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 $\text{NED}(\mathbf{G}(\tau_2)) \neq \text{NED}(\mathbf{G}(\tau_1))$ 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 $\text{NED}(\mathbf{G}(\tau_2)) \neq \text{NED}(\mathbf{G}(\tau_1))$ 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

$$\frac{\tau_1(t|\omega_1)}{\tau_2(s_i|\omega_1)} \neq \frac{\tau_1(t|\omega^*)}{\tau_2(s_i|\omega^*)}. \quad (3)$$

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

$$\begin{aligned}\tau_2(s_i|\omega_1) &= c_1 \cdot \tau_2(s_i|\omega_2) = c_1 \cdot c_2 \cdot \tau_2(s_i|\omega_3), \\ \tau_1(t|\omega_1) &= c_1 \cdot \tau_1(t|\omega_2) = c_1 \cdot c_2 \cdot \tau_1(t|\omega_3).\end{aligned}\tag{4}$$

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

$$\begin{aligned}\tau_2(s_i|\omega_2) &= c_2 \cdot \tau_2(s_i|\omega_3) = c_2 \cdot c_3 \cdot \tau_2(s_i|\omega_4), \\ \tau_1(t|\omega_2) &= c_2 \cdot \tau_1(t|\omega_3) = c_2 \cdot c_3 \cdot \tau_1(t|\omega_4),\end{aligned}\tag{5}$$

and the first equality in Equation (5) coincides with the second equality in Equation (4). Namely, Equations (4) and (5) either hold for every signal s_i , or hold for the same signal s_1 . Repeatedly following the same procedure, we get that

$$\tau_2(s_i|\omega_1) = c_1 \cdot \tau_2(s_i|\omega_2) = \dots = [\prod_{k \geq 1} c_k] \cdot \tau_2(s_i|\omega^*),\tag{6}$$

$$\tau_1(t|\omega_1) = c_1 \cdot \tau_1(t|\omega_2) = \dots = [\prod_{k \geq 1} c_k] \cdot \tau_1(t|\omega^*).\tag{7}$$

Dividing Equation (7) by Equation (6), 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 (3), 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 posterior profiles match), which implies the third condition, thus concluding the proof. $\square$

A.10 Proof of Theorem 5

Proof. One direction is straightforward. Assume, to the contrary, that Oracle 1 dominates Oracle 2, but F_1 does not refine F_2 in some CKC. Denote this CKC by C_1 , and consider the set of all games where the payoffs of all players are zero in every $\omega \notin C_1$ independently of their actions. Thus, Oracle 1 dominates Oracle 2 in every game restricted to C_1 , although F_1 does not refine F_2 in C_1 . This contradicts Theorem 4.

Moving on to the second part, assume to the contrary that F_1 refines F_2 in every CKC, but Oracle 1 does not dominate Oracle 2. Therefore, there exists a strategy τ_2 such that Oracle 1 cannot produce the same distribution over posteriors as τ_2 .²⁰ The proof now splits to 4 steps.

Step 1: Mimicking sub-strategies.

We start by defining the notion of a sub-strategy, that resembles a strategy, but with induced probabilities that may sum to less than 1. Formally, a *partial distribution* $\tilde{p}$ is a non-negative function from a finite subset of S to $[0, 1]$ such that $\sum_{s \in S} \tilde{p}(s) \leq 1$. A partial distribution differs from a distribution as the probabilities need not sum to 1. Let $\tilde{\Delta}(S)$ be the set of partial

²⁰Observe that the condition that Oracle 1 can generate the same distribution over posterior profiles as Oracle 2 implies that Oracle 1 dominates Oracle 2. To see this, consider any game and any signaling strategy τ . Since the players' strategies depend on the profile of posteriors, we can then abstract away from the underlying private and public information and assume that the players play a Bayesian game with a given probability distribution over the profiles of posteriors, which can be generated by both Oracles.

distributions on S , and define a *sub-strategy* $\underline{\tau} : \Omega \rightarrow \tilde{\Delta}(S)$ as an F_1 -measurable function from Ω to the set of partial distributions on S . That is, $\underline{\tau}(s|\omega) \geq 0$ and $\sum_s \underline{\tau}(s|\omega) \leq 1$, for every ω and s . Evidently, every F_1 -measurable strategy is a sub-strategy.

For every sub-strategy $\underline{\tau}$ and every $p \in (\Delta(\Omega))^n$, let $\mathbf{P}_{\underline{\tau}}(p)$ be the probability that $\underline{\tau}$ yields the posterior p , i.e.,

$$\mathbf{P}_{\underline{\tau}}(p) = \sum_{\substack{(\omega, s): \underline{\tau}(s|\omega) > 0, \\ \text{and } (\mu_{\underline{\tau}|\omega, s}^i)_{i \in N} = p}} \mu(\omega) \underline{\tau}(s|\omega). \quad (8)$$

Similarly, define $\mathbf{P}_{\tau_2}(p)$ for every posterior p given the stated strategy τ_2 . We say that a sub-strategy $\underline{\tau}$ *mimics* τ_2 if

$$\mathbf{P}_{\underline{\tau}}(p) \leq \mathbf{P}_{\tau_2}(p), \text{ for every } p \in (\Delta(\Omega))^n. \quad (9)$$

Hence, a sub-strategy $\underline{\tau}$ mimics τ_2 if, for every posterior p , the probability that $\underline{\tau}$ generates p does not exceed the probability that τ_2 generates it. Note that the null sub-strategy (i.e., $\underline{\tau}(s|\omega) = 0$ for every ω and s) also mimics τ_2 .

Consider any sub-strategy $\underline{\tau}$ that mimics τ_2 . Because τ_2 generates a finite set $\text{Post}(\tau_2)$ of possible posteriors, there exists a finite number of combinations of posteriors (which does not exceed $2^{|\text{Post}(\tau_2)|}$) that every signal of $\underline{\tau}$ supports. So, if some sub-strategy uses more than $2^{|\text{Post}(\tau_2)|}$ signals, we can apply the pigeonhole principle to deduce that the additional signals support similar combinations of posteriors as some other signals. Therefore, for every such additional signal s , there exists another signal s' and a constant $c > 0$ such that $\underline{\tau}(s|\omega) = c \underline{\tau}(s'|\omega)$ for every ω , and we can unify the two signals into one. We can thus assume that there exists a finite set of signals $\underline{S}$, such that every mimicking sub-strategy (i.e., that mimics τ_2) uses only signals from $\underline{S}$.

Step 2: Optimal sub-strategies.

Let $A_{\underline{\tau}}$ be the set of sub-strategies that mimic τ_2 . Note that the set of sub-strategies supported on $\underline{S}$ is compact, and the (inequality) mimicking condition, $\mathbf{P}_{\underline{\tau}}(p) \leq \mathbf{P}_{\tau_2}(p)$ for every $p \in (\Delta(\Omega))^n$, remains valid when considering a converging sequence of sub-strategies. Thus, $A_{\underline{\tau}}$ is also compact.

Consider the function $H(\underline{\tau}) = \sum_{p \in \text{Post}(\tau_2)} \mathbf{P}_{\underline{\tau}}(p)$ defined from $A_{\underline{\tau}}$ to $[0, 1]$. As a piecewise linear function of τ , it is continuous, so $\underline{\tau}_{1.0} = \arg\max_{\underline{\tau} \in A_{\underline{\tau}}} H(\underline{\tau})$ is well-defined. If $H(\underline{\tau}_{1.0}) = 1$, then $\underline{\tau}_{1.0}$ is an F_1 -measurable strategy that mimics τ_2 . This contradicts the original premise (that Oracle 1 cannot induce the same distribution over posteriors as τ_2), so assume to the contrary that $\underline{\tau}_{1.0}$ is a proper sub-strategy and $H(\underline{\tau}_{1.0}) < 1$. If that is the case (i.e., if $H(\underline{\tau}_{1.0}) < 1$), there exists a posterior $p^* \in \text{Post}(\tau_2)$ so that $\mathbf{P}_{\underline{\tau}_{1.0}}(p^*) < \mathbf{P}_{\tau_2}(p^*)$.

Step 3: Partially supported and connected posteriors.

For every posterior $p \in \text{Post}(\tau_2)$, let $A_p = \{\omega \in \Omega : p^i(\omega) > 0 \text{ for some player } i\}$ be the set of states on which p is strictly positive, contained in some CKC denoted C_p . We say that a posterior $p \in \text{Post}(\tau_2)$ is *partially supported* (PS) if $\mathbf{P}_{\underline{\tau}_{1.0}}(p) < \mathbf{P}_{\tau_2}(p)$, otherwise we say that p is *fully supported* (FS). Let us now prove a few supporting claims related to PS posteriors.

Claim 1: If p is PS, then $\sum_s \underline{\tau}_{1.0}(s|\omega) < 1$ for every state $\omega \in A_p$.

Proof. Fix a posterior p and a state ω_0 such that $(\mu_{\tau|\omega_0,s}^i)_{i \in N} = p$ for some signal s and $\tau \in \{\underline{\tau}_{1.0}, \tau_2\}$. There exists a constant α_{p,ω_0} , independent of s and τ , such that $\alpha_{p,\omega_0} \mu(\omega_0) \tau(s|\omega_0) = \sum_{\omega \in A_p \setminus \{\omega_0\}} \mu(\omega) \tau(s|\omega)$. This follows from the fact that, in order to induce the posterior p , the probabilities induced by τ must maintain the same proportions along the different states in A_p , independently of either the strategy or the signal. Otherwise, the induced posterior would not match p . Thus, Equation (8) could be re-formulated as follows,

$$\begin{aligned} \mathbf{P}_{\tau}(p) &= \sum_{(\omega,s): (\mu_{\tau|\omega,s}^i)_{i \in N} = p} \mu(\omega) \tau(s|\omega) \\ &= \sum_{s: (\mu_{\tau|\omega_0,s}^i)_{i \in N} = p} \mu(\omega_0) \tau(s|\omega_0) + \sum_{\substack{(\omega,s): \omega \in A_p \setminus \{\omega_0\}, \\ \text{and } (\mu_{\tau|\omega,s}^i)_{i \in N} = p}} \mu(\omega) \tau(s|\omega) \\ &= (1 + \alpha_{p,\omega_0}) \mu(\omega_0) \sum_{s: (\mu_{\tau|\omega_0,s}^i)_{i \in N} = p} \tau(s|\omega_0), \end{aligned}$$

which translates to

$$\sum_{s: (\mu_{\tau|\omega_0,s}^i)_{i \in N} = p} \tau(s|\omega_0) = \frac{\mathbf{P}_{\tau}(p)}{(1 + \alpha_{p,\omega_0}) \mu(\omega_0)}.$$

Summing over all $p \in \text{Supp}(\tau_2)$, we get

$$\sum_s \tau(s|\omega_0) = \frac{1}{\mu(\omega_0)} \sum_{\substack{p: (\mu_{\tau|\omega_0,s}^i)_{i \in N} = p, \\ \text{for some } s}} \frac{\mathbf{P}_\tau(p)}{(1 + \alpha_{p,\omega_0})}. \quad (10)$$

Note that the RHS holds for either $\tau_{1,0}$ or τ_2 .

Now assume, by contradiction, that p_0 is a PS posterior and $\sum_s \tau_{1,0}(s|\omega_0) = 1$ for some state $\omega_0 \in A_{p_0}$. Using Equation (10), for both τ_2 and $\tau_{1,0}$, we get

$$\begin{aligned} 1 &= \sum_s \tau_2(s|\omega_0) = \frac{1}{\mu(\omega_0)} \sum_{\substack{p: (\mu_{\tau_2|\omega_0,s}^i)_{i \in N} = p, \\ \text{for some } s}} \frac{\mathbf{P}_{\tau_2}(p)}{(1 + \alpha_{p,\omega_0})} \\ 1 &= \sum_s \tau_{1,0}(s|\omega_0) = \frac{1}{\mu(\omega_0)} \sum_{\substack{p: (\mu_{\tau_{1,0}|\omega_0,s}^i)_{i \in N} = p, \\ \text{for some } s}} \frac{\mathbf{P}_{\tau_{1,0}}(p)}{(1 + \alpha_{p,\omega_0})}, \end{aligned}$$

which implies that

$$\sum_{\substack{p: (\mu_{\tau_2|\omega_0,s}^i)_{i \in N} = p, \\ \text{for some } s}} \frac{\mathbf{P}_{\tau_2}(p)}{(1 + \alpha_{p,\omega_0})} = \sum_{\substack{p: (\mu_{\tau_{1,0}|\omega_0,s}^i)_{i \in N} = p, \\ \text{for some } s}} \frac{\mathbf{P}_{\tau_{1,0}}(p)}{(1 + \alpha_{p,\omega_0})} < \sum_{\substack{p: (\mu_{\tau_2|\omega_0,s}^i)_{i \in N} = p, \\ \text{for some } s}} \frac{\mathbf{P}_{\tau_2}(p)}{(1 + \alpha_{p,\omega_0})},$$

where the strict inequality follows from the fact that $\mathbf{P}_{\tau_{1,0}}(p) \leq \mathbf{P}_{\tau_2}(p)$ for every posterior p , with a strict inequality for $p = p_0$. This yields a contradiction and the result follows. $\square$

Claim 2: If $\sum_s \tau_{1,0}(s|\omega) < 1$ for some state ω , then there exists a PS posterior p such that $\omega \in A_p$.

Proof. Assume, to the contrary, that $\sum_s \tau_{1,0}(s|\omega_0) < 1$ for some state ω_0 , and every posterior

p such that $\omega_0 \in A_p$ is FS. Using Equation (10), we deduce that

$$\begin{aligned}
1 &= \sum_s \tau_2(s|\omega_0) \\
&= \frac{1}{\mu(\omega_0)} \sum_{\substack{p: (\mu_{\tau_2}^i|\omega_0, s)_{i \in N=p}, \\ \text{for some } s}} \frac{\mathbf{P}_{\tau_2}(p)}{(1 + \alpha_{p, \omega_0})} \\
&= \frac{1}{\mu(\omega_0)} \sum_{\substack{p: (\mu_{\tau_{1.0}}^i|\omega_0, s)_{i \in N=p}, \\ \text{for some } s}} \frac{\mathbf{P}_{\tau_{1.0}}(p)}{(1 + \alpha_{p, \omega_0})} \\
&= \sum_s \tau_{1.0}(s|\omega_0) < 1,
\end{aligned}$$

where the first equality follows from the fact that τ_2 is a strategy, the second and forth equations follow from Equation (10), the third equality follows from the fact that every posterior p such that $\omega_0 \in A_p$ is FP, and the last inequality is by assumption. We thus reach a contradiction and the result follows. $\square$

We will use Claims 1 and 2 to extend $\tau_{1.0}$, and show that it cannot be a maximum of H . For this purpose we need to define the notion of connected posteriors. Formally, we say that two posteriors $p, p' \in \text{Post}(\tau_2)$ are *connected* if there exist two states $(\omega, \omega') \in A_p \times A_{p'} \subseteq C_p \times C_{p'}$, where $C_p \neq C_{p'}$ are two distinct CKCs, such that $F_1(\omega) = F_1(\omega')$. Equivalently, in such a case, we refer to C_p and $C_{p'}$ as *connected*, as well. Let (ω, ω') and $F_1(\omega)$ be the *connection* and *connecting set* of p and p' , respectively.²¹ We can now relate the notion of connected posteriors to PS ones, through the following claim.

Claim 3: Fix a PS posterior p and $\omega \in A_p$. Then, for every connection (ω, ω') , there exists a PS posterior p' such that $\omega' \in A_{p'} \cap F_1(\omega)$.

Proof. Let p be a PS posterior with a connection (ω, ω') and $F_1(\omega) = F_1(\omega')$. Using Claim 1, if p is PS, then $\sum_s \tau_{1.0}(s|\omega) < 1$ for every $\omega \in A_p$, so the F_1 -measurability constraint implies that $\sum_s \tau_{1.0}(s|\omega') < 1$. Thus, according to Claim 2, there exists a PS posterior p' such that $\omega' \in A_{p'}$, as needed. $\square$

²¹Equivalently, we refer to (ω, ω') and $F_1(\omega)$ as the *connection* and *connecting set* of the CKCs C_p and $C_{p'}$.

Step 4: Extending $\underline{\tau}_{1,0}$.

Recall that p^* is a PS posterior. Let V be the set of all CKCs C_l such that there exists a sequence of PS posteriors $(p^*, p_1, \dots, p_l)$ where every two successive posteriors are connected and $A_{p_l} \subseteq C_l$. Assume that V also contains C_{p^*} . Let $E \subseteq V^2$ be the set of couples (C, C') such that C and C' are connected, and denote by $\mathcal{P}^*$ the set of all PS connected posteriors that generate V . Clearly, (V, E) is a connected graph and we can use it to construct a sub-strategy $\underline{\tau}$ which mimics τ_2 and $\text{Post}(\underline{\tau}) = \mathcal{P}^*$. The proof proceeds by induction on the number of vertices in V .

Preliminary step: $|V| = 1$. Assume that C_{p^*} is the unique CKC in V . Because $p^* \in \text{Post}(\tau_2)$, there exists a signal s^* and state $\omega \in C_{p^*}$ such that $\tau_2(s^*|\omega) > 0$ and $(\mu_{\tau_2|\omega, s^*}^i)_{i \in N} = p^*$. Define the sub-strategy $\underline{\tau}_{1,1}(s|\omega) = \tau_2(s^*|\omega)$ for every $\omega \in A_{p^*}$. Recall that F_1 refines F_2 in every CKC, therefore $\underline{\tau}_{1,1}$ is well defined. Moreover, it is a sub-strategy that mimics τ_2 and $\text{Post}(\underline{\tau}_{1,1}) = \mathcal{P}^*$, as needed.

Induction step: $|V| = m$. Assume that for every graph (V, E) where $|V| = m$, there exists a sub-strategy $\underline{\tau}_{1,m}$ that mimics τ_2 , and $\text{Post}(\underline{\tau}_{1,m}) = \mathcal{P}^*$.

Induction proof for $|V| = m + 1$. Assume that $|V| = m + 1$. The distance between C_{p^*} and every vertex (i.e., every CKC) in V is defined by the shortest path between the two vertices. Denote by C_{m+1} the vertex in (V, E) with the longest path from C_{p^*} .

We argue that C_{m+1} has exactly one connecting set with the other vertices. Otherwise, assume that there are at least two connecting sets. If the two originate from the same CKC in V , then we get an F_1 -loop, which cannot exist. Thus, we can assume that the two sets originate from different CKCs, denoted C and C' . Since (V, E) is a connected graph, there exists a path from C_{p^*} to each of these CKCs. Consider the two sequences of connecting sets for these two paths. If the two are pairwise disjoint, then we have an F_1 -loop from C_{p^*} to C_{m+1} , which again yields a contradiction. So the sequences must coincide at some stage. Take a truncation of the sequences from the last stage in which they coincide until C_{m+1} . The origin of the two paths are connected CKCs (sharing the same connecting set), denoted C_l and C_{l+1} , so we now have two pairwise disjoint sequences between these two connected CKCs till C_{m+1} , thus generating an F_1 -loop. Therefore, we conclude that there is exactly one connecting set, denoted A , between

C_{m+1} and the other CKCs in V .

Consider a refinement of F_1 where A is partitioned into two disjoint sets, $A_1 = A \setminus C_{m+1}$ and $A_2 = A \cap C_{m+1}$. In such a case, $|V| = m$ and, according to the induction step, there exists a mimicking sub-strategy $\underline{\tau}_{1,m}$ supported on every PS connected posterior in $\mathcal{P}^*$ other than the ones related to the CKC C_{m+1} . Let p_{m+1} denote a PS posterior such that $A_2 \subset A_{p_{m+1}} \subseteq C_{m+1}$. In case there is more than one PS posterior, the proof works similarly because every additional posterior shares the same connecting set A .

According to the induction step, $\text{Post}(\underline{\tau}_{1,m}) = \mathcal{P}^* \setminus \{p_{m+1}\}$, so we need to extend this sub-strategy to support p_{m+1} as well. Since $p_{m+1} \in \text{Post}(\tau_2)$, there exists a signal, denoted s^* w.l.o.g., and states $\omega \in A_{p_{m+1}} \subseteq C_{m+1}$ such that $\tau_2(s^*|\omega) > 0$ and $(\mu_{\tau_2|\omega, s^*}^i)_{i \in N} = p_{m+1}$. Moreover, because C_{m+1} is not connected (neither directly, nor indirectly) to the other CKCs in V under the refined F_1 , we can assume that $\sum_s \underline{\tau}_{1,m}(s|A_1) > \sum_s \underline{\tau}_{1,m}(s|A_2)$. Otherwise, we can re-scale $\underline{\tau}_{1,m}$ in the different *unconnected* elements of the refined F_1 . Hence, we can also assume that there exists a signal, again denoted s^* w.l.o.g., such that $\underline{\tau}_{1,m}(s^*|A_1) > 0 = \underline{\tau}_{1,m}(s^*|A_2)$.

Define the following function

$$\underline{\tau}_{1,m+1}(s|\omega) = \begin{cases} c_m \underline{\tau}_{1,m}(s|\omega), & \text{for every } (\omega, s) \text{ s.t. } \underline{\tau}_{1,m}(s|\omega) > 0, \\ c_2 \underline{\tau}_2(s^*|\omega), & \text{for every } (\omega, s) \text{ s.t. } \omega \in A_{p_{m+1}}, s = s^*, \end{cases}$$

where the parameters $c_m > 0$ and $c_2 > 0$ are chosen to ensure that $\underline{\tau}_{1,m+1}(s^*|A_1) = \underline{\tau}_{1,m+1}(s^*|A_2)$, thus sustaining the F_1 -measurability constraint across the connecting set A , and that $\underline{\tau}_{1,m+1}$ remains a sub-strategy that mimics τ_2 (ensuring that $\sum_s \underline{\tau}(s|\omega) \leq 1$ for every s and ω and the that Inequality (9) holds). In conclusion, we constructed a sub-strategy that mimics τ_2 and whose support is $\mathcal{P}^*$, and this concludes the induction.

Let $\underline{\tau}_{1*}$ be the sub-strategy that mimics τ_2 and $\mathbf{P}_{\underline{\tau}_{1*}}(p) > 0$ if and only if $p \in \mathcal{P}^*$. Assume that $\underline{\tau}_{1*}$ only uses signals in some set S^* , that are not used by $\underline{\tau}_{1,0}$ (i.e., $S^* \cap \underline{S} = \emptyset$). Define

the following sub-strategy

$$\tau_{2.0}(s|\omega) = \begin{cases} \tau_{1.0}(s|\omega), & \text{for every } (\omega, s) \text{ s.t. } \tau_{1.0}(s|\omega) > 0, \\ c\tau_{1*}(s|\omega), & \text{for every } (\omega, s) \text{ s.t. } \tau_{1*}(s|\omega) > 0, \end{cases}$$

where c is a constant. Since $\tau_{1*}(s|\omega)$ supports only PS posteriors of $\tau_{1.0}$, for every state ω where there exists a PS posterior p of $\tau_{1*}(s|\omega)$ such that $\omega \in A_p$, it follows from Claim 1 that $\sum_{s \in \underline{S}} \tau_{1.0}(s|\omega) < 1$. Therefore, by choosing c sufficiently small, we can ensure that $\sum_{s \in \underline{S} \cup S^*} \tau_{2.0}(s|\omega) = \sum_{s \in \underline{S}} \tau_{1.0}(s|\omega) + c \sum_{s \in S^*} \tau_{1*}(s|\omega) < 1$. Hence, for the extended strategy $\tau_{2.0}(s|\omega)$, we can guarantee that for every $\omega \in \Omega$, $\sum_{s \in \underline{S} \cup S^*} \tau_{2.0}(s|\omega) \leq 1$. We conclude that $\tau_{2.0}$ is a sub-strategy that mimics τ_2 and $H(\tau_{2.0}) > H(\tau_{1.0})$ due to the extension over PS posteriors. This contradicts the definition of $\tau_{1.0}$ as a mimicking sub-strategy the maximizes H . We can thus conclude that $H(\tau_{1.0}) = 1$, and $\tau_{1.0}$ is an F_1 -measurable strategy that mimics τ_2 , as needed. $\square$

A.11 Proof of Proposition 5

Proof. **iii** $\Rightarrow$ **i**. Suppose that $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$ is not F_2 -balanced. It means that there is a partition $\{A, B\}$ s.t. $\#(A \rightarrow B) \neq \#(B \rightarrow A)$. Define

$$f(\omega) = \begin{cases} 1, & \text{if } \omega \in A, \\ 2, & \text{if } \omega \in B. \end{cases}$$

We obtain,

$$\prod_{i=1}^m \frac{f(\omega_i)}{f(\bar{\omega}_i)} = \left(\frac{1}{2}\right)^{\#(A \rightarrow B)} \cdot 2^{\#(B \rightarrow A)} \neq 1.$$

This contradicts **iii**.

i $\Rightarrow$ **ii**. Assume **i**. For every i , let $D_i = \{\omega_j; \omega_j \in F_2(\omega_i)\} \cup \{\bar{\omega}_j; \bar{\omega}_j \in F_2(\omega_i)\}$ be the set which contains all the states in the loop that share the same information set of F_2 as ω_i . Condition **i** implies that for every ω_i , the partition $A = D_i$ and $B = (D_i)^c$ satisfies $\#(A \rightarrow B) = \#(B \rightarrow A)$. Note that $|\{\omega_j; \omega_j \in F_2(\omega_i)\}| = \#(A \rightarrow B) + \#(A \rightarrow A)$, and

$|\{\bar{\omega}_j; \bar{\omega}_j \in F_2(\omega_i)\}| = \#(B \rightarrow A) + \#(A \rightarrow A)$, where $\#(A \rightarrow A) = |\{i \in \{1, \dots, m\}; \omega_i \in A, \bar{\omega}_i \in A\}|$. It follows from $\#(A \rightarrow B) = \#(B \rightarrow A)$ that

$$|\{\omega_j; \omega_j \in F_2(\omega_i)\}| = |\{\bar{\omega}_j; \bar{\omega}_j \in F_2(\omega_i)\}| \quad (11)$$

for every ω_i .

Define $J = \{i; \omega_i \in F_2(\bar{\omega}_i)\}$. We show the rest of the states are decomposed into F_2 -loops. Specifically, we show that if a finite set $S = \{(\omega_j, \bar{\omega}_j); \bar{\omega}_j \notin F_2(\omega_j)\}$, not necessarily an F_1 -loop, satisfies Eq. (11) for every $\omega_i \in S$, then it is covered by F_2 -loops.

When $|S| = 2$, Eq. (11) implies that this is an F_2 -loop. We now assume the induction hypothesis: if Eq. (11) is satisfied for a set $S = \{(\omega_j, \bar{\omega}_j)\}$ and for every $\omega_i \in S$, and S contains less than or equal to m pairs, then it is covered by F_2 -loops. We proceed by showing this statement for sets S containing $m + 1$ pairs.

We start at an arbitrary pair, say $(\omega_1, \bar{\omega}_1)$, and show that it belongs to an F_2 -loop. Once this F_2 -loop is formed, the states outside of this loop satisfy Eq. (11) for every ω_i outside of this loop. By the induction hypothesis, this set is covered by F_2 -loops.

Due to Eq. (11), there is at least one $\bar{\omega}_j$ such that $\bar{\omega}_j \in F_2(\omega_1)$. Consider now the two pairs, $(\omega_j, \bar{\omega}_j, \omega_1, \bar{\omega}_1)$. If this is a loop, Eq. (11) remains true when applied to the states out of this loop. The induction hypothesis completes the argument. Otherwise, there is $\bar{\omega}_k$ where $k \neq 1, j$, such that $\bar{\omega}_k \in F_2(\omega_j)$. Consider now the three pairs, $(\omega_k, \bar{\omega}_k, \omega_j, \bar{\omega}_j, \omega_1, \bar{\omega}_1)$. If this is an F_2 -loop, the other states satisfy Eq. (11), and as before, this set is covered by F_2 -loops. However, if this is not an F_2 -loop, Eq. (11) remains true, we annex another pair and continue this way until we obtain an F_2 -loop. This loop might cover the entire set, but if not, the remaining states are, by the induction hypothesis, covered by F_2 -loops. This shows **ii**.

ii $\Rightarrow$ **iii**. Let $f : \{\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m\} \rightarrow (0, \infty)$ be a positive and F_2 -measurable function. Suppose that $I_1, \dots, I_r$ is a partition of $\{1, \dots, m\}$, and for each $t = 1, \dots, r$, the set $\left((\omega_i, \bar{\omega}_i)\right)_{i \in I_t}$ is an F_2 -loop. Since, $\left((\omega_i, \bar{\omega}_i)\right)_{i \in I_t}$ is an F_2 -loop,

$$\prod_{i \in I_t} \frac{f(\omega_i)}{f(\bar{\omega}_i)} = 1,$$

which implies that

$$\prod_{i=1}^m \frac{f(\omega_i)}{f(\bar{\omega}_i)} = \prod_{t=1}^r \prod_{i \in I_t} \frac{f(\omega_i)}{f(\bar{\omega}_i)} = 1.$$

This proves **iii**. □

A.12 Proof of Proposition 6

Proof. Fix an F_i -loop $L_i = \left((\omega_j, \bar{\omega}_j) \right)_{j \in I}$ where $I = \{1, 2, \dots, m\}$. Let C_j denote the CKC that contains every pair $(\omega_j, \bar{\omega}_j)$.

Proof for first statement: Assume that L_i intersects the same CKC at least twice, so that $C_{l_1} = C_{l_2}$, where $l_1 < l_2$, is such CKC. Because L_i is a loop, the two pairs $(\omega_{l_1}, \bar{\omega}_{l_1})$ and $(\omega_{l_2}, \bar{\omega}_{l_2})$ that are in this CKC cannot be adjacent in the loop L_i , i.e., $l_1 \neq l_2 \pm 1$. Define the following sub-loop of L_i by omitting every state from $\bar{\omega}_{l_1}$ to ω_{l_2} . Formally, $L'_i = (\omega_1, \bar{\omega}_1, \dots, \bar{\omega}_{l_1-1}, \omega_{l_1}, \bar{\omega}_{l_2}, \omega_{l_2+1}, \dots, \omega_m, \bar{\omega}_m)$. This is a well-defined sub-loop of L_i (as $\omega_{l_1}, \bar{\omega}_{l_2} \in C_{l_1}$ while all other parts of the sub-loop match those of L_i), which implies that L_i is not irreducible. Note that the part we truncated from the loop L_i also forms a sub-loop, namely $L''_i = (\omega_{l_2}, \bar{\omega}_{l_1}, \omega_{l_1+1}, \bar{\omega}_{l_1+1}, \dots, \omega_{l_2-1}, \bar{\omega}_{l_2-1})$.

Proof for second statement: Assume, by contradiction, that L_i is irreducible, yet it has a pair of states $(\omega_l, \bar{\omega}_l)$ such that $\bar{\omega}_l \in F_i(\omega_l)$. This implies that $\{\bar{\omega}_{l-1}, \omega_l, \bar{\omega}_l, \omega_{l+1}\} \subseteq F_i(\omega_l) = F_i(\omega_{l+1})$. We can assume that $C_{l-1} \neq C_{l+1}$, otherwise the first statement suggests that L_i is not irreducible. So, define the following sub-loop of L_i by $L'_i = \left((\omega_j, \bar{\omega}_j) \right)_{j \in I \setminus \{l\}}$. Note that L'_i is a well-defined sub-loop, as $C_{l-1} \neq C_{l+1}$ and $\bar{\omega}_{l-1} \in F_i(\omega_{l+1})$, thus contradicting the irreducible property.

Proof for third statement: Assume, w.l.o.g., that $F_i(\omega_1) \neq F_i(\bar{\omega}_1)$. If L_i intersects the same CKC twice, then we can follow the proof of the first statement, truncate the loop, and take a sub-loop that has an informative pair of states and intersects every CKC at most once. Thus, w.l.o.g., assume that L_i intersect every CKC at most once. Denote the set of informative pairs by $I^c = \{j : F_i(\omega_j) \neq F_i(\bar{\omega}_j)\}$ and define the following ordered sub-loop of L_i by $L'_i = \left((\omega_j, \bar{\omega}_j) \right)_{j \in I^c}$. In simple terms, L'_i is generated from L_i by truncating all non-informative pairs $(\omega_j, \bar{\omega}_j)$, where $F_i(\omega_j) = F_i(\bar{\omega}_j)$, similarly to the process used in the proof of

the second statement. Focusing on L'_i , note that: (i) all pairs are pairwise disjoint; (ii) every CKC is crossed at most once; (iii) $\omega_{j+1} \in F_i(\bar{\omega}_j)$ as we removed only non-informative pairs; and (iv) $\omega_j \neq \bar{\omega}_j$ are both in the same CKC as in the original loop. Hence, L'_i is a well-defined loop and an F_i -fully-informative sub-loop of L_i .

Proof of fourth statement: If the loop L_i is irreducible, then the statement holds. Otherwise, it is not irreducible and we will prove by induction on the number of pairs m in L_1 . If $m = 2$, then L_i is irreducible. If $m = 3$ and L_i is not irreducible, then it has a sub-loop with two pairs. Assume w.l.o.g. that this sub-loop is based on the states $\{\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2\}$. It cannot be that $F_i(\bar{\omega}_1) = F_i(\bar{\omega}_2)$, because that would make $(\omega_2, \bar{\omega}_2)$ a non-informative pair. So the sub-loop is $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2)$ such that $F_i(\omega_1) = F_i(\bar{\omega}_2)$, but $F_i(\omega_1) = F_i(\bar{\omega}_3)$ and $F_i(\bar{\omega}_2) = F_i(\omega_3)$, so the pair $(\omega_3, \bar{\omega}_3)$ is non-informative.

Assume the statement holds for $m = k$ pairs, and consider an L_i loop with $k + 1$ pairs. If the loop intersects the same CKC more than once, we can split it to two sub-loops (as previously done), and use the induction hypothesis for each. Hence, we can assume that the loop does not intersect the same CKC twice.

Because the loop is not irreducible, there are two states ω_{i_1} and $\bar{\omega}_{i_2}$ that are not adjacent in the loop (so $i_1 \geq i_2 + 2$), yet $F_i(\omega_{i_1}) = F_i(\bar{\omega}_{i_2})$. The last equality also suggests that $F_i(\bar{\omega}_{i_1-1}) = F_i(\omega_{i_2+1})$. If $i_1 = i_2 + 2$, then there exists only one pair between the two states. This implies that the pair $(\omega_{i_2+1}, \bar{\omega}_{i_2+1}) = (\omega_{i_1-1}, \bar{\omega}_{i_1-1})$ is non-informative, contradicting the fact that L_i is F_i -fully-informative. So we conclude that $i_1 \geq i_2 + 3$. Define the following two loops $L'_i = (\omega_{i_1}, \bar{\omega}_{i_1}, \dots, \omega_{i_2}, \bar{\omega}_{i_2})$ and $L''_i = (\omega_{i_2+1}, \bar{\omega}_{i_2+1}, \dots, \omega_{i_1-1}, \bar{\omega}_{i_1-1})$, where the ordering of states follows the original loop L_i . These are two well-defined F_i -loops with less than $k + 1$ pairs each, so the induction hypothesis holds and the result follows.

If L_i does not intersect the same CKC more than once and does not have at least 4 states in the same partition element, then it is irreducible.

Proof of fifth statement: If the loop has a non-informative pair $\omega_j \in F_i(\bar{\omega}_i)$, then it contains 4 states from the same partition element, so assume that the loop is F_i -fully-informative and that it does not intersect the same CKC more than once. Thus, we need to prove that it has at least 4 states in the same partition element of F_i .

Consider the strict sub-loop L_i^- of L_i . It consists of pairs, taken from the original loop. Because L_i does not intersect the same CKC more than once, all the pairs of L_i^- are a strict subset of the pairs of L_i . This implies that some pairs were omitted from L_i when generating L_i^- , so assume w.l.o.g. that the pair $\{\omega_1, \bar{\omega}_1\}$ is not included in L_i^- . This implies that one pair $\{\omega_j, \bar{\omega}_j\}$ precedes in L_i^- a different one that it precedes in L_i . That is, $F_i(\bar{\omega}_j) = F_i(\omega_{j+1})$ according to L_i , whereas $F_i(\bar{\omega}_j) = F_i(\omega_k)$ where $k \neq j+1$, according to L_i^- . But also $F_i(\omega_k) = F_i(\bar{\omega}_{k-1})$ according to L_i . Thus, $\{\bar{\omega}_j, \omega_{j+1}, \omega_k, \bar{\omega}_{k-1}\}$ are in the same partition element of L_i , as stated and the result follows. $\square$

A.13 Proof of Theorem 6

Proof. Suppose that Oracle 1 dominates Oracle 2. If there exists a CKC in which F_1 does not refine F_2 , Theorem 4 states that Oracle 1 does not dominate Oracle 2 in that CKC. In other words, there exists τ_2 defined on this CKC, such that for every τ_1 , it follows that $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$. We extend the definition of τ_2 to the entire state space in an arbitrary way, and still for every τ_1 , it follows that $\text{Post}(\tau_1) \not\subseteq \text{Post}(\tau_2)$, and we can use Proposition 3 accordingly.

We proceed to show that any F_1 -loop is F_2 -balanced, which is equivalent to the existence of a cover by loops of F_2 . Suppose, to the contrary, that an F_1 -loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$ is not F_2 -balanced. This means that there is an F_2 -measurable partition $\{A, B\}$ of these states such that Eq. (2) is not satisfied. We define an F_2 -measurable signaling function that obtains two signals, α and β . Over the states of the loop, let

$$\tau_2(\alpha|\omega) = \begin{cases} x, & \text{if } \omega \in A, \\ y, & \text{if } \omega \in B, \end{cases} \quad (12)$$

and $\tau_2(\beta|\omega) = 1 - \tau_2(\alpha|\omega)$. On other states, τ_2 is defined arbitrarily. The numbers $x, y \in (0, 1)$ are chosen so that $\frac{\ln x - \ln y}{\ln(1-x) - \ln(1-y)}$ is irrational.

Claim 1: If $\text{Post}(\tau_1) \subseteq \text{Post}(\tau_2)$, then any signal of τ_1 induces the same posteriors as α does or as β does in every CKC.

Claim 2: For any signal s of τ_1 and for any i , $\frac{\tau_1(s|\omega_i)}{\tau_1(s|\bar{\omega}_i)} \in \{\frac{x}{y}, \frac{1-x}{1-y}, \frac{y}{x}, \frac{1-y}{1-x}\}$. Therefore,

$$\prod_{i=1}^m \frac{\tau_1(s|\omega_i)}{\tau_1(s|\bar{\omega}_i)} = \left(\frac{x}{y}\right)^{\ell_1} \cdot \left(\frac{1-x}{1-y}\right)^{\ell_2} \cdot \left(\frac{y}{x}\right)^{k_1} \cdot \left(\frac{1-y}{1-x}\right)^{k_2},$$

where $\ell_1 + \ell_2 = |\{i; \omega_i \in A \text{ and } \bar{\omega}_i \in B\}|$ and $k_1 + k_2 = |\{i; \omega_i \in B \text{ and } \bar{\omega}_i \in A\}|$.

Claim 3: For any signal s of τ_1 , $\prod_{i=1}^m \frac{\tau_1(s|\omega_i)}{\tau_1(s|\bar{\omega}_i)} = 1$.

We therefore obtain $(\frac{x}{y})^{\ell_1} (\frac{1-x}{1-y})^{\ell_2} (\frac{y}{x})^{k_1} (\frac{1-y}{1-x})^{k_2} = 1$. We conclude that there are whole numbers, say $\ell = \ell_1 - k_1$ and $k = k_2 - \ell_2$ such that $(\frac{x}{y})^\ell = (\frac{1-x}{1-y})^k$. Since $\frac{\ln x - \ln y}{\ln(1-x) - \ln(1-y)} = \frac{\ln \frac{x}{y}}{\ln \frac{1-x}{1-y}}$ is irrational, $\ell = k = 0$, implying that Eq. (2) is satisfied. This is a contradiction, so every F_1 -loop is F_2 -balanced.

Moving on to the third part of the theorem, fix an irreducible F_1 -loop L_1 , and consider an irreducible cover by a unique F_2 -loop L_2 , i.e., L_2 covers L_1 and both are irreducible w.r.t. the relevant partition. Note that if L_2 is also order-preserving, it implies that it *matches* L_1 .

Assume, by contradiction, that L_2 is not order-preserving and the two loops do not match one another. Denote $L_1 = (\omega_1, \bar{\omega}_1, \dots, \omega_m, \bar{\omega}_m)$ and $L_2 = (\omega_1, \bar{\omega}_1, \omega_{i_2}, \bar{\omega}_{i_2}, \dots, \omega_{i_m}, \bar{\omega}_{i_m})$. Thus, there exist indices $k > j > 1$ such that ω_k precedes ω_j in L_2 . In simple terms, it implies that though L_2 consists of the same pairs as L_1 , the ordering of pairs throughout the two loops differs, as suggested in Footnote 16.

Since the two loops are irreducible, it follows from Proposition 6 that they intersect every CKC at most once and that both are fully-informative. Moreover, for every state ω in every loop L_i , every set $F_i(\omega)$ contains two states from the loop L_i (otherwise, the loop is not irreducible). So, one can define an F_i -measurable function τ_i such that $\tau_i(s|\omega_l) = \tau_i(s|\bar{\omega}_{l-1}) \neq \tau_i(s|\omega_{l'})$ for every $\omega_l \neq \omega_{l'}$ in the loop.

To simplify the exposition, partition the states of L_2 into three disjoint sets: the set $A_1^2 = \{\bar{\omega}_1, \dots, \omega_k\}$ contains all the states of L_2 from $\bar{\omega}_1$ till ω_k (following the order of L_2), $A_k^2 = \{\bar{\omega}_k, \dots, \omega_j\}$ contains all the states of L_2 from $\bar{\omega}_k$ till ω_j , and $A_j^2 = \{\bar{\omega}_j, \dots, \omega_1\}$ which contains all remaining states of L_2 . Follow a similar process with L_1 , so that $A_1^1 = \{\bar{\omega}_1, \dots, \omega_j\}$ contains all the states of L_1 from $\bar{\omega}_1$ till ω_j (following the order of L_1), $A_j^1 = \{\bar{\omega}_j, \dots, \omega_k\}$ contains all the states of L_1 from $\bar{\omega}_j$ till ω_k , and $A_k^1 = \{\bar{\omega}_k, \dots, \omega_1\}$ which contains all remaining states of

L_1 .

Denote by C_l the CKC of the pair $(\omega_l, \bar{\omega}_l)$. Fix two distinct signals s_1 and s_2 , and define the signaling function τ_2 as follows:

$$\tau_2(s_1|\omega) = 1 - \tau_2(s_2|\omega) = \begin{cases} p_1, & \text{if } \omega \in A_1^2 = \{\bar{\omega}_1, \dots, \omega_k\}, \\ p_2, & \text{if } \omega \in A_k^2 = \{\bar{\omega}_k, \dots, \omega_j\}, \\ p_3, & \text{if } \omega \in A_j^2 = \{\bar{\omega}_j, \dots, \omega_1\}, \\ p_4, & \text{if } \omega \in \Omega \setminus \bigcup_{i=1,j,k} A_i^2, \end{cases}$$

where the probabilities $\{p_1, p_2, p_3, p_4\}$ are chosen as in the strategy defined in Equation (1). Because the loop is irreducible, intersects every CKC at most once and F_2 -fully-informative, τ_2 is a well-defined F_2 -measurable function.

The result of Lemma 2 holds in every CKC of the loop (though with different probabilities). So given a CKC C_l , if there exists τ_1 such that $\text{Post}(\tau_1) \subseteq \text{Post}(\tau_2)$, 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 $\omega \in C_l$. Therefore, in every CKC C_l and for every signal t , there exists a signal s such that $\frac{\tau_2(s|\omega_l)}{\tau_2(s|\bar{\omega}_l)} = \frac{\tau_1(t|\omega_l)}{\tau_1(t|\bar{\omega}_l)}$. Fix such a strategy τ_1 .

Notice that in every CKC $C_l \neq C_1, C_j, C_k$ and for every signal $s \in \{s_1, s_2\}$, we get $\tau_2(s|\omega_l) = \tau_2(s|\bar{\omega}_l)$. Thus, $\frac{\tau_1(t|\omega_l)}{\tau_1(t|\bar{\omega}_l)} = 1$ for every t and every $l \neq i, j, k$. This implies that for every feasible signal t restricted to the loop L_1 ,

$$\tau_1(t|\omega) = \begin{cases} a_t, & \text{if } \omega \in A_1^1 = \{\bar{\omega}_1, \dots, \omega_j\}, \\ b_t, & \text{if } \omega \in A_j^1 = \{\bar{\omega}_j, \dots, \omega_k\}, \\ c_t, & \text{if } \omega \in A_k^1 = \{\bar{\omega}_k, \dots, \omega_1\}, \end{cases}$$

where $a_t, b_t, c_t \in (0, 1]$. Evidently, the parameters a_t, b_t and c_t can vary across the feasible signals.

In addition, Lemma 2 states that in every CKC, $\tau_1(t|\omega)$ is proportional to $\tau_2(s_i|\omega)$ for some

signal $s_i \in \{s_1, s_2\}$. This yields the following constraints:

$$\begin{aligned}\frac{\tau_1(t|\omega_1)}{\tau_1(t|\bar{\omega}_1)} &= \frac{c_t}{a_t} = \frac{\tau_2(s_i|\omega_1)}{\tau_2(s_i|\bar{\omega}_1)} \in \left\{ \frac{p_3}{p_1}, \frac{1-p_3}{1-p_1} \right\}, \\ \frac{\tau_1(t|\omega_j)}{\tau_1(t|\bar{\omega}_j)} &= \frac{a_t}{b_t} = \frac{\tau_2(s_i|\omega_j)}{\tau_2(s_i|\bar{\omega}_j)} \in \left\{ \frac{p_2}{p_3}, \frac{1-p_2}{1-p_3} \right\}, \\ \frac{\tau_1(t|\omega_k)}{\tau_1(t|\bar{\omega}_k)} &= \frac{b_t}{c_t} = \frac{\tau_2(s_i|\omega_k)}{\tau_2(s_i|\bar{\omega}_k)} \in \left\{ \frac{p_1}{p_2}, \frac{1-p_1}{1-p_2} \right\}.\end{aligned}$$

Because the two loops cover one another and specifically because L_2 is F_1 -covered, Proposition 5 states that $\prod_{l=1}^m \frac{\tau_1(t|\omega_{i_l})}{\tau_1(t|\bar{\omega}_{i_l})} = 1$, which leaves only two possibilities for the ratios $\{\frac{c_t}{a_t}, \frac{a_t}{b_t}, \frac{b_t}{c_t}\}$ above: either they equal $\{\frac{p_3}{p_1}, \frac{p_2}{p_3}, \frac{p_1}{p_2}\}$ respectively, or $\{\frac{1-p_3}{1-p_1}, \frac{1-p_2}{1-p_3}, \frac{1-p_1}{1-p_2}\}$. This follows from the uniqueness of the ratios, as stated in Lemma 2. Note that this must hold for every feasible signal t of τ_1 across the loop.

$\tau_1(t \omega)$	t_1	t_2
ω_1	$\lambda_1 c_1$	$\lambda_2 c_2$
$\bar{\omega}_1$	$\lambda_1 a_1$	$\lambda_2 a_2$
ω_j	$\lambda_1 a_1$	$\lambda_2 a_2$
$\bar{\omega}_j$	$\lambda_1 b_1$	$\lambda_2 b_2$
ω_k	$\lambda_1 b_1$	$\lambda_2 b_2$
$\bar{\omega}_k$	$\lambda_1 c_1$	$\lambda_2 c_2$

Figure 27: The structure of τ_1 restricted to the states $\{\omega_1, \bar{\omega}_1, \omega_j, \bar{\omega}_j, \omega_k, \bar{\omega}_k\}$, where $\frac{c_1}{a_1} = \frac{p_3}{p_1}$, $\frac{b_1}{c_1} = \frac{p_1}{p_2}$, $\frac{c_2}{a_2} = \frac{1-p_3}{1-p_1}$ and $\frac{b_2}{c_2} = \frac{1-p_1}{1-p_2}$ and $\lambda_1, \lambda_2 > 0$.

Thus, if we focus on the states $\{\omega_1, \bar{\omega}_1, \omega_j, \bar{\omega}_j, \omega_k, \bar{\omega}_k\}$ and group together all signals t with the same distribution on these states, then for some positive constants $\lambda_1, \lambda_2 > 0$ we get the strategy defined in Figure 27. Plugging in the relevant ratios yields the probabilities given in Figure 28.

Recall that the rows must sum to 1, so that τ_1 is a well-defined strategy. So, we get the

$\tau_1(t \omega)$	t_1	t_2
ω_1	$\lambda_1 c_1$	$\lambda_2 c_2$
$\bar{\omega}_1$	$\lambda_1 c_1 \frac{p_1}{p_3}$	$\lambda_2 c_2 \frac{1-p_1}{1-p_3}$
ω_j	$\lambda_1 c_1 \frac{p_1}{p_3}$	$\lambda_2 c_2 \frac{1-p_1}{1-p_3}$
$\bar{\omega}_j$	$\lambda_1 c_1 \frac{p_1}{p_2}$	$\lambda_2 c_2 \frac{1-p_1}{1-p_2}$

Figure 28: The structure of τ_1 restricted to the states $\{\omega_1, \bar{\omega}_1, \omega_j, \bar{\omega}_j\}$, where probabilities are presented in terms of c_1, c_2, λ_1 and λ_2 .

following system of linear equations, in which $(x, y) = (\lambda_1 c_1, \lambda_2 c_2)$ and:

$$\begin{aligned}
x + y &= 1, \\
\frac{p_1}{p_3}x + \frac{1-p_1}{1-p_3}y &= 1, \\
\frac{p_1}{p_2}x + \frac{1-p_1}{1-p_2}y &= 1,
\end{aligned}$$

which does not have a solution since p_1, p_2, p_3 are required to be distinct. Thus, we conclude that the loops must sustain the same ordering of pairs, and therefore coincide as needed. This concludes the third and final part of the theorem. $\square$

A.14 Proof of Theorem 7

Proof. We first define an auxiliary set $\bar{\Omega}$, which groups together states that are in the same partition element of F_2 within CKCs. Formally, define the set $\bar{\Omega}$ such that $\eta(\omega') \in \bar{\Omega}$ if and only if $\eta(\omega') = \{\omega \in \Omega : \omega, \omega' \in C_j, F_2(\omega) = F_2(\omega')\}$. Accordingly, define the partition $\bar{F}_2$ to be discrete in every CKC, such that $\bar{F}_2(\eta(\omega)) = \bar{F}_2(\eta(\omega'))$ if and only if $F_2(\omega) = F_2(\omega')$. Note that $\bar{F}_2$ is essentially a projection of F_2 onto $\bar{\Omega}$. In addition, $\bar{F}_1$ is defined as follows: (i) discrete in every CKC, similarly to $\bar{F}_2$; (ii) $\bar{F}_1(\eta(\omega)) = \bar{F}_1(\eta(\omega'))$ if ω and ω' are not in the same CKC, and there exist $\bar{\omega} \in \eta(\omega)$ and $\bar{\omega}' \in \eta(\omega')$ such that $F_1(\bar{\omega}) = F_1(\bar{\omega}')$; and (iii) $\bar{F}_1$ forms a partition (i.e., given (i) and (ii), if two elements of $\bar{F}_1$ contain the same state $\eta(\omega)$, they are unified into one element).

We now prove that $\bar{F}_1 = \bar{F}_2$ in every CKC and that there are no $\bar{F}_1$ -loops. Thus, by Theorem 5, any $\bar{F}_2$ -measurable strategy $\bar{\tau}_2$ (which, extended to Ω , is also F_2 -measurable) can be imitated by an $\bar{F}_1$ -measurable strategy $\bar{\tau}_1$.

Step 1: $\overline{F}_1 = \overline{F}_2$ in every CKC.

By definition, $\overline{F}_2$ refines $\overline{F}_1$, so we need to prove that $\overline{F}_1$ also refines $\overline{F}_2$ in every CKC. Assume, by contradiction, that $\overline{F}_1(\eta(\omega)) = \overline{F}_1(\eta(\omega'))$ where ω and ω' are in the same CKC, whereas $\overline{F}_2(\eta(\omega)) \neq \overline{F}_2(\eta(\omega'))$. This suggests that $F_2(\omega) \neq F_2(\omega')$, which implies that $F_1(\omega) \neq F_1(\omega')$. According to the construction of $\overline{F}_1$, we conclude that the equality $\overline{F}_1(\eta(\omega)) = \overline{F}_1(\eta(\omega'))$ followed from the partition-formation stage described in (iii) above, through at least one other CKC. Thus, there exists an F_1 -loop which connects a state in $\eta(\omega)$ with a state in $\eta(\omega')$. Without loss of generality, assume these states are ω and ω' . Because every F_1 -loop is F_2 -non-informative, it follows that $F_2(\omega) = F_2(\omega')$, a contradiction.

Step 2: There are no $\overline{F}_1$ -loops.

An $\overline{F}_1$ -loop implies that an F_1 -loop exists. By construction, all Ω states in every CKC are F_2 -equivalent (i.e., grouped together according to F_2). Because every F_1 -loop is F_2 -non-informative, it implies that the loop consists of only one $\overline{\Omega}$ state in every CKC, and not two. This contradicts the definition of a loop.

Step 3: $\overline{F}_1$ can mimic $\overline{F}_2$.

Fix a strategy τ_2 , and let $\overline{\tau}_2$ be the projected strategy on $\overline{\Omega}$. Because $\overline{F}_1 = \overline{F}_2$ in every CKC and there are no $\overline{F}_1$ -loops, there exists an $\overline{F}_1$ -measurable strategy $\overline{\tau}_1$ that imitates $\overline{\tau}_2$. Therefore, one can lift $\overline{\tau}_1$ to Ω to create τ_1 , whose projection onto $\overline{\Omega}$ matches $\overline{\tau}_1$. Thus, the strategy τ_1 imitates τ_2 , as needed. $\square$

A.15 Proof of Proposition 7

Proof. Denote the two CKCs by C_1 and C_2 . One part of the statement follows directly from Theorem 6, so assume that F_1 refines F_2 in every CKC and any F_1 -loop is F_2 -balanced. If there are no F_1 -loops, then the result follows from Theorem 5, so assume there exists at least one F_1 -loop, and every such loop is F_2 -balanced.

Take any F_1 -loop $(\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2)$ with four states. We argue that either it is also an F_2 -loop or it is F_2 -non-informative. Otherwise, we can assume (without loss of generality) that $F_2(\omega_1) \neq F_2(\overline{\omega}_i)$, for every $i = 1, 2$. So, there are only two possibilities left: either $F_2(\omega_1) = F_2(\omega_2)$ or $F_2(\omega_1) \neq F_2(\omega_2)$. If $F_2(\omega_1) = F_2(\omega_2)$, then there exists an F_2 -measurable partition of the four

states such that $A = \{\omega_1, \omega_2\}$ and $B = \{\overline{\omega_1}, \overline{\omega_2}\}$, which is not balanced. Otherwise, there exists another non-balanced F_2 -measurable partition of the form $A = \{\omega_1\}$ and $B = \{\overline{\omega_1}, \omega_2, \overline{\omega_2}\}$. In any case, we get a contradiction.

The proof now splits into two cases: either there exists an F_1 -loop $(\omega_1, \overline{\omega_1}, \omega_2, \overline{\omega_2})$ and an index i such that $F_2(\omega_i) \neq F_2(\overline{\omega_i})$, or every such loop is F_2 -non-informative. If indeed every such loop is F_2 -non-informative, Theorem 7 states that Oracle 1 dominates Oracle 2, so we need only focus on the former.

Assume that there exists an F_1 -loop $(\omega_1, \overline{\omega_1}, \omega_2, \overline{\omega_2})$ and an index i such that $F_2(\omega_i) \neq F_2(\overline{\omega_i})$. Denote this couple by $\{\omega_1, \overline{\omega_1}\} \subseteq C_1$. The previous conclusion implies that it is also an F_2 -loop. We claim that, under these conditions, every τ_2 is F_1 -measurable. Note that F_1 refines F_2 in every CKC, so we need to verify that for every $(\omega, \overline{\omega}) \in C_1 \times C_2$ such that $F_1(\omega) = F_1(\overline{\omega})$, it follows that $F_2(\omega) = F_2(\overline{\omega})$.

Take $(\omega, \overline{\omega}) \in C_1 \times C_2$ such that $F_1(\omega) = F_1(\overline{\omega})$. If $\omega = \omega_1$ or $\omega = \overline{\omega_1}$, then $(\omega, \overline{\omega})$ are part of the previously stated F_2 -loop, so $F_2(\omega) = F_2(\overline{\omega})$. Otherwise, we can construct two new F_1 -loops $(\omega, \overline{\omega}, \omega_1, \overline{\omega_2})$ and $(\omega, \overline{\omega}, \omega_2, \overline{\omega_1})$. Because $F_2(\omega_1) \neq F_2(\overline{\omega_1})$, either $F_2(\omega) \neq F_2(\omega_1)$ or $F_2(\omega) \neq F_2(\overline{\omega_1})$. The previous conclusion again implies that $(\omega, \overline{\omega})$ are a part of an F_2 -loop, so $F_2(\omega) = F_2(\overline{\omega})$, as needed. $\square$

A.16 Proof of Theorem 8

Proof. We start by assuming that F_1 and F_2 are equivalent. According to Theorem 6, every F_i refines F_{-i} in every CKC, and every F_i -loop is covered by F_{-i} -loops. Fix an irreducible F_i -loop with at least 6 states, denoted L_i , and consider a cover by F_{-i} -loops. There are two possibilities: either the cover constitutes a single loop, or else. If the cover contains a shorter loop, say L'_{-i} , then that loop is not F_i -covered because L_i is irreducible, and this contradicts Theorem 6. Moreover, the cover cannot have non-informative pairs where $F_{-i}(\omega_i) = F_{-i}(\overline{\omega_i})$, because the two partitions match one another in every CKC and L_i is irreducible. So, the cover consists of a single irreducible F_{-i} -loop, and Theorem 6 states that it is order-preserving. Thus, L_i and L_{-i} coincide as stated.

Moving to the other direction, assume that F_i refines F_{-i} in every CKC, that any F_i -loop

has a cover of F_{-i} -loops, and every irreducible F_i -loop with at least 6 states is an irreducible F_{-i} -loop. Let us prove that Oracle 1 dominates Oracle 2 (and the reverse dominance follows symmetrically).

We start with two simple observations. First, in case F_1 has no loops, then the statement follows from previous results, so assume F_1 has loops. Second, we say that two CKCs C_1 and C_2 are connected if there exist $\omega_1 \in C_1$ and $\omega_2 \in C_2$ such that $F_1(\omega_1) = F_1(\omega_2)$. If there exists a CKC C which is not connected to any other CKC (i.e., for every $\omega \in C$, the partition element $F_1(\omega) \subseteq C$), then Oracle 1 dominates Oracle 2 conditional on that CKC and independently of all other CKCs. Thus, without loss of generality, we can assume that all CKCs are connected, either directly or sequentially.

For this part, we will need to define the notion of *type-2 irreducible* loops, which are fully-informative loops that do not have four states in the same information set of the relevant F_i .

Definition 9. *Let L_i be an F_i -loop. We say that the loop is type-2 irreducible if it does not have four states in the same information set (i.e., partition element) of F_i .*

We shall use this notion of type-2 irreducible F_1 -loops as building blocks upon which every F_2 -measurable τ_2 is also F_1 -measurable. For that purpose, we start by proving in the following Claim 2 that every type-2 irreducible F_1 -loop is also an F_2 -loop. Next, we will extend this measurability result to every set of type-2 irreducible F_1 -loops that intersect the same CKCs, and finally derive it to all CKCs that these loops intersect. This sets of CKCs, to be later defined as *clusters*, will be the basic sets upon which every F_2 -measurable strategy is also F_1 -measurable.

Claim 2. *Every type-2 irreducible F_1 -loop L_1 is an F_2 -loop.*

Proof. If L_1 is irreducible, then it is also an irreducible F_2 -loop, and the result holds. Thus assume that L_1 is not irreducible. Using the fifth result in Proposition 6, we deduce that L_1 intersects the same CKC more than once. Using the proof of the first result in Proposition 6, we can decompose L_1 to two disjoint strict sub-loops of F_1 . This can be done repeatedly, so that L_1 is decomposed to sub-loops that do not intersect the same CKC more than once. This

implies that every such loop is type-2 irreducible. Thus, every such sub-loop is irreducible, and so it is also an F_2 -loop.

Note that the decomposition process occurs *within* every relevant CKC C and that $F_1|_C = F_2|_C$. That is, once there are two pairs of the same loop within the same CKC, we can decompose the loop into two disjoint loops by rearranging these four states. So, one can reverse the process and recombine the sub-loops of F_2 to regenerate the original loop L_1 , which is now also an F_2 -loop, as needed. $\square$

Once we dealt with individual type-2 irreducible loops, we move to loops that intersect the same CKC. For that purpose, we need to prove the following supporting, general Claim 3 which states that every F_i -fully-informative loop L_i can be decomposed to type-2 irreducible F_i -loops.

Claim 3. *Every F_i -fully-informative loop L_i that is not type-2 irreducible can be decomposed to type-2 irreducible F_i -loops.*

Proof. The proof is done by induction on the number of pairs m in L_i . If $m = 2$, then it is irreducible, as needed. Assume that the statement holds for $m = k$, and consider a loop with $k + 1$ pairs. If it is not type-2 irreducible, then it has four different states $\{\bar{\omega}_j, \omega_{j+1}, \bar{\omega}_l, \omega_{l+1}\}$ in the same information set of F_i , where $l > j+1$ and $l+1 < j$ so that the two pairs are not adjacent in the original loop L_i (otherwise, the loop has a non-informative pair). Note that additional connection may exists, but in any case ω_{j+1} is in the same partition element as $\bar{\omega}_j$, and the same holds for $\bar{\omega}_l$ and ω_{l+1} . Consider the loops $(\omega_j, \bar{\omega}_j, \omega_{l+1}, \bar{\omega}_{l+1}, \omega_{l+2}, \bar{\omega}_{l+2}, \dots, \omega_{j-1}, \bar{\omega}_{j-1})$ and $(\omega_l, \bar{\omega}_l, \omega_{j+1}, \bar{\omega}_{j+1}, \omega_{j+2}, \bar{\omega}_{j+2}, \dots, \omega_{l-1}, \bar{\omega}_{l-1})$. The two sub-loops are based on the original loop, other than the first pair, see Figure 29




Figure 29: A fully-informative loop that is not type-2 irreducible, with four states in the same information set of F_i . The red rectangle denotes the same partition element of F_i , and the green edges denote the additional states of the original loop.

Each of these sub-loops is F_i -fully-informative, and have strictly less than k pairs. Thus, the induction hypothesis holds, and they are either type-2 irreducible, or can be separately decomposed to type-2 irreducible loops, so the result follows.

Note that even without the induction hypothesis, we can repeat the decomposition process, so that all the connections of the original loop that are based on information sets of F_i with no more than two states (in the loop) are kept in one of the sub-loops. $\square$

Using Claim 3, we now prove in the following Claim 4, that every F_2 -measurable strategy on two type-2 irreducible F_1 -loops with a joint CKC (i.e., pass through the same CKC) is F_1 -measurable.

Claim 4. *Fix two type-2 irreducible F_1 -loops L_1 and L'_1 that share at least one CKC. Then, every $\tau_2|_{L_1 \cup L'_1}$ is F_1 -measurable.*

Proof. Fix two type-2 irreducible F_1 -loop L_1 and L'_1 , and assume that they share at least one CKC. Denote $L_1 = (\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_m, \bar{\omega}_m)$ and $L'_1 = (\omega'_1, \bar{\omega}'_1, \omega'_2, \bar{\omega}'_2, \dots, \omega'_{m'}, \bar{\omega}'_{m'})$. Assume, by contradiction, that there exists a strategy $\tau_2|_{L_1 \cup L'_1}$ which is not F_1 -measurable. As already proven, each of these loops is also an F_2 -loop, so the measurability constraint implies that there exist $\omega \in L_1$ and $\omega' \in L'_1$ such that $F_2(\omega) \neq F_2(\omega')$ whereas $F_1(\omega) = F_1(\omega')$. Because F_1 and F_2 match one another in every CKC, this suggests that ω and ω' are in two different CKCs. Denote a shared CKC by C_j in which there are the pairs $(\omega_j, \bar{\omega}_j)$ and $(\omega'_j, \bar{\omega}'_j)$ taken from L_1

and L'_1 respectively. Note that the two pairs may coincide, as well as contain one of the states ω and ω' , but not both (because the two are in different CKCs). See Figure 30




Figure 30: Two type-2 irreducible loops of F_1 that share at least one CKC.

Let us now compose a type-2 irreducible F_1 loop, using the fact that $F_1(\omega) = F_1(\omega')$. Without loss of generality, assume that $\omega = \omega_1$ and $\omega' = \bar{\omega}'_1$, and that ω_1 is not in C_j . Moreover, it cannot be the case that ω_1 and $\bar{\omega}'_1$ are both in the same loop, say L_1 , because L_1 is also an F_2 -loop and that would imply that either $F_2(\omega) = F_2(\omega')$ in case $\bar{\omega}'_1 = \bar{\omega}_m$, or that L_1 is not a type-2 irreducible loop in case $\bar{\omega}'_1 \neq \bar{\omega}_m$. Also, it must be that $F_1(\bar{\omega}'_1) = F_1(\omega^*)$ where $\omega^* \in L_1$ if and only if $\omega^* \in \{\omega_1, \bar{\omega}_m\}$, otherwise L_1 is not type-2 irreducible.

We now split the proof to four possibilities:

- $\bar{\omega}'_1 \in C_j$.
- $\bar{\omega}'_1 \notin C_j$ and $|\{\omega_j, \bar{\omega}_j\} \cap \{\omega'_j, \bar{\omega}'_j\}| = 0, 1, 2$.

Assume that $\bar{\omega}'_1 \in C_j$. Consider the loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_j, \bar{\omega}'_1)$. This loop matches L_1 up to state ω_j and $F_1(\omega_1) = F_1(\bar{\omega}'_1)$. Thus, it is a well-defined type-2 irreducible F_1 -loop, hence also an F_2 -loop. Therefore, $F_2(\omega_1) = F_2(\bar{\omega}'_1)$ and we reach a contradiction.

Moving on to the next possibility, assume that $\bar{\omega}'_1 \notin C_j$ and $|\{\omega_j, \bar{\omega}_j\} \cap \{\omega'_j, \bar{\omega}'_j\}| = 0$. Consider the loop $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \omega_j, \bar{\omega}'_j, \omega'_{j+1}, \bar{\omega}'_{j+1}, \dots, \omega'_1, \bar{\omega}'_1)$. If ω_j and $\bar{\omega}'_j$ are in different partition elements of F_1 , then this is a well-defined F_1 -fully-informative loop. If the two states are in the same partition element, then we can omit this pair from the loop and get a shorter

loop (in terms of pairs). This process could be done repeatedly, until we get a well-defined F_1 -fully-informative loop which starts with ω_1 and ends with $\bar{\omega}'_1$. If it is a type-2 irreducible F_1 -loop, then it is also an F_2 -loop, and $F_2(\omega_1) = F_2(\bar{\omega}'_1)$. Thus, assume that it is not type-2 irreducible, which implies that it has at least four states in the same partition element of F_1 . These four states include neither ω_1 nor $\bar{\omega}'_1$, because that would imply that either L_1 or L'_1 is not type-2 irreducible. Now we can apply Claim 3, to decompose this F_1 -fully-informative loop to type-2 irreducible F_1 -loops, where at least one maintains the connection between ω_1 nor $\bar{\omega}'_1$ (see the comment at the end of the proof of Claim 3). We thus conclude that it is also an F_2 -loop and $F_2(\omega_1) = F_2(\bar{\omega}'_1)$.

The next possibility is that $\bar{\omega}'_1 \notin C_j$ and $|\{\omega_j, \bar{\omega}_j\} \cap \{\omega'_j, \bar{\omega}'_j\}| = 1$. If either $\omega'_j \in \{\omega_j, \bar{\omega}_j\}$ or $\bar{\omega}'_j = \bar{\omega}_j$, then we can follow a similar proof as in the previous case where $|\{\omega_j, \bar{\omega}_j\} \cap \{\omega'_j, \bar{\omega}'_j\}| = 0$, so assume that $\bar{\omega}'_j = \omega_j$. In that case, we can re-define the previous loop by omitting ω_j and $\bar{\omega}'_j$ to get $(\omega_1, \bar{\omega}_1, \omega_2, \bar{\omega}_2, \dots, \bar{\omega}_{j-1}, \omega'_{j+1}, \bar{\omega}'_{j+1}, \dots, \omega'_1, \bar{\omega}'_1)$. Again, this is either a well-defined F_1 -fully-informative loop, or could be reduced to such a loop. Applying the same arguments as before, we conclude that there exists a type-2 irreducible F_1 -loop which maintains the connection between ω_1 nor $\bar{\omega}'_1$, so it is also an F_2 -loop and $F_2(\omega_1) = F_2(\bar{\omega}'_1)$.

The last possibility is that $\bar{\omega}'_1 \notin C_j$ and $|\{\omega_j, \bar{\omega}_j\} \cap \{\omega'_j, \bar{\omega}'_j\}| = 2$, but in that case the analysis in the previous possibilities holds, and we reach the same conclusion that $F_2(\omega_1) = F_2(\bar{\omega}'_1)$, as needed.²² $\square$

Next, we extend the result of Claim 4 to more than two loops. Specifically, we say that two loops L_i and L'_i are *connected* if either they share at least one CKC, or there exists a sequence of loops starting with L_i and ending with L'_i where each two consecutive loops share at least one CKC.

Claim 5. *Consider a set A of type-2 irreducible and connected F_1 -loops, i.e., every two loops are connected by one of these type-2 irreducible loops. Then, every F_2 -measurable $\tau_2|_A$ is F_1 -measurable.*

²²Note that the proof of Claim 4 also holds if ω and ω' are not in the original L_1 and L'_1 loops, respectively, but are simply states in different CKCs that these loops intersect. That is, if ω and ω' are in different CKCs that L_1 and L'_1 intersect and $F_1(\omega) = F_1(\omega')$, we can construct an F_1 -fully-informative loop that starts with ω and ends with ω' in a similar manner as before, and eventually conclude that $F_2(\omega) = F_2(\omega')$.

Proof. Let us prove this by induction on the number of loops. The case of two loops is proved in Claim 4, so assume the statement holds for m loops, and consider a set of $m + 1$ type-2 irreducible and connected F_1 -loops. Further assume, by contradiction, that there exists an F_2 -measurable strategy over this set that is not F_1 -measurable. Thus, there exists ω and ω' such that $F_2(\omega) \neq F_2(\omega')$ whereas $F_1(\omega) = F_1(\omega')$. Evidently, ω and ω' are in different loops and different CKCs. Denote the loops of ω and ω' by L_1 and L'_1 , respectively.

If L_1 and L'_1 are connected directly (through a joint CKC) or through at most m loops (including L_1 and L'_1), then the induction hypothesis holds and every F_2 -measurable strategy this set of loops is F_1 -measurable, implying that $F_2(\omega) = F_2(\omega')$. Thus, assume that L_1 and L'_1 are connected through a sequence of all the $m + 1$ loops (including L_1 and L_{m+1}). Note that ω' cannot be in the same partition element as any other state from this set of loops, other than ω , the state connected to ω in L_1 , and the state connected to ω' in L'_1 . Otherwise, either one of these loops is not type-2 irreducible, or the F_2 -measurability constraints with every intermediate loop is met (by the induction hypothesis) and again we get that $F_2(\omega) = F_2(\omega')$.

Thus, we can now follow the same stages as in the proof of Claim 4 and generate an F_1 -fully-informative loop based on the sequence of loops connecting L_1 and L'_1 (as well as ω and ω'), which starts with ω_1 and ends with $\bar{\omega}'_1$. In this case, Claim 3 holds and we get a type-2 irreducible F_1 -loop, which starts with ω_1 and ends with $\bar{\omega}'_1$, that is also an F_2 -loop. We therefore conclude that $F_2(\omega) = F_2(\omega')$ and the induction follows accordingly. $\square$

After we established that every F_2 -measurable strategy over a set of connected loops is F_1 -measurable, let us extend this result to all the CKCs that these loops intersect. For that purpose, let A be a maximal set of connected loops, where every two are connected, and let C_A be the set of all CKCs that intersect one of these loops (that is, every CKC contains a pair of states from one of these loops). We refer to every C_A as a *cluster*. We argue that every F_2 -measurable strategy over a cluster C_A is F_1 -measurable. To see this, recall Footnote 22 which states that the proof of Claim 4 holds for every ω and ω' in two different CKCs that intersect two connected loops L_1 and L'_1 , respectively. Namely, for every two such states ω and ω' where $F_1(\omega) = F_1(\omega')$, it follows that $F_2(\omega) = F_2(\omega')$. So, as argued in the proof of Claim 5, we conclude that every F_2 -measurable strategy over a cluster is F_1 -measurable.

Observation 3. *Every F_2 -measurable strategy over a cluster is F_1 -measurable.*

Once we have established that every F_2 -measurable strategy over a cluster is F_1 -measurable, let us consider a partition Ω^* of Ω into clusters and individual CKCs that are not part of clusters. Note that *any* two elements of the partition Ω^* jointly intersect at most one partition element of F_1 , otherwise the two components would be in the same cluster. To see this, consider the different possible intersections of elements in Ω^* . If both elements A_1 and A_2 are CKCs, then any two different partition elements of F_1 that intersect both A_1 and A_2 would form a type-2 irreducible F_1 -loop. Otherwise, one of these elements is a cluster, say A_1 , and it follows from previous proofs that for every ω and ω' that belong to the same cluster (but in different CKCs) and $F_1(\omega) = F_1(\omega')$, then one can form an F_1 -fully-informative loop that starts with ω and ends with ω' . Thus, in case ω and ω' are in cluster A_1 and in different partition elements of F_1 that intersect A_2 (whether A_2 is a CKC or another cluster), one can form an F_1 -fully-informative loop that intersects A_1 and A_2 . Using Claim 3, we can conclude that A_1 and A_2 belong to the same cluster. This result is summarized in the following observation.

Observation 4. *Fix two elements $A_1, A_2 \in \Omega^*$. Then, there exists at most one partition element $F_1(\omega)$ of F_1 such that $F_1(\omega) \cap A_1$ and $F_1(\omega) \cap A_2$ are non-empty sets.*

We would now want to prove that Oracle 1 can mimic every F_2 -measurable strategy defined over Ω^* . For this purpose, we present the following Lemma 3 which relates to the F_2 -measurability constraints over different sets of CKCs, that are not in the same cluster (i.e., they are not connected by type-2 irreducible F_1 -loops).

Lemma 3. *Fix two disjoint sets $A_1, A_2 \subseteq \Omega$ that do not intersect the same CKCs, and denote $A = A_1 \cup A_2$. Assume that:*

- *For every i and for every F_2 -measurable $\tau_2|_{A_i}$, there exists an F_1 -measurable $\tau_1^i|_{A_i}$, such that $\mu_{\tau_1}|_{A_i} = \mu_{\tau_2}|_{A_i}$.*
- *For every $\omega_1, \omega'_1 \in A_1$ and $\omega_2, \omega'_2 \in A_2$ such that $F_1(\omega_1) = F_1(\omega_2)$ and $F_1(\omega'_1) = F_1(\omega'_2)$, it follows that $F_1(\omega_1) = F_1(\omega'_1)$.*

Then, for every $\tau_2|_A$, there exists $\tau_1|_A$ such that $\mu_{\tau_1}|_{A_i} = \mu_{\tau_2}|_{A_i}$ for every $i = 1, 2$.

Proof. Fix $\tau_2|_A$ and $\tau_1^i|_{A_i}$ where $i = 1, 2$, such that $\mu_{\tau_2}|_{A_i} = \mu_{\tau_1^i}|_{A_i}$ for every i . Define the sets $\tilde{A}_i = \{\omega_i \in A_i : \exists \omega_{-i} \in A_{-i}, F_1(\omega_i) = F_1(\omega_{-i})\}$ for every $i = 1, 2$. The second condition of the claim implies that all the states in $\tilde{A}_1 \cup \tilde{A}_2$ are in the same partition element of F_1 . To see this, fix $\omega_1 \in \tilde{A}_1$ and, by definition, there exists a state $\omega_2 \in \tilde{A}_2$ such that $F_1(\omega_1) = F_2(\omega_2)$. If there exists another $\omega'_1 \in \tilde{A}_1$, it is either connected to ω_2 (i.e., $F_1(\omega'_1) = F_1(\omega_2)$), or to some $\omega'_2 \in \tilde{A}_2$, and in that case the condition implies that $F_1(\omega_1) = F_1(\omega'_1)$. The same holds for every $\omega_2 \in \tilde{A}_2$.

For every $i = 1, 2$, let S_i be the signals induced by $\tau_1^i|_{A_i}$. Define the following strategy τ_1 :

$$\tau_1((s_1, s_2)|\omega) = \begin{cases} \tau_1^1(s_1|\omega)\tau_1^2(s_2|\tilde{A}_2), & \text{if } \omega \in A_1, (s_1, s_2) \in S_1 \times S_2, \\ \tau_1^1(s_1|\tilde{A}_1)\tau_1^2(s_2|\omega), & \text{if } \omega \in A_2, (s_1, s_2) \in S_1 \times S_2. \end{cases}$$

One can easily verify that $\sum_{(s_1, s_2)} \tau_1((s_1, s_2)|\omega) = 1$ for every ω , so τ_1 is indeed a strategy.

Let us now prove that τ_1 is F_1 -measurable and $\mu_{\tau_1}|_A = \mu_{\tau_2}|_A$. If we restrict τ_1 to A_i , it is clearly F_1 -measurable as $\tau_1^{-i}(s_{-i}|\tilde{A}_{-i})$ is fixed for every $\omega \in A_i$ and $s_i \in S_i$. Thus, consider $\tau_1((s_1, s_2)|\omega)$ where $\omega \in \tilde{A}_1$. All the states in $\tilde{A}_1 \cup \tilde{A}_2$ are in the same partition element of F_1 , so for every $(\omega_1, \omega_2) \in \tilde{A}_1 \times \tilde{A}_2$ we get

$$\begin{aligned} \tau_1((s_1, s_2)|\omega_1) &= \tau_1^1(s_1|\omega_1)\tau_1^2(s_2|\tilde{A}_2) \\ &= \tau_1^1(s_1|\tilde{A}_1)\tau_1^2(s_2|\tilde{A}_2) \\ &= \tau_1^1(s_1|\tilde{A}_1)\tau_1^2(s_2|\omega_2) \\ &= \tau_1((s_1, s_2)|\omega_2), \end{aligned}$$

and the F_1 -measurability condition holds. Moreover, for every $\omega_i, \omega'_i \in A_i$ and for every (s_1, s_2) such that $\tau_1^i(s_i|\omega) > 0$ where $\omega \in \{\omega_1, \omega'_1\}$, it follows that

$$\frac{\tau_1((s_1, s_2)|\omega_i, A_i)}{\tau_1((s_1, s_2)|\omega'_i, A_i)} = \frac{\tau_1^i(s_i|\omega_i)}{\tau_1^i(s_i|\omega'_i)},$$

which implies that conditional on A_i , τ_1 yields the same distribution over posteriors profiles as τ_1^i , thus mimicking τ_2 on every A_i , as needed. $\square$

We can thus finalize the proof using induction on the number of elements in Ω^* . Until now, we established in Observation 3, Observation 4 and Lemma 3 that, given either $|\Omega^*| = 1$ or $|\Omega^*| = 2$, then for every F_2 -measurable strategy $\tau_2|_{\Omega^*}$, there exists $\tau_1|_{\Omega^*}$ such that $\mu_{\tau_1}|_A = \mu_{\tau_2}|_A$ for every $A \in \Omega^*$. Assume this holds for $|\Omega^*| = k \geq 2$, and consider $|\Omega^*| = k + 1$.

Denote the elements of Ω^* by $A_1, A_2, \dots, A_k, A_{k+1}$. If there exists only one partition element of F_1 that intersects A_{k+1} and at least one A_i for $i \leq k$, then Lemma 3 holds and the result follows. Thus, assume there are at least two different partition elements $F_1(\omega) = F_1(\omega_1)$ and $F_1(\omega') = F_1(\omega_2)$ of F_1 such that $\omega, \omega' \in A_{k+1}$ and $\omega_i \in A_i$ for every $i = 1, 2$.

The proof now splits into two parts: either A_1 and A_2 are connected (i.e., there exists a sequence of partition elements of F_1 that sequentially intersect elements in $\Omega^* \setminus A_{k+1}$, starting with A_1 and ending with A_2) or A_1 and A_2 are unconnected. If they are unconnected, we can apply Lemma 3 for A_1 and A_{k+1} and then use the induction hypothesis, so assume they are connected.

Whether A_{k+1} is a CKC or a cluster and assuming that A_1 and A_2 are connected, we argue that there exists a type-2 irreducible F_1 -loop that include ω and ω' , implying that A_{k+1} is part of a cluster with other elements in Ω^* . To see this, recall whenever ω and ω' belong to the same cluster and $F_1(\omega) = F_1(\omega')$, then there exists an F_1 -fully-informative loop that start with ω and ends with ω' . So consider such a sequence of states $l_{\omega \rightarrow \omega'} = (\omega, \dots, \omega')$, which would have been an F_1 -loop had $F_1(\omega) = F_1(\omega')$.

Next, fix the entire path of connections of elements in Ω^* that starts with A_1 and ends with A_2 . Again, the connection between A_1 and A_2 implies that there exists a sequence of states $l_{\omega_1 \rightarrow \omega_2} = (\omega_1, \dots, \omega_2)$ in $\Omega^* \setminus A_{k+1}$, that would have been an F_1 -loop had $F_1(\omega_1) = F_1(\omega_2)$. Hence, consider the sequence of states $l = (\omega, \dots, \omega', \omega_2, \dots, \omega_1)$ which forms an informative F_1 -loop, because $F_1(\omega) \neq F_1(\omega')$. Using Proposition 6 and Claim 3, we know that this loop has a type-2 irreducible F_1 -sub-loop that contains ω and ω' . Thus, A_{k+1} is in the same cluster as other elements in Ω^* , thus contradicting the assumption that $|\Omega^*| = k + 1$. $\square$