Comparison of Oracles*

David Lagziel† Ehud Lehrer‡ Tao Wang§

May 15, 2026

Abstract

We analyze incomplete-information games where an oracle publicly shares information with privately informed players. One oracle dominates another if, in every game, it can match the set of equilibrium outcomes induced by the latter. Distinct characterizations of dominance and equivalence (mutual dominance) are provided for deterministic and stochastic signaling functions. The analysis highlights the role of common-knowledge components and develops a theory of information loops, thereby extending the seminal work of Blackwell (1951) to strategic environments and Aumann (1976)'s theory of common knowledge.

JEL codes: C72, D82, D83.

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

*For their valuable comments, the authors wish to thank participants of the University of Toronto Economic Theory Workshop, the Durham University Economics Seminar, the Adam Smith Business School Micro theory seminar of Glasgow University, INSEAD EPS seminar, The School of Economics seminar at the University of Edinburgh, Western University Economics seminar, the Rationality Center Game Theory Seminar, the Technion Game Theory Seminar, the Bar-Ilan University Theoretical Economics Seminar, the Tel Aviv University Game Theory 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. 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: davidlag@bgu.ac.il. Visiting Professor, Department of Economics, University of Toronto, 2025–2027.

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

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

1 Introduction

Public information in strategic environments does more than reveal facts: it changes how privately informed agents reason about one another. As a result, forecasts, ratings, disclosures, and recommendations may alter equilibrium behavior even when the intermediary itself has only partial information. This raises a fundamental comparative question: when is one information intermediary more informative than another? The question is central to environments involving rating agencies, media outlets, central banks, prediction markets, and online platforms, all of which shape strategic behavior by selectively revealing information to heterogeneously informed agents.

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

Formally, an oracle holds a partition of the state space and communicates through a *signaling function* that is measurable with respect to the oracle's partition. Any such signaling scheme is a Blackwell experiment (see Blackwell, 1951), so an oracle can be viewed as *a generator of experiments* compatible with its information. Each experiment, combined with the players' private partitions, induces a "guided game" with its own set of equilibria. Notably, the oracle need not know what is common knowledge among the players.

The primary objective is to compare two oracles in terms of their ability to induce equilibria across all games. Oracle 1 is said to *dominate* Oracle 2 if, for every game and every experiment available to Oracle 2, Oracle 1 can generate an experiment that induces exactly the same set of equilibrium distributions over state-action profile pairs.

Throughout the paper, we keep the players' private information fixed and vary only the oracles' signaling strategies; we also adopt a stringent notion of dominance of matching the set of equilibrium outcome distributions induced by the other oracle. This makes the comparison independent of any equilibrium-selection rule or objective the oracle might pursue, and contrasts with alternative notions that require the comparison to hold *for every* configuration of the players' private information. Further discussion of these modeling assumptions and the challenges they create is given in Section 3.

Dominance has two complementary components: a local component and a global one. The local part is governed by the players' *common knowledge components* (CKCs), the minimal commonly known events (see Aumann, 1976). The global part is governed by a new structure, referred to as *information loops*: closed paths that link different CKCs through atoms of the oracle's partition. Figure 1 illustrates an F_i-loop, induced by the partition F_i of Oracle i, and consisting of three ordered pairs $((\omega_1, \overline{\omega}_1), (\omega_2, \overline{\omega}_2), (\omega_3, \overline{\omega}_3))$. The presence of two distinct states in every CKC throughout the loop is fundamental: information transmission relies on controlling likelihood ratios across such pairs, and the loops are defined precisely because they constrain this ability.

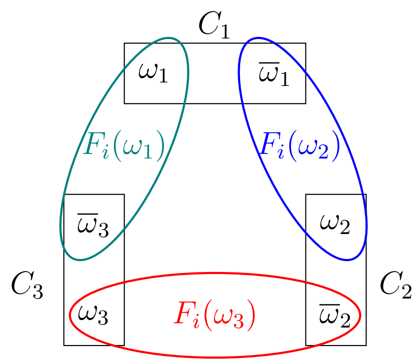

Figure 1: An illustration of an F_i-loop, which intersects three CKCs C_1, C_2 and C_3 with two states in each, using three distinct F_i atoms. F_i denotes the partition of Oracle i, where $i = 1, 2$.

Information loops capture a layer of structure beyond common knowledge. They arise only through interactions with an external agent who is not part of the common-knowledge structure, and they determine whether the local mimicking constructions inside each CKC can be glued into a single oracle-measurable experiment. As a consequence, refinement of partitions within every CKC is not sufficient for dominance: the more informative oracle may still face global compatibility constraints, namely, trade-offs in information transmission across CKCs.

Our analysis begins with the study of loop geometry. We identify several essential properties of information loops, the central one being the existence of an *order-preserving cover*. An F_1-loop admits an order-preserving F_2-cover if its pairs, excluding those that F_2 does not distinguish, can be partitioned into F_2-sub-loops, each preserving the cyclic order inherited from the original F_1-loop. This property, along with several others studied in Section 4, underlies our two main results.

The first main result, Theorem 1, characterizes *equivalent* oracles, defined through mutual dominance. The equivalence problem is the cleanest benchmark: mutual dominance forces both oracles to agree locally within every CKC and globally on their loop structure. Specifically, Oracles 1 and 2, with partitions F_1 and F_2 respectively, are equivalent if and only if F_1 and F_2 refine each other within every CKC and every F_i-loop admits an order-preserving F_{-i}-cover, for $i = 1, 2$.

We then turn to one-sided dominance, Theorem 2, our second main result and a loop-cover theorem. If Oracle 1 dominates Oracle 2, then F_1 refines F_2 inside every CKC, every

F_1-loop is covered by F_2-loops, and whenever the cover is unique it must preserve the cyclic order of the original loop. Conversely, if F_1 refines F_2 inside every CKC, every F_1-loop has an order-preserving F_2-cover, and the F_1-loops are separated, then Oracle 1 dominates Oracle 2. Thus, dominance is governed by local refinement together with global loop compatibility.

These results immediately yield several benchmark cases. If the state space consists of a single CKC, or if there are no F_1-loops at all, the global loop condition is vacuous and dominance collapses to refinement within every CKC (Corollary 1).

Finally, we focus on deterministic experiments. A deterministic experiment assigns a signal to each atom of the oracle's partition, and in this restricted setting dominance collapses to a *jointly more informative* (JMI) condition: Oracle 1 dominates Oracle 2 if and only if, for every deterministic experiment of Oracle 2, Oracle 1 can match the players' joint posterior beliefs once their private information is taken into account (Proposition 3). The refinement condition in the stochastic results does not follow from the deterministic JMI criterion (Example 1), so the stochastic theory is not simply a stochastic recasting of the deterministic one.

1.1 Motivation and applications

Although Section 2 develops a complete motivating example, it is useful to discuss how our model applies more broadly, starting from the opening example in Kamenica and Gentzkow (2011). Recall that a fully informed prosecutor seeks to persuade an uninformed judge to convict a defendant. To achieve this, the prosecutor commits to a signaling structure designed to induce conviction with some probability.

Taking this example a step further, in many judicial systems the prosecutor must persuade not only a judge but also a *jury*.[1] Moreover, the prosecutor is typically *not* fully informed, while neither the judge nor the jury is completely uninformed. Each holds personal intuitions, prior beliefs, and subjective perspectives. In practice, the prosecutor faces informational constraints analogous to those encountered by the oracle in our framework. In this analogy, the prosecutor is the oracle, the judge and jury are the players, and their private intuitions and priors correspond to the players' private information. Our partial order ranks prosecutorial information systems (e.g., evidence-disclosure regimes) by the sets of equilibrium verdicts they can induce across all such games. The same logic applies to the problem of belief polarization in our Section 2 example: different "oracles" (fact checkers) can be ranked by the equilibrium depolarization patterns they can implement.

[1] For example, the Sixth Amendment to the United States Constitution guarantees criminal defendants a speedy and public trial by an impartial jury.

This perspective highlights why the partial ordering of oracles is a central comparative notion in such environments: standard tools (e.g., concavification) may fail under partial and asymmetric information. In a similar set-up, Lagziel and Lehrer (2025) show that the feasible set of joint posterior beliefs is typically non-convex given a non-binary state space, when both the oracle and the players possess private and partial information. Consequently, a partial ordering of "experiment generators" becomes a powerful methodological tool for analyzing these problems. Practically, this means that in settings with partially informed judges/juries or central banks/markets, one can compare alternative information systems (transparency regimes, disclosure formats, types of forecasts) by asking which oracle dominates another, without solving a fully-fledged information-design problem.

1.2 Relation to literature

Blackwell (1951, 1953) provides the classic comparison of experiments for a single decision maker: one information structure dominates another if it yields weakly higher expected utility in every decision problem. We extend this framework along two dimensions. First, we move from decision problems to incomplete-information games, where the objects being compared are equilibria of "guided games" induced by an oracle's signals. Second, instead of a fixed experiment, an oracle is a generator of experiments compatible with its partition. Our dominance notion is therefore about the ability of one oracle to replicate the equilibrium outcome distributions that another oracle can induce, rather than optimizing a particular decision maker's payoff.

Brooks et al. (2024) strengthen Blackwell by requiring robustness to arbitrary auxiliary signals and decision problems and characterize *strong Blackwell dominance* between two signals. Their analysis compares two information sources (signals) that are robust to any external information source and decision problem. In contrast, we fix the players' private information structures and the prior, and compare oracles that can implement any experiment measurable with respect to their partitions. Our dominance relation is specific to a given configuration of players' information and to strategic interaction, not to a universal set of decision problems.

A related literature compares information structures in games and establishes partial orders. 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, 2013) analyze signaling and mediation in common-interest games and show how variants of Blackwell garbling characterize outcome equivalence. Likewise, Bergemann and Morris (2016) characterize dominance among two information structures through the concept of individual sufficiency, an extension of Blackwell's notion of garbling to n-player games. A common feature of these papers is that they compare *fixed* information structures (typically private signals) and use versions of Blackwell's garbling to capture dominance. We instead hold players' private information fixed and compare oracles that provide *public signals*, subject to an oracle-measurability constraint (the oracle cannot condition on players' private signals). For this reason, oracle dominance is not meant to coincide with, or simply specialize, Bergemann and Morris (2016) to public signals: even under a "public-signal-only" restriction, their order is a pairwise comparison of two fixed information structures, whereas our is a comparison over the menus of feasible public experiments generated by different oracle partitions. Dominance in our setting is driven by the implementable joint posterior beliefs, constrained by CKCs and information loops.

Another strand studies mediators in incomplete-information games who correlate players' actions through private recommendations, often without adding information about the realized state; see Forges (1993) and Gossner (2000) among others. In these models, the mediator's role is to coordinate actions so as to implement correlated equilibria. Our oracles differ in two respects: they provide *public signals* rather than private recommendations, and their role is not to coordinate actions but to change beliefs. As a result, dominance in our setting cannot be reduced to the richness of the correlated-equilibrium correspondence.

Our departure from existing dominance notions (see, e.g., Gossner, 2000 and Bergemann and Morris, 2016) lies in restricting information provision to *public disclosure*, given players' private information and the Oracle's information. This restriction imposes global measurability and consistency constraints that are absent in other frameworks. If players' private information is trivial, if the Oracle is allowed to issue private recommendations, or if the Oracle is agnostic about players' information and the comparison ranges over all possible information structures, these measurability constraints vanish, and our notion of dominance collapses to standard refinement or Blackwell-type comparisons. The conceptual contribution of this paper is to isolate and clarify the role of public disclosure in shaping the interaction between fixed players' private information and the Oracle's information, and to show how this interaction fundamentally alters the set of attainable equilibrium outcomes.

Our use of common knowledge components is rooted in the epistemic foundations of games. Aumann (1976) defines common knowledge and the induced partition into common knowledge components. Monderer and Samet (1989), Mertens and Zamir (1985), and Brandenburger and Dekel (1993) clarify how hierarchies of beliefs and type spaces encode such information. Our model builds on these studies by fixing the partition structures while varying only the oracle's public experiment. The novel constraints we study arise from *global* measurability across CKCs (via loops), not from additional complexity in private belief hierarchies. Information loops then formalize how public measurability links different CKCs and constrains the set of posterior profiles that an oracle can generate.

The dominance problem, and in particular the analysis of order-preserving covers, can be partially translated into the study of Eulerian directed multigraphs. We employ the main results and techniques of Cooper and Okur (2025) to show that the unique-cover condition corresponds to the case in which distinct loops intersect in at most one atom. To the best of our knowledge, however, no general theory currently exists for the order-preservation property.

1.3 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, the dominance relation, and the posterior-mimicry perspective used throughout the analysis. Section 4 introduces common knowledge components, local refinement, information loops, covers, and order preservation. Section 5 gives the complete characterization of equivalent oracles. Section 6 studies one-sided stochastic dominance and states the loop-cover dominance theorem. Section 7 collects the benchmark cases: unique CKC, acyclic/no-loop stochastic structures, non-informative loops, the two-CKC case, and deterministic dominance. Section 8 concludes.

2 Fact checking and belief depolarization: an example

The problem of polarization, whether affective, ideological, or identity-driven, has received substantial attention in both public discourse and academic research (e.g., Arieli et al., 2021; Arieli et al., 2024, and Ikan et al., 2025). Technological progress has not mitigated this problem. Selective exposure to information and the ease of coordinated disinformation may instead amplify disagreement, making fact checking and depolarization particularly challenging. The following stylized example illustrates how our framework captures these forces.

Consider two individuals (or populations) who receive information from different sources about an unknown state $\omega \in \Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ with common prior $\mu = (\frac{1}{16}, \frac{4}{16}, \frac{5}{16}, \frac{6}{16})$. Their private information is represented by the partitions $\Pi_1 = \{\{\omega_1, \omega_3, \omega_4\}, \{\omega_2\}\}$ and $\Pi_2 = \{\{\omega_1\}, \{\omega_2, \omega_3, \omega_4\}\}$. To abstract from actions and payoffs, polarization is measured by the expected ℓ^1-disagreement $D(\tau) = \mathbb{E}[\|\mu^1_\tau - \mu^2_\tau\|_1]$, where μ^i_τ denotes player i's posterior given her private information and the public signal generated by τ.[2] If no public information

[2]The $D(\tau)$ is essentially the "β-polarization" in Arieli et al., 2024, when $\beta = 1$. The objective $D(\tau)$ is

is provided, that is, if the public signal is constant, each player conditions only on her private information. A direct computation yields $D(\tau^{\text{Const}}) = \frac{407}{480} \approx 0.848$.

Now consider a fact checker, referred to as Oracle 1, who seeks to reduce disagreement. Oracle 1 observes the public partition $F_1 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. While a perfectly informed oracle could eliminate disagreement, in practice verifiable evidence is often limited. In this example, F_1 represents the informational constraint faced by the fact checker. If Oracle 1 fully reveals the realized F_1-cell, then $D(\tau_1^{\text{Full}}) = \frac{97}{120} \approx 0.808$, so revealing F_1 reduces disagreement relative to having no oracle.

Next, define an F_1-measurable experiment with two public signals s_1, s_2 by

$$\tau_1^{\text{Noisy}}(s_1 \mid \omega) = \begin{cases} 0, & \text{if } \omega \in \{\omega_1, \omega_3\}, \\ \frac{3}{4}, & \text{if } \omega \in \{\omega_2, \omega_4\}, \end{cases} \qquad \tau_1^{\text{Noisy}}(s_2 \mid \omega) = 1 - \tau_1^{\text{Noisy}}(s_1 \mid \omega).$$

This F_1-measurable experiment is a garbling of τ_1^{Full} and hence Blackwell-inferior to full revelation of F_1.[3] Nevertheless, a direct computation gives $D(\tau_1^{\text{Noisy}}) = \frac{31}{40} = 0.775 < D(\tau_1^{\text{Full}})$, so full revelation is strictly dominated for the purpose of minimizing expected disagreement.

Consider instead Oracle 2, endowed with the public partition $F_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$. Notice that F_2 neither refines F_1 nor is refined by it, so neither oracle Blackwell-dominates the other. If Oracle 2 fully reveals the realized F_2-cell, then for every $\omega \in \Omega$, we get $\Pi_1(\omega) \cap F_2(\omega) = \Pi_2(\omega) \cap F_2(\omega)$. Thus the players condition on the same event at every state and therefore hold identical posterior beliefs, implying $D(\tau_2^{\text{Full}}) = 0$. In summary,

$$D(\tau^{\text{Const}}) > D(\tau_1^{\text{Full}}) > D(\tau_1^{\text{Noisy}}) > D(\tau_2^{\text{Full}}) = 0.$$

This example highlights several themes that recur throughout our analysis. First, players' private information substantially complicates the analysis, and the problem of polarization more generally, to the extent that even garbling need not be "inferior" to full revelation. Second, the oracle's objective in this example, and more generally in any strategic environment we consider, is to control the distribution of joint posterior beliefs, which encodes the entire hierarchy of beliefs. This objective is distinct from, and typically more restrictive

used only as a convenient summary statistic for posterior disagreement. The example is intended to illustrate informational feasibility and dominance comparisons, rather than to formulate an optimization problem for the oracle. One may microfound this objective via a coordination game in which equilibrium payoffs depend on the players' joint posterior beliefs, so that less disagreement is beneficial.

[3] Let τ_1^{Full} reveal the realized F_1-atom, with signals t for $\{\omega_1, \omega_3\}$ and f for $\{\omega_2, \omega_4\}$. Define a kernel K by $K(s_1 \mid t) = 0$, $K(s_2 \mid t) = 1$, and $K(s_1 \mid f) = \frac{3}{4}$, $K(s_2 \mid f) = \frac{1}{4}$. Then $\tau_1^{\text{Noisy}}(s \mid \omega) = \sum_{r \in \{t,f\}} K(s \mid r)\, \tau_1^{\text{Full}}(r \mid \omega)$ for all ω, so τ_1^{Noisy} is a garbling of τ_1^{Full}.

than, coordination via correlated equilibrium in the sense of Forges (1993). Third, this example also motivates the general question: given players' partitions (possibly with multiple CKCs), when does Oracle 1 dominate Oracle 2 in the sense that for every experiment τ_2 there exists an experiment τ_1 inducing the same joint distribution of posterior beliefs (same polarization)?

3 The model and posterior mimicry

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 public 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 $N = \{1, 2, \ldots, n\}$ 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 set of actions 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 $u_i : \Omega \times A \to \mathbb{R}$, which maps states and action profiles to real-valued payoffs.[4]

To extend the basic game into a guided game, we introduce an oracle who provides public information before players choose their actions. The oracle is endowed with a partition F of the state space Ω, and a countable set S of possible signals. A *signaling strategy* of the oracle is an F-measurable experiment $\tau : \Omega \to \Delta(S)$, where $\Delta(S)$ is the set of probability distributions on S with finite support. We denote by $\tau(s|\omega)$ the probability that the oracle sends signal s when the realized state is ω.

The guided game evolves as follows. First, the oracle publicly announces an experiment τ. Then, a state $\omega \in \Omega$ is drawn according to a common prior $\mu \in \Delta(\Omega)$ with full support. 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, a signal $s \in S$ is drawn according to $\tau(\omega)$ and is publicly announced.

Let the join[5] $\Pi_i \vee F'$ denote the updated information (i.e., partition) of player i given Π_i and some partition F'. Let $\mu^i_{\tau|\omega,s} = \mu(\cdot|\Pi_i(\omega), \tau, s) \in \Delta(\Omega)$ denote player i's posterior belief

[4]The underlying state space Ω represents payoff-relevant fundamentals only, referred to in Maschler et al. (2013) as the states of nature. Agents' information is modeled by their information partitions, and payoffs do not depend directly on agents' information or beliefs. The states of nature together with these partitions, when they are common knowledge among all players, determine the players' types, which are characterized by the full hierarchy of beliefs, à la Harsanyi.

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

after observing $\Pi_i(\omega)$ and a realized signal s according to $\tau(\omega)$, and let $\mu_{\tau,s} = \{(\mu^i_{\tau|\omega,s})_{i\in N} : \omega \in \Omega \text{ s.t. } \tau(s|\omega) > 0\}$ be the set of *joint posteriors* associated with τ and a signal s, across all relevant states. The joint posteriors capture each player's belief about the realized state and their beliefs about others' beliefs, as well as higher-order beliefs. We use μ_τ to denote the distribution over all joint posteriors induced by τ across all signals, and use $\text{Post}(\tau) = \text{Supp}(\mu_\tau)$ to denote its support. Thus, every experiment τ yields an incomplete-information game $G(\tau) = (N, (A_i)_{i\in N}, \mu_\tau, (u_i)_{i\in N})$. When there is no risk of ambiguity, we denote the incomplete-information game without τ by G.

3.1 Partial ordering of oracles

To discuss the role of the oracle in the current framework, we define a relevant solution concept, referred to as a *Guided equilibrium*, which incorporates the oracle's strategy. Formally, let $\sigma_i : \Pi_i \times S \to \Delta(A_i)$ be a strategy of player i. A tuple $(\tau, \sigma_1, \ldots, \sigma_n)$ is a *Guided equilibrium* if $(\sigma_1, \ldots, \sigma_n)$ 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 $\text{NED}(\text{G}(\tau)) \subseteq \Delta(\Omega \times \text{A})$ 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 experiment of Oracle j by F_j and τ_j, respectively.

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

Intuitively, Oracle 1 dominates Oracle 2 if, whatever experiment Oracle 2 uses and whatever game is played, Oracle 1 can choose an experiment that yields exactly the same set of equilibrium outcome distributions. Note that a direct comparison of the games' equilibria is problematic because the players' strategies depend on the oracles' experiments.

Three points are worth noting here. First, one could consider defining dominance between oracles in a more robust manner by allowing players' information structures to vary over a set of possible partitions. If dominance were required to hold uniformly over all such partitions, the comparison problem would become rather simple. In particular, the set of admissible information structures would include the case of trivial private information. In that case,

[6] Note that a Nash equilibrium $(\sigma_i^*, ..., \sigma_n^*)$ induces a probability distribution over $\Omega \times A$. Specifically, fix ω and an action profile a, the probability of (ω, a) under the equilibrium strategy $(\sigma_i^*, ..., \sigma_n^*)$ and the experiment 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, $\text{NED}(\text{G}(\tau))$ is a subset of $\Delta(\Omega \times A)$.

the common-knowledge structure admits no loops and, by our results, oracle dominance reduces to partition refinement. The substantive challenge in our framework arises precisely because players' information partitions are fixed and predetermined, so that dominance must be assessed under binding measurability and common-knowledge constraints.

Second, Definition 1 compares oracles via equality of equilibrium distributions rather than inclusion. Analogously, one can define dominance in terms of equilibrium payoff sets, either via equality or inclusion. Outcome–based dominance implies the corresponding payoff–based notions, but the former is more natural when the objective is to control actions and aggregate outcomes, whereas payoff–based notions suit environments that focus on agents' utilities. We adopt this stronger notion to avoid imposing a specific equilibrium-selection rule (which may diverge from the Pareto frontier), and to keep the dominance relation independent of any objective functions that might be assigned to the oracles in a parallel setup.

Third, another way to compare oracles is to treat them as strategic players in some sender–receiver game, assign payoff functions, and say that Oracle 1 is more informative than Oracle 2 if it obtains (weakly) higher equilibrium payoffs in every such game. This approach faces several difficulties: (i) equilibria are typically multiple (so oracle payoffs depend on an arbitrary selection rule); (ii) it presumes that the oracle knows the players' private information partitions (and more generally their type structure); and (iii) it ties the comparison of information structures to particular oracle objectives. By contrast, our dominance notion compares oracles through the equilibrium outcome distributions they induce for a *fixed* configuration of players' information, which is common knowledge among the players but not necessarily to the oracle. This lets us analyze the oracles' information partitions, and the associated notions, independently of any objectives they may have.

The analysis below repeatedly uses the following posterior-mimicry perspective. For any experiment τ, the induced guided game depends on τ only through the distribution μ_τ over joint posterior profiles. Thus, to prove that one oracle dominates another, we construct an experiment that reproduces that distribution. Conversely, when such posterior mimicry fails, proper-scoring-rule games separate the two oracles, and the induced equilibrium distributions cannot coincide in every game.[7] This is summarizes in the following lemma (all proofs are deferred to Appendix A).

Lemma 1 (Posterior mimicry)**.** *Oracle* 1 *dominates Oracle* 2 *if and only if for every experiment* τ_2 *there exists an experiment* τ_1 *such that* $\mu_{\tau_1} = \mu_{\tau_2}$*.*

[7] In the extended version of the paper, we construct a two-stage finite game that can replace these proper-scoring-rule games.

4 Information loops

Our characterization of dominance has two components: a local one, comparing the oracles' partitions within each CKC, and a global one, capturing how the two partitions link different CKCs. To analyze the global component, we develop a theory of *information loops*, presented below.

4.1 Common knowledge components and local refinement

To define the local part of stochastic dominance, recall the notion of a common knowledge component. Following Aumann (1976), an event E is a *common knowledge component* (CKC) of the partitions $\Pi_1, \Pi_2, \ldots, \Pi_n$ if it is an element in the meet $\bigwedge_{i=1}^{n} \Pi_i$, which is the finest common coarsening of all the partitions. That is, a CKC is a minimal event (inclusion-wise) that is common knowledge among the players. Within each CKC, all posteriors and hence equilibrium payoffs are determined solely by the information available inside that component, so players' expected payoffs and the oracle's impact can be analyzed CKC-by-CKC.

For a partition F and a CKC C, write $F|_C$ for the restriction of F to C, namely the collection of non-empty sets $A \cap C$ with $A \in F$. We say that F_1 *refines* F_2 in the CKC C if $F_1|_C$ refines $F_2|_C$, and we say that F_1 refines F_2 in every CKC if this condition holds for every common knowledge component. This local refinement condition is the within-component part of dominance. The remaining difficulty is global: a single F_1-measurable experiment must be chosen across all CKCs, and this creates additional measurability constraints.

When there is a single CKC, dominance is essentially local: refinement inside that component is enough to reproduce the relevant posterior distributions. With several CKCs, this local argument no longer settles the problem. Atoms of the oracle partition may connect different CKCs, and the mimicking experiment must be measurable across all these links simultaneously. The compatibility constraints generated by such cross-component links are the source of the loop restrictions below. Appendix B.1 gives examples showing that, with several CKCs, a mimicking oracle may need to use a different signal space than the oracle it mimics, even when the two oracles have the same information inside each CKC.

4.2 Information loops and covers

We now define the loop objects that create global constraints across CKCs. Assume that $C_1, \ldots, C_l$ are mutually exclusive CKCs such that $\Omega = \bigcup_{j=1}^{l} C_j$. A key aspect of our analysis is that atoms of an oracle partition may connect different CKCs. A sequence of such connections can form a closed path, which we call an information loop, or simply a loop.

Definition 2. *An F_i-loop is a sequence $(\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2, \ldots, \omega_m, \overline{\omega}_m)$, with $m \geq 2$, where indices are taken modulo m, such that*

- $\omega_j, \overline{\omega}_j \in C_{r_j}$ *and* $\omega_j \neq \overline{\omega}_j$ *for all* $j = 1, \ldots, m$;[8]
- $\omega_{j+1} \in F_i(\overline{\omega}_j)$ *for all* $j = 1, \ldots, m$;
- $C_{r_j} \neq C_{r_{j+1}}$ *for all* $j = 1, \ldots, m$;
- *the sets* $\{\overline{\omega}_j, \omega_{j+1}\}$ *are pairwise disjoint for all* $j = 1, \ldots, m$.

To understand information loops, one can view the CKCs as the vertices of a multigraph. An edge connects two CKCs if there exist ω_{j+1} and $\overline{\omega}_j$ that 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 $m + 1 \equiv 1$). The restriction $C_{r_j} \neq C_{r_{j+1}}$ focuses attention on the genuinely cross-CKC links, because those are the links that create global compatibility constraints. An illustrative loop is given in Figure 1 above.

The reason loops matter is immediate from measurability. If τ_i is F_i-measurable and $(\omega_1, \overline{\omega}_1, \ldots, \omega_m, \overline{\omega}_m)$ is an F_i-loop, then for every fully supported signal s,

$$\prod_{j=1}^{m} \frac{\tau_i(s \mid \omega_j)}{\tau_i(s \mid \overline{\omega}_j)} = 1. \tag{1}$$

Indeed, $\overline{\omega}_j$ and ω_{j+1} belong to the same F_i-atom, so the denominator at one step equals the numerator at the next step. Thus an F_1-loop imposes a product restriction that every F_1-measurable experiment must satisfy. The central question is when every F_2-measurable experiment that is locally mimicked (within every CKC) also respects these loop restrictions. We express the answer through the following notion of *covers*.

Definition 3. *An F_i-loop $(\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2, \ldots, \omega_m, \overline{\omega}_m)$ is F_{-i}-*covered *if the set $\{1, \ldots, m\}$ can be partitioned into disjoint sets of indices $J, I_1, \ldots, I_r$ such that $\big((\omega_j, \overline{\omega}_j)\big)_{j \in I_t}$ is an F_{-i}-loop for each $t = 1, \ldots, r$ (also referred to as an F_{-i}-*sub-loop*) and $J = \{j : \omega_j \in F_{-i}(\overline{\omega}_j)\}$ is the set of F_{-i}-*non-informative *pairs.*[9] *The cover is *order-preserving *if every F_{-i}-sub-loop in the cover follows the same cyclic ordering of pairs as the original F_i-loop, up to a cyclic rotation.*

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

[9] The order of the pairs $(\omega_j, \overline{\omega}_j)$ in the F_{-i}-loop need not coincide with their order in the original F_i-loop. For instance, an F_1-loop $(\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2, \omega_3, \overline{\omega}_3)$ might be covered by the F_2-loop $(\omega_1, \overline{\omega}_1, \omega_3, \overline{\omega}_3, \omega_2, \overline{\omega}_2)$. Whenever $F_{-i}(\omega_j) = F_{-i}(\overline{\omega}_j)$, Oracle $-i$ cannot distinguish between ω_j and $\overline{\omega}_j$, hence uninformative.

In simple terms, an F_{-i}-cover decomposes the F_i-loop into F_{-i}-loops, together with pairs for which Oracle $-i$ cannot provide additional information. Figure 2 illustrates the definition. Panel (a) shows an F_1-loop (consisting of $((\omega_j, \overline{\omega}_j))_{j=1,\ldots,4}$) covered by two order-preserving F_2-loops. Panel (b) shows a case in which a candidate F_2-loop does not cover the original F_1-loop, because it does not use the same ordered CKC-pairs ($\overline{\omega}_1$ is linked to $\overline{\omega}_3$ instead of ω_3). Panel (c) shows a cover that is not order-preserving (the pair $(\omega_3, \overline{\omega}_3)$ appears before $(\omega_2, \overline{\omega}_2)$). Our main dominance result (Theorem 2) states that order-preservation is necessary for dominance whenever the cover is unique, and that order-preserving covers are sufficient under separated loops.

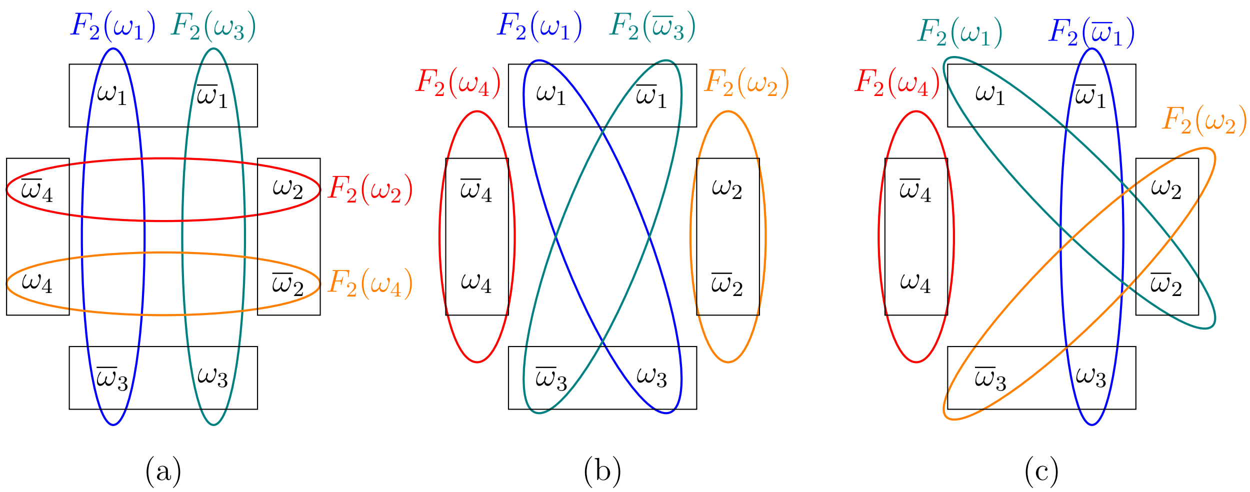

Figure 2: Two states connected by a colored line are in the same information set of F_2. In (a), the F_2-sub-loops that cover the F_1-loop are order-preserving, i.e., following the ordering of pairs in the original F_1-loop, whereas the sub-loop in (c) is not order-preserving. (b) illustrates a case where $(\omega_1, \overline{\omega}_1, \overline{\omega}_3, \omega_3)$ forms an F_2 loop, but it is not an F_2-sub-loop of the original F_1-loop.

4.3 Irreducible and informative loops

The ability to classify covers, and specifically order-preserving covers, is rather difficult. It appears that there is currently no relevant theory for this classification. To face this challenges, we decompose loops and covers into smaller objects, referred to as *irreducible loops*, and use these objects as building blocks in our main results. Formally, an F_i-loop is irreducible if it does not contain a smaller F_i-loop using only states from the original loop.

Definition 4. *Let $L_i = (\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2, \ldots, \omega_m, \overline{\omega}_m)$ be an F_i-loop. We say that L_i is* irreducible *if there is no strict subset of $\{\omega_j, \overline{\omega}_j : j = 1, \ldots, m\}$ that forms an F_i-loop. A cover is* irreducible *if every loop in the cover is irreducible.*

We also distinguish loops according to whether another oracle separates the two states that lie in the same CKC-pair. The non-informative case will be used in Corollary 2 as a sufficient condition for dominance.

Definition 5. *An F_i-loop $(\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2, \ldots, \omega_m, \overline{\omega}_m)$ is F_k-*non-informative *if $F_k(\omega_j) = F_k(\overline{\omega}_j)$ for every j. The loop is F_k-*fully informative *if $F_k(\omega_j) \neq F_k(\overline{\omega}_j)$ for every j, and it is F_k-*informative *if this inequality holds for some j.*

If an F_1-loop is F_2-non-informative, then the loop does not pose a mimicking challenge for Oracle 1, as Oracle 2 cannot provide any relevant information in any of the pairs. Algebraically, every F_2-measurable likelihood ratio $\frac{\tau_2(s|\omega_j)}{\tau_2(s|\overline{\omega}_j)}$ equals one. Thus the F_1-loop product restriction in Eq. (1) is automatically satisfied. More generally, the next proposition records the elementary loop facts that allow us to reduce the analysis to irreducible and fully informative loops.

Proposition 1. *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 into 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 Proposition 1, together with Proposition 4, which establishes additional properties of loops and is proved in Appendix A.2, in the next two sections: first to characterize equivalent oracles, and then to establish necessary and sufficient conditions for one-sided dominance.

5 Equivalent oracles

We begin the substantive comparison of oracles with the two-sided problem, characterizing necessary and sufficient conditions for each oracle to dominate the other. Formally, Oracle 1 is *equivalent* to Oracle 2, written $F_1 \sim F_2$, if $F_i \succeq_{\text{NE}} F_{-i}$ for $i = 1, 2$.

The loop geometry developed above is sharpest in this setting. Equivalence requires that the two partitions agree within every CKC and satisfy the global loop constraints jointly generated by the two partitions. In particular, every irreducible F_i-loop must be covered by F_{-i}-loops. This condition is necessary but not sufficient: equivalence further requires that every irreducible F_i-loop with at least six states is itself an irreducible F_{-i}-loop, which implies order-preservation. (The minimum number of states in a loop is four, and every such loop is automatically order-preserving, so the six-state threshold is the first at which the condition becomes non-vacuous.) Theorem 1 provides this characterization.

Theorem 1. *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 *and any F_i-loop has an order-preserving F_{-i}-cover.*

The equivalence loop condition reflects a symmetry requirement on measurability: if one oracle is constrained by an information loop, the other must be constrained by a corresponding loop. The proof of the theorem leans heavily on irreducibility, through the final statement of Proposition 1: every reducible loop either intersects some CKC more than once or contains at least four states within a single partition element of F_i. Under equivalence, F_1 and F_2 coincide within every CKC, so neither case generates an asymmetric constraint. In the first case, the loop decomposes into disjoint sub-loops, and the associated constraints apply identically on both sides. In the second case, the four states in a common F_i-atom either generate a new loop, which is then covered (with a decomposition matching that on the other oracle's side), or fail to generate any loop at all, so no binding constraint arises.

The proof of Theorem 1 also uses a refinement of this idea. An F_i-loop is *type-2 irreducible* if no four of its states lie in the same partition element of F_i. This condition is strictly weaker than irreducibility, since a type-2 irreducible loop may still intersect the same CKC multiple times and hence admit a decomposition into sub-loops. Even so, we show that F_1 and F_2 coincide along every type-2 irreducible loop.

From type-2 irreducible loops we construct the basic building blocks of the analysis, called *clusters*. Two type-2 irreducible loops are said to be connected if they intersect a common CKC, and the transitive closure of this relation partitions the loops into disjoint families. Each such family determines a cluster, namely, the set of CKCs collectively intersected by its loops. We prove that the oracles' partitions agree within every cluster. Clusters therefore form the basic structure on which equivalence is established, after which the analysis extends to simpler inter-cluster connections involving only a single partition element of F_i.

6 One-sided dominance

We can now state the main one-sided dominance result. The theorem has two parts: the first gives necessary conditions for dominance in general, and the second gives a sufficient condition when the F_1-loops are *separated*. We say that the F_1-loops are separated if no two of them share either a CKC or an F_1-atom. Together, the two parts isolate the two forces underlying dominance: local refinement within CKCs and global compatibility of loop constraints.

Theorem 2. *Fix partitions F_1 and F_2.*

(a) Necessity. *If Oracle* 1 *dominates Oracle* 2*, then:*

   (i) F_1 refines F_2 in every CKC*;*

   (ii) every F_1-loop has a cover by F_2-loops; and

   (iii) if an F_1-loop has a unique F_2-cover, then the cover is order-preserving.

(b) Sufficiency. *If F_1 refines F_2 in every* CKC*, every F_1-loop has an order-preserving F_2-cover, and all F_1-loops are separated, then Oracle* 1 *dominates Oracle* 2*.*

The necessity part shows that a dominant oracle must satisfy both local and global constraints. The local constraint is that F_1 refines F_2 within every CKC. The global constraint is that every F_1-loop is covered by F_2-loops. In particular, when an F_1-loop admits a unique cover, that cover must preserve the cyclic order, since otherwise the likelihood ratios that must be matched around the loop would become inconsistent with F_1-measurability. Building on Cooper and Okur (2025) on Eulerian directed multigraphs, we show that the unique-cover condition corresponds to the case in which distinct loops intersect in at most one atom.

The sufficiency part shows that local refinement together with order-preserving covers implies dominance whenever the F_1-loops are separated. Order-preserving covers ensure consistency of the cyclic likelihood-ratio constraints along each loop, while the separation condition allows the resulting local mimicking constructions to be glued across loops. The proof also relies on a tree-mimicry argument: informative paths create no compatibility obstruction, and the only obstruction arises from closing cycles.

7 Acyclic benchmarks and deterministic experiments

The preceding sections characterize the two ingredients of oracle comparisons: local refinement within each CKC and the global loop geometry. We conclude the body of the paper with two benchmark environments in which the general constraints simplify. In unique-CKC and acyclic stochastic structures, the global ingredient is vacuous and dominance reduces to local refinement. Under deterministic signaling, the loop geometry simplifies into a joint-partition condition.

7.1 Benchmark cases

Theorem 2 turns into a full characterization in two benchmark cases: a unique CKC and no information loops. In both cases, dominance collapses to local refinement. (The proof follows directly from Theorem 2, thus omitted.)

Corollary 1 (Unique-CKC and acyclic benchmarks)**.** *The following statements hold.*

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

  (i) F_1 *refines* F_2*;*

  (ii) $F_1 \succeq_{\mathrm{NE}} F_2$*;*

  (iii) For every τ_2*, there exists* τ_1*, so that* $\mathrm{Post}(\tau_1) \subseteq \mathrm{Post}(\tau_2)$*;*

  (iv) For every τ_2*, there exists* τ_1*, so that* $\mathrm{Post}(\tau_1) = \mathrm{Post}(\tau_2)$*;*

  (v) For every τ_2*, there exists* τ_1*, so that* $\mu_{\tau_1} = \mu_{\tau_2}$*.*

(b) 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 first part is the unique-CKC characterization. The second part is the acyclic benchmark. The proof of the acyclic benchmark constructs a graph whose vertices are CKCs and whose edges connect CKCs intersected by a common atom of F_1. The no-loop assumption makes each connected component of this graph a tree. On a tree, one can assign one independent signal coordinate to each edge and then glue the local mimicking strategies without generating any cyclic compatibility condition.

Next, we record two useful special cases. The first is the case of non-informative loops, in which the loop constraints imposed by F_1 are already respected by F_2. The second is the two-CKC case, where the balance/cover condition is both necessary and sufficient.

Corollary 2 (Non-informative loops)**.** *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*.*

Proposition 2. *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.*

The two-CKC case is especially tractable. Every loop consists of two CKC-pairs, and a cover leaves only two possibilities: either the loop is F_2-non-informative, or it is itself an F_2-loop. The first case is covered by Corollary 2; the second allows Oracle 1 to satisfy the same loop constraint as Oracle 2. In both cases, the order-preservation requirement is vacuous.

7.2 Deterministic experiments

This section focuses on dominance when oracles are restricted to deterministic experiments. Throughout the section, we consider only deterministic experiments, namely maps $\tau_i : F_i \to S$ for each Oracle i. We identify each such experiment with the partition of Ω that it induces. Equivalently, a deterministic experiment of Oracle i is a coarsening of F_i.

The characterization is based on the ability of one oracle to *match* the players' joint posterior beliefs, for any given deterministic experiment of the other oracle. More formally, we say that Oracle 1 is *jointly more informative* (JMI) than Oracle 2 if, for every deterministic experiment of Oracle 2, there exists a deterministic experiment of Oracle 1 that induces the same posterior partition for every player.

Definition 6. *Oracle* 1 *is* jointly more informative than Oracle 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*.*

Evidently, dominance does not require an oracle to possess information that every player has, even if that information is available to the other oracle. The same relation can also be expressed directly in terms of the oracles' coarse partitions.

Observation 1. *Oracle* 1 *is jointly more informative than Oracle* 2 *if and only if, for every coarsening* F_2' *of* F_2*,*[10] *there exists a coarsening* F_1' *of* F_1 *such that* $\Pi_i \vee F_1' = \Pi_i \vee F_2'$ *for every player* i*.*

The partial order generated by the JMI notion need not coincide with the usual refinement order. For example, consider the trivial case in which the players have perfect information. Then every oracle is jointly more informative than every other oracle, independently of their partitions. Nevertheless, in Section 7.3.1, we show that, when there is a unique CKC and every oracle is JMI than the other, then their partitions coincide.

One can also bridge the gap between JMI and refinement by allowing the players' partitions to vary.[11] If JMI is required to hold for every profile of players' partitions, then, in particular, it must hold when all players have the trivial partition. In that case, Oracle 1 must be able to match every deterministic experiment τ_2, so F_1 must refine F_2.

7.2.1 Dominance and JMI

The main result of this subsection shows that, when oracles use deterministic experiments, dominance is equivalent to being jointly more informative.

[10] Equivalently, the σ-field generated by F_2' is a sub-σ-field of the σ-field generated by F_2.

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

Proposition 3. *Assume that oracles are deterministic. Then, Oracle* 1 *dominates Oracle* 2 *if and only if Oracle* 1 *is JMI than Oracle* 2*.*

Intuitively, the "if" direction is immediate: if Oracle 1 can replicate the information of Oracle 2 for every deterministic experiment, the induced incomplete-information games have the same joint posteriors and hence the same equilibrium outcome distributions. For the reverse direction, the proof constructs a game, based on proper scoring rules employing the Kullback–Leibler divergence, with a unique equilibrium which maps to the players' joint posterior belief. Thus, any failure of JMI for Oracle 1 relative to Oracle 2 generates an equilibrium distribution that Oracle 1 cannot mimic.[12]

7.3 JMI and refinement inside CKCs

The stochastic results above are based on refinement inside CKCs. The deterministic benchmark is based on JMI. The following example clarifies that these two notions are distinct, even though they coincide in special cases.

Example 1. JMI *does not imply refinement in every* CKC*.*

To see that JMI does not imply refinement in every CKC, consider the information structure described in Figure 3. Both oracles can either withhold all information or fully disclose their information, thereby ensuring that all players become fully informed of the realized state. In fact, these are all the possible experiments of Oracle 2. On the other hand, Oracle 1 can also 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. Thus, Oracle 1 is JMI than Oracle 2, while neither of the two partitions is finer than the other.

Another aspect of this example, which resonates with the key insight of the general analysis, is that there exists a stochastic experiment τ_2 that Oracle 1 cannot imitate. Specifically, consider the stochastic experiment τ_2 given in Figure 4. One can verify that there exists no τ_1 that yields the same profiles of posteriors as the stated experiment τ_2, and this hinges on the fact that F_1 does not refine F_2.

The key issue is that in the deterministic case each state is associated with a *unique* public signal, so JMI guarantees coincidence of the entire profile of posteriors and hence of the induced Bayesian game. Under stochastic experiment, however, each state can generate multiple signals with *different weights*, so the same partitions can induce different joint posteriors. This richer structure is not fully captured by players' interim partitions (i.e.,

[12] In the extended version of the paper, we construct a finite game where failure to meet the JMI condition leads to different equilibrium payoffs.

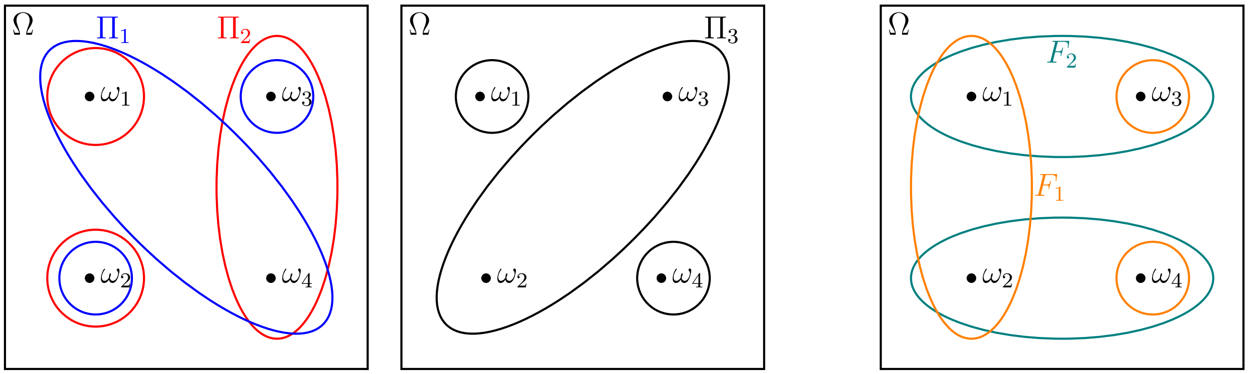

(a) The players' information

(b) The oracles' information

Figure 3: 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 JMI than F_2 whereas the converse is not true, because Oracle 2 cannot replicate the partition $F_1' = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}\}$.

given any deterministic information conveyed by the oracles), and creates both within-CKC and across-CKC difficulties that require stronger conditions than JMI.

$\tau_2(s\|\omega)$	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 4: A stochastic F_2-measurable experiment of Oracle 2.

Example 2. *Refinement in every* CKC *does not imply* JMI*.*

To demonstrate that refinement in every CKC does not imply JMI, 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 5. Observe that in every CKC, F_1 refines F_2.

Now consider a completely revealing, deterministic experiment 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 an experiment 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 experiment 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 JMI 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 JMI.

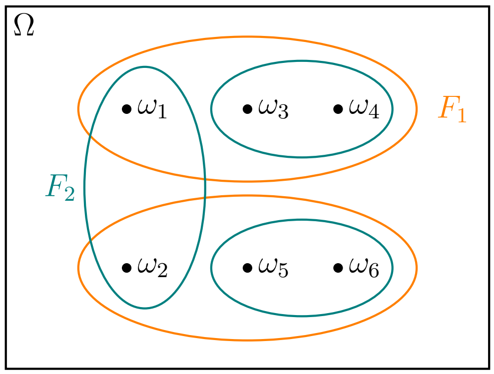

Figure 5: Refinement in every CKC does not imply JMI. Suppose $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_4, \omega_5\}, \{\omega_3, \omega_6\}\}$ and $\Pi_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, \{\omega_5, \omega_6\}\}$. There are two CKCs, $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4, \omega_5, \omega_6\}$. 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 jointly more informative than F_2.

7.3.1 Two-sided JMI implies equivalence in every CKC

Though we substantiated that an JMI 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 corollary provides this equivalence by stating that, given a specific CKC, both oracles dominate each other if and only if their partitions coincide.

Corollary 3. *Fix a unique* CKC*. Then, Oracle i is* JMI *than Oracle $-i$ for every i if and only if $F_1 = F_2$.*

By applying the result within each CKC, the corollary asserts that the partitions F_1 and F_2 are equivalent in every CKC if and only if they are mutually JMI within that CKC, given any *fixed* set of players' partitions. 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 the stochastic analysis above.

8 Conclusion

This paper develops a comparative theory of generators of public information in incomplete-information games. We study an external oracle endowed with a partition of the state space who provides public signals to players holding heterogeneous private information. Signals must be measurable with respect to the oracle's partition, so each signaling scheme induces a Blackwell experiment. Fixing players' information structures and the prior, we introduce a dominance relation over oracles based on their ability to replicate each other's sets of equilibrium outcome distributions across all games. This extends Blackwell (1951) from decision problems to incomplete-information games and from single experiments to generators of experiments. Oracle dominance further connects to Aumann's theory of common knowledge through the central roles of CKCs and information loops.

The main conceptual object is the information loop. Local refinement inside each CKC is necessary for stochastic dominance, but it is not sufficient once the oracle's partition connects several CKCs. Loops create global measurability constraints: the local experiments that mimic another oracle inside different CKCs must be glued into a single oracle-measurable public experiment. Covers and order-preserving covers describe when these loop constraints can be reproduced by the other oracle.

Using this geometry, we characterize oracle equivalence (mutual dominance) via two-sided refinement and order-preserving covers (Theorem 1). For one-sided stochastic dominance, we prove that dominance implies refinement within each CKC and that every loop of the dominating oracle must be covered by loops of the dominated one; when the relevant cover is unique, it must preserve cyclic order (Theorem 2). Conversely, refinement together with order-preserving covers is sufficient under separated-loop structures. The benchmark cases then follow naturally: with a unique CKC, or with no F_1-loops, dominance reduces to refinement inside CKCs (Corollary 1); non-informative loops and the two-CKC case provide additional tractable environments.

For deterministic experiments, dominance is equivalent to being jointly more informative: Oracle 1 can match every player's posterior partition induced by any deterministic signaling scheme of Oracle 2 (Proposition 3). This deterministic result is a useful boundary case, but the stochastic analysis shows that public information in strategic environments cannot be ranked solely by local informativeness. Global measurability constraints across CKCs, captured by information loops, determine which joint beliefs and hence equilibrium outcomes are implementable.

References

Arieli, I., Y. Babichenko, and F. Sandomirskiy (2024). Feasible conditional belief distributions. arXiv:2307.07672.

Arieli, I., Y. Babichenko, F. Sandomirskiy, and O. Tamuz (2021, September). Feasible Joint Posterior Beliefs. *Journal of Political Economy 129*(9), 2546–2594. Publisher: The University of Chicago Press.

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.

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.

Brandenburger, A. and E. Dekel (1993, February). Hierarchies of Beliefs and Common Knowledge. *Journal of Economic Theory 59*(1), 189–198.

Brooks, B., A. Frankel, and E. Kamenica (2024, September). Comparisons of Signals. *American Economic Review 114*(9), 2981–3006.

Cooper, J. and U. Okur (2025, April). Multigraphs with Unique Partition into Cycles. arXiv:2504.08083 [math.CO].

Forges, F. (1993, November). Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision 35*(3), 277–310.

Good, I. J. (1952, January). Rational Decisions. *Journal of the Royal Statistical Society: Series B (Methodological) 14*(1), 107–114.

Gossner, O. (2000, January). Comparison of Information Structures. *Games and Economic Behavior 30*(1), 44–63.

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

Kamenica, E. and M. Gentzkow (2011, October). Bayesian Persuasion. *American Economic Review 101*(6), 2590–2615.

Lagziel, D. and E. Lehrer (2025, October). Constrained Mediation: Bayesian Implementability of Joint Posteriors. arXiv:2510.20986 [econ].

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.

Maschler, M., S. Zamir, and E. Solan (2013). *Game theory.* Cambridge University Press.

Mertens, J. F. and S. Zamir (1985, March). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory 14*(1), 1–29.

Monderer, D. and D. Samet (1989, June). Approximating common knowledge with common beliefs. *Games and Economic Behavior 1*(2), 170–190.

Peski, M. (2008, March). Comparison of information structures in zero-sum games. *Games and Economic Behavior 62*(2), 732–735.

Savage, L. J. (1971). Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association 66*(336), 783–801.

A Proofs

Appendix A is organized as follows. We first isolate three primitive tools: posterior mimicry, local refinement inside a single CKC, and the graph-theoretic form of loop compatibility. The proofs of the main results then cite these primitive tools when needed.

A.1 Posterior mimicry and local refinement

Lemma (1)**.** *Oracle* 1 *dominates Oracle* 2 *if and only if for every experiment* τ_2 *there exists an experiment* τ_1 *such that* $\mu_{\tau_1} = \mu_{\tau_2}$*.*

Proof. If $\mu_{\tau_1} = \mu_{\tau_2}$, then in every game the two guided games have the same distribution over interim belief profiles. Hence they generate the same set of equilibrium outcome distributions.

Conversely, consider the auxiliary game G_{KL} in which each player i reports a belief $p_i \in \Delta(\Omega)$ and receives payoff $u_i(p_i, \omega) = \log p_i(\omega)$ for $\omega \in \Omega$.[13] This logarithmic scoring rule is strictly proper (see, e.g.,Good (1952) and Savage (1971), among others), so the unique best reply of player i is her posterior belief. Thus, for any oracle experiment τ, the induced equilibrium outcome distribution is the distribution μ_τ over joint posterior profiles. Dominance in this particular game therefore implies that, for every τ_2, some F_1-measurable τ_1 satisfies $\mu_{\tau_1} = \mu_{\tau_2}$. □

The following separation device is used only for necessity statements. It is a compact version of the proportionality argument used throughout the text.

Lemma 2 (Proportionality lemma)**.** *Fix two distinct signals $\{s_1, s_2\}$ and assume that the partition $F_2 = \{A_1, A_2, \ldots, A_m\}$ has m elements. Let $p_1, p_2, \ldots, p_m$ be m distinct probabilities such that all ratios of two distinct numbers from the set $\mathbb{A} = \{p_j, 1 - p_j : j = 1, 2, \ldots, m\}$ are pairwise different.*[14] *Define*

$$\tau_2(s_1|A_j) = 1 - \tau_2(s_2|A_j) = p_j, \quad \forall 1 \leq j \leq m. \tag{2}$$

Assume a unique CKC. *If* Post$(\tau_1) \subseteq$ Post(τ_2), *then for every signal $t \in$* Supp(τ_1) *there exists a signal $s \in \{s_1, s_2\}$ and a constant $c_t > 0$ such that $\tau_1(t|\omega) = c_t \tau_2(s|\omega)$ for every $\omega \in \Omega$.*

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

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

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 $(\omega_1, \omega_2, \omega_3, \ldots, \omega^*)$ such that every two adjacent states are in the same partition element for some player. Using the definition of τ_2, it follows that in every joint posterior $(\mu^l_{\tau_2|\omega,s_i})_{l \in N} \in \mathrm{Post}(\tau_2)$, the coordinates relating to $\Pi_l(\omega)$ are strictly positive (for every player l and every signal s_i). Because $\mathrm{Post}(\tau_1) \subseteq \mathrm{Post}(\tau_2)$, it follows that $\tau_1(t|\omega) > 0$ for every $\omega \in \{\omega_1, \omega_2, \ldots, \omega^*\}$, as well.

[13]The "KL" stands for the *Kullback–Leibler divergence* as this scoring rule builds on relative entropy.

[14]To achieve this, one can consider m distinct prime numbers $r_1 < r_2 < \cdots < r_m$. 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 $p_j \in \mathbb{T}_j \setminus \mathbb{T}_{j-1}$.

Using Bayes' rule and for every $\omega, \omega' \in \Pi_l(\omega'')$, we get

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

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 $\mu^l_{\tau_2|\omega'',s_i}(\omega) \cdot \mu^l_{\tau_2|\omega'',s_i}(\omega') > 0$, there exists a unique $c \in \{\frac{x}{y} : \ x, y \in \mathbb{A}\} \cup \{1\}$ such that

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

In case $c = 1$, then the last equation holds for every signal s_i.

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

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

and

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

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

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

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 $(\omega_1, \omega_2, \omega_3, \ldots, \omega^*)$ to get

$$\tau_2(s_i|\omega_1) = c_1 \cdot \tau_2(s_i|\omega_2) = \cdots = [\Pi_{k\geq 1} c_k] \cdot \tau_2(s_i|\omega^*), \tag{4}$$

$$\tau_1(t|\omega_1) = c_1 \cdot \tau_1(t|\omega_2) = \cdots = [\Pi_{k\geq 1} c_k] \cdot \tau_1(t|\omega^*). \tag{5}$$

Dividing Equation (5) by Equation (4), 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). $\square$

The last supporting lemma states that refinement within every CKC is a necessary condition for dominance.

Lemma 3 (Single-CKC refinement)**.** *Assume that C is a CKC. The following are equivalent:*

(i) $F_1|_C$ *refines* $F_2|_C$*;*

(ii) for every $F_2|_C$-measurable experiment τ_2, there exists an $F_1|_C$-measurable experiment τ_1 inducing the same distribution over posterior profiles conditional on C.

Consequently, if $F_1 \succeq_{\text{NE}} F_2$, then F_1 refines F_2 in every CKC.

Proof. If $F_1|_C$ refines $F_2|_C$, every $F_2|_C$-measurable experiment is already $F_1|_C$-measurable, so the claim is immediate. Conversely, suppose $F_1|_C$ does not refine $F_2|_C$. Then there are $\omega, \bar{\omega} \in C$ with $F_1(\omega) = F_1(\bar{\omega})$ but $F_2(\omega) \neq F_2(\bar{\omega})$. Apply Lemma 2 inside C. Any $F_1|_C$-measurable mimicker would have to assign proportional likelihoods to the separating F_2-signal, but F_1-measurability forces equal likelihoods at ω and $\bar{\omega}$, a contradiction. The final statement in the lemma follows from Lemma 1 applied to experiments supported on a single CKC. □

A.2 Loop algebra

The cover condition has a useful algebraic characterization. The next definition is mainly a test for coverability: it asks whether every binary F_{-i}-measurable split of the loop has the same number of transitions in the two directions.

Definition 7. *An F_i-loop $(\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2, \ldots, \omega_m, \overline{\omega}_m)$ is F_{-i}-*balanced *if for every pair of disjoint F_{-i}-measurable sets A and B such that $\cup_{j=1}^{m}\{\omega_j, \overline{\omega}_j\} \subseteq A \cup B$, it follows that*

$$\#(A \to B) := \big|\{j : \omega_j \in A \text{ and } \overline{\omega}_j \in B\}\big| = \big|\{j : \omega_j \in B \text{ and } \overline{\omega}_j \in A\}\big| =: \#(B \to A). \quad (6)$$

Figure 6(a) depicts an F_1-loop with three CKCs. Panel (b) shows a binary F_2-measurable split of this loop into $A = \{\omega_1, \omega_2, \omega_3\}$ and $B = \{\overline{\omega}_1, \overline{\omega}_2, \overline{\omega}_3\}$. Since $\#(A \to B) = 3$ and $\#(B \to A) = 0$, the loop is not F_2-balanced and therefore, by Proposition 4, not F_2-covered.

Proposition 4 (Cover criterion)**.** *Let $(\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2, \ldots, \omega_m, \overline{\omega}_m)$ be an F_1-loop. The following statements are equivalent:*

i. *The loop is F_2-balanced;*

ii. *The loop is F_2-covered;*

iii. *For every F_2-measurable function $f : \{\omega_1, \overline{\omega}_1, \omega_2, \overline{\omega}_2, \ldots, \omega_m, \overline{\omega}_m\} \to (0, \infty)$,*

$$\prod_{j=1}^{m} \frac{f(\omega_j)}{f(\overline{\omega}_j)} = 1.$$

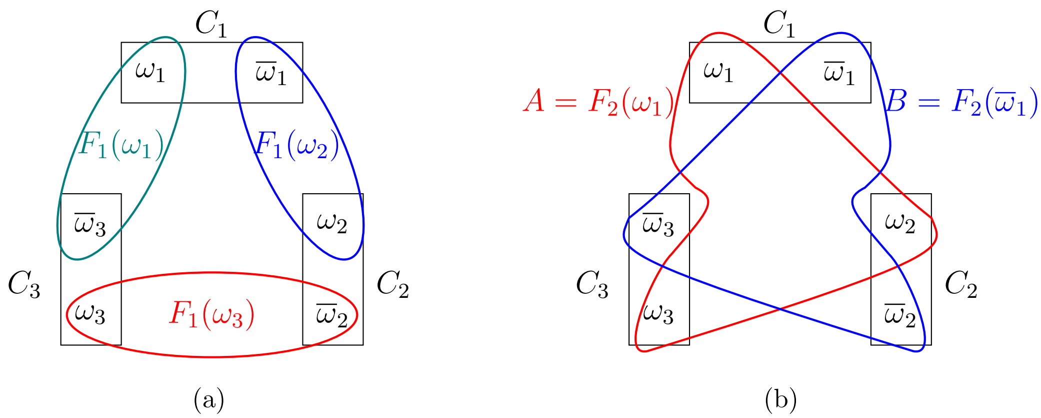

Figure 6: Figure (a) depicts an F_1-loop with three CKCs and six states overall. Figure (b) illustrates how the F_1-loop, presented in (a), is non-balanced with respect to F_2. Namely, F_2 has two elements $A = \{\omega_1, \omega_2, \omega_3\}$, and $B = \{\overline{\omega}_1, \overline{\omega}_2, \overline{\omega}_3\}$ such that the number of transitions from A to B are 3, while the reverse equals 0. A broader discussion on this figure and motivation is given in Appendix B.2.

The third part of Proposition 4, combined with Eq. (1), shows that covers are exactly the configurations under which every F_2-measurable likelihood-ratio product is neutral around the loop, and are therefore necessary for dominance.

Proof. **iii** $\Rightarrow$ **i**. Suppose that $L = (\omega_1, \bar{\omega}_1, \ldots, \omega_m, \bar{\omega}_m)$ is not F_2-balanced. Then there is an F_2-measurable partition $\{A, B\}$ of the states of the loop such that $\#(A \to B) \neq \#(B \to A)$. Define $f(\omega) = \mathbf{1}_{\{\omega \in A\}} + 2 \cdot \mathbf{1}_{\{\omega \in B\}}$. Then $\prod_{j=1}^{m} \frac{f(\omega_j)}{f(\bar{\omega}_j)} = \frac{1}{2}^{\#(A \to B)} 2^{\#(B \to A)} \neq 1$, contradicting **iii**.

i $\Rightarrow$ **ii**. Assume that L is F_2-balanced. For any F_2-atom D that meets the loop, apply balancedness to the partition $\{D, D^c\}$. Since $|\{j : \omega_j \in D\}| = \#(D \to D) + \#(D \to D^c)$ and $|\{j : \bar{\omega}_j \in D\}| = \#(D^c \to D) + \#(D \to D)$, balancedness implies $|\{j : \omega_j \in D\}| = |\{j : \bar{\omega}_j \in D\}|$. Call this condition atom-balance.

Let $J = \{j : \omega_j \in F_2(\bar{\omega}_j)\}$. These indices are already covered by the trivial part of the cover. Remove them from the list. Since each removed pair has both states in the same F_2-atom, atom-balance continues to hold for the remaining pairs.

If no pairs remain, we are done. Otherwise choose one remaining pair $(\omega_{j_1}, \bar{\omega}_{j_1})$. Since $j_1 \notin J$, $\omega_{j_1} \notin F_2(\bar{\omega}_{j_1})$. By atom-balance, the atom $F_2(\bar{\omega}_{j_1})$ contains some state ω_{j_2} from a remaining pair. Thus $\omega_{j_2} \in F_2(\bar{\omega}_{j_1})$. Repeating the same argument gives a sequence $j_1, j_2, \ldots$ such that $\omega_{j_{r+1}} \in F_2(\bar{\omega}_{j_r})$ for every r. Since there are finitely many remaining indices, some index eventually repeats. The first closed segment obtained in this way is an F_2-loop.

Remove the indices of this F_2-loop. Removing an F_2-loop preserves atom-balance, because from each F_2-atom it removes the same number of ω-states and $\bar{\omega}$-states. Repeating the procedure decomposes all remaining pairs into F_2-loops. Together with J, this gives an F_2-cover of L.

ii $\Rightarrow$ **iii**. Assume that L is F_2-covered. Then $\{1, \dots, m\}$ is partitioned into $J, I_1, \dots, I_r$, where $J = \{j : \omega_j \in F_2(\bar{\omega}_j)\}$ and each I_t forms an F_2-loop. Let f be positive and F_2-measurable. If $j \in J$, then $f(\omega_j) = f(\bar{\omega}_j)$, so $\frac{f(\omega_j)}{f(\bar{\omega}_j)} = 1$. If I_t is an F_2-loop, then $\prod_{j \in I_t} \frac{f(\omega_j)}{f(\bar{\omega}_j)} = 1$, because the denominator at each step equals the numerator at the next step around the loop. Hence $\prod_{j=1}^{m} \frac{f(\omega_j)}{f(\bar{\omega}_j)} = 1$. This proves **iii**. □

A.3 Proof of Proposition 1

Proof. We prove the five claims in order.

First, suppose L_i contains two non-adjacent pairs $(\omega_a, \bar{\omega}_a)$ and $(\omega_b, \bar{\omega}_b)$ in the same CKC. By switching the partners inside that CKC, the two outside arcs between the pairs close to two shorter F_i-loops. Hence, the original loop has a strict sub-loop and is not irreducible.

Second, suppose L_i is irreducible, has at least three CKC-pairs, and contains a non-informative pair $(\omega_\ell, \bar{\omega}_\ell)$ with $F_i(\omega_\ell) = F_i(\bar{\omega}_\ell)$. If deleting this pair from the loop creates two neighboring pairs in the same CKC, the first paragraph applies. Otherwise, the remaining pairs still form an F_i-loop, contradicting irreducibility. Thus an irreducible loop with at least six states is F_i-fully informative.

Third, if L_i is F_i-informative, remove all non-informative pairs and, if a CKC appears more than once, split the loop as in the first paragraph and keep a component containing an informative pair. Iterating gives an F_i-fully informative sub-loop.

Fourth, let L_i be F_i-fully informative. If it is not irreducible, choose a strict sub-loop. The omitted pairs form one or more complementary arcs; each arc closes because the sub-loop uses an F_i-connection (the same F_i-tom on both sides of the arc) that skips over it. Hence, L_i decomposes into shorter fully informative F_i-loops. Repeating the operation terminates and yields irreducible F_i-loops.

Finally, assume L_i is not irreducible and does not intersect any CKC more than once. Let L' be a strict sub-loop. Since no CKC is repeated, L' omits at least one pair of the original loop. Therefore some F_i-connection in L' skips over an omitted arc: for some j and $k \neq j+1$, $F_i(\bar{\omega}_j) = F_i(\omega_k)$. In the original loop, $F_i(\bar{\omega}_j) = F_i(\omega_{j+1})$ and $F_i(\bar{\omega}_{k-1}) = F_i(\omega_k)$. Hence, the four states $\{\bar{\omega}_j, \omega_{j+1}, \omega_k, \bar{\omega}_{k-1}\}$ lie in the same F_i-atom. □

A.4 Necessity of loop covers

Lemma 4 (Cover necessity). *If $F_1 \succeq_{\text{NE}} F_2$, then every F_1-loop is F_2-covered.*

Proof. Suppose, to the contrary and using Proposition 4, that an F_1-loop is not F_2-balanced. This means that there is an F_2-measurable partition $\{A, B\}$ of these states such that Eq.

(6) is not satisfied. Define an F_2-measurable experiment that obtains two signals, α and β,

$$\tau_2(\alpha|\omega) = \begin{cases} x, & \text{if } \omega \in A, \\ y, & \text{if } \omega \in B, \end{cases} \tag{7}$$

and $\tau_2(\beta|\omega) = 1 - \tau_2(\alpha|\omega)$. On other states (outside the loop), 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 $\mathrm{Post}(\tau_1) \subseteq \mathrm{Post}(\tau_2)$, 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, $\frac{\tau_1(s|\omega_i)}{\tau_1(s|\overline{\omega}_i)} \in \left\{\frac{x}{y}, \frac{1-x}{1-y}, \frac{y}{x}, \frac{1-y}{1-x}\right\}$. Therefore,

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

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

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

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. (6) is satisfied. A contradiction. ☐

Lemma 5 (Unique covers are order-preserving)**.** *Assume $F_1 \succeq_{\mathrm{NE}} F_2$. If an F_1-loop has a unique F_2-cover, then that cover is order-preserving.*

Proof. Let $I = \{1, \ldots, m\}$. For each $i \in I$, write $\alpha(i) = F_2(\omega_i)$ and $\beta(i) = F_2(\bar{\omega}_i)$. Form the directed atom multigraph H whose vertices are the relevant F_2-atoms and whose edge e_i goes from $\alpha(i)$ to $\beta(i)$. Thus an F_2-sub-loop in the cover is exactly a directed cycle in H, while indices in $J = \{i : \omega_i \in F_2(\bar{\omega}_i)\}$ are self-loops and play no role in order preservation.

Since the F_2-cover is unique, each connected component of H has a unique partition of its non-self-loop edges into directed cycles. Fix such a connected component K. By Theorem 2.8 of Cooper and Okur (2025), K is a bridgeless cactus digraph and its unique partition is the set $B(K)$ of directed cycles of K. Moreover, the proof of Lemma 2.7 of Cooper and Okur (2025) implies that, for every directed cycle $\beta \in B(K)$, each weakly connected component of $K \setminus E(\beta)$ meets β in exactly one vertex.

Fix a directed cycle $\beta \in B(K)$, and let $U \subseteq I$ be the set of indices of the edges of β. The corresponding F_2-sub-loop is irreducible (otherwise the cover is not unique). Suppose, toward a contradiction, that this sub-loop is not order-preserving relative to the cyclic order of L_1. After cyclically renaming the relevant indices, write the selected three indices as $1 < j < k$ in the F_1-order, with ω_k preceding ω_j in the F_2-order of β.

Partition the states of β into three arcs. Let A_1^2 be the states of β from $\bar{\omega}_1$ to ω_k, following the F_2-order of β; let A_k^2 be the states from $\bar{\omega}_k$ to ω_j; and let A_j^2 be the remaining states, from $\bar{\omega}_j$ to ω_1. Similarly, partition the states of L_1 into the three F_1-arcs A_1^1, from $\bar{\omega}_1$ to ω_j; A_j^1, from $\bar{\omega}_j$ to ω_k; and A_k^1, from $\bar{\omega}_k$ to ω_1.

Choose $p_1 = \frac{1}{7}$, $p_2 = \frac{2}{7}$, $p_3 = \frac{4}{7}$, and $p_4 = \frac{1}{2}$. Define a binary F_2-measurable experiment τ_2 with distinct signals s_1, s_2 by setting $\tau_2(s_1|\omega) = p_r$ whenever the F_2-atom of ω meets A_r^2, for $r \in \{1, k, j\}$, where $p_k := p_2$ and $p_j := p_3$. For each weak component C of $K \setminus E(\beta)$, assign to every atom of C the same probability as its unique attachment atom on β. On atoms outside K, set $\tau_2(s_1|\cdot) = p_4$. Finally set $\tau_2(s_2|\cdot) = 1 - \tau_2(s_1|\cdot)$. This is well-defined and F_2-measurable, and by the construction, for every $\ell \notin \{1, j, k\}$ and every $s \in \{s_1, s_2\}$, we have $\tau_2(s|\omega_\ell) = \tau_2(s|\bar{\omega}_\ell)$.

By dominance and Lemma 1, there is an F_1-measurable experiment τ_1 with $\text{Post}(\tau_1) \subseteq \text{Post}(\tau_2)$. Fix a signal $t \in \text{Supp}(\tau_1)$. By Lemma 2, in each CKC of the loop the likelihood ratio of t under τ_1 is equal to the likelihood ratio of either s_1 or s_2 under τ_2. Hence $\tau_1(t|\omega_\ell) = \tau_1(t|\bar{\omega}_\ell)$ for every $\ell \notin \{1, j, k\}$. Together with F_1-measurability along L_1, this implies that, on the states of L_1, the likelihood vector of t is constant on the three F_1-arcs: write these constants as a_t on A_1^1, b_t on A_j^1, and c_t on A_k^1.

At the three exceptional CKCs, the proportionality lemma gives $\frac{c_t}{a_t} \in \{\frac{p_3}{p_1}, \frac{1-p_3}{1-p_1}\}$, $\frac{a_t}{b_t} \in \{\frac{p_2}{p_3}, \frac{1-p_2}{1-p_3}\}$, and $\frac{b_t}{c_t} \in \{\frac{p_1}{p_2}, \frac{1-p_1}{1-p_2}\}$. With the chosen probabilities, these three pairs of possible ratios are $\{4, \frac{1}{2}\}$, $\{\frac{1}{2}, \frac{5}{3}\}$, and $\{\frac{1}{2}, \frac{6}{5}\}$, respectively. Since $\frac{c_t}{a_t} \cdot \frac{a_t}{b_t} \frac{b_t}{c_t} = 1$, the only possible choices are either $\left(4, \frac{1}{2}, \frac{1}{2}\right)$ or $\left(\frac{1}{2}, \frac{5}{3}, \frac{6}{5}\right)$.

Call signals of the first kind type 1, and signals of the second kind type 2. Let x and y denote the total probability mass of types 1 and 2 at ω_1 respectively. The row-sum constraint at ω_1 gives $x + y = 1$. The row-sum constraint at $\bar{\omega}_1$ gives $\frac{p_1}{p_3}x + \frac{1-p_1}{1-p_3}y = 1$, namely $\frac{1}{4}x + 2y = 1$. Together with $x + y = 1$, this implies $x = 4/7$. The row-sum constraint at $\bar{\omega}_j$ gives $\frac{p_1}{p_2}x + \frac{1-p_1}{1-p_2}y = 1$, namely $\frac{1}{2}x + \frac{6}{5}y = 1$. Together with $x + y = 1$, this implies $x = 2/7$. This contradiction shows that β must be order-preserving. Since $\beta \in B(K)$ was arbitrary, every directed cycle in the unique cycle partition of K is order-preserving. Since K was arbitrary, the unique F_2-cover of L_1 is order-preserving. □

A.5 Constructive sufficiency tools

Definition 8 (Posterior-equivalent extension of a binary experiment)**.** *Let $a, b \in \Delta(S_1)$ be two probability laws on a finite signal set S_1. Let $q, q' \in \Delta(S_2)$ be two probability laws on* *another finite signal set S_2. We say that the ordered pair (q, q') is a* posterior-equivalent extension *of (a, b) if one of the following holds:*

1. *There exist a finite set R, a probability law $\nu \in \Delta(R)$, and an identification $S_2 = S_1 \times R$ such that $q(s, r) = a(s)\nu(r)$ and $q'(s, r) = b(s)\nu(r)$ for all $(s, r) \in S_1 \times R$.*

2. *$a = b$ and $q = q'$.*

In case (1), the coordinate r is common noise. In case (2), both experiments are uninformative in the two-state cell, and hence induce the same posterior as each other with total probability one.

Lemma 6 (Tree mimicry)**.** *Let $T = (V, E)$ be a finite directed tree. For each edge $e = (u, v) \in E$, fix a finite signal set S_e and two laws $\alpha_e, \beta_e \in \Delta(S_e)$. Then there exist a finite signal set S_T and laws $q_v \in \Delta(S_T)$, $v \in V$, such that for every edge $e = (u, v)$, the ordered pair (q_u, q_v) is a posterior-equivalent extension of (α_e, β_e).*

Proof. Let $S_T = \prod_{e \in E} S_e$. For an edge $e = (u, v)$, deleting e leaves two connected components, so denote by T_e^- the component containing u and by T_e^+ the component containing v. For $x \in V$, define $q_x \in \Delta(S_T)$ by

$$q_x((s_f)_{f \in E}) = \prod_{e \in E} \begin{cases} \alpha_e(s_e), & x \in T_e^-, \\ \beta_e(s_e), & x \in T_e^+. \end{cases}$$

This is a probability law because it is a product of probability laws. Fix $e = (u, v)$. If $f \neq e$, then u and v lie in the same component of $T \setminus \{f\}$, so the f-coordinate has the same law under q_u and q_v. On coordinate e, the law is α_e under q_u and β_e under q_v. Hence (q_u, q_v) is (α_e, β_e) multiplied by the same law on the coordinates $E \setminus \{e\}$, and is therefore a posterior-equivalent extension. □

Lemma 7 (Single-loop construction)**.** *Let $L = (\omega_1, \bar{\omega}_1, \ldots, \omega_m, \bar{\omega}_m)$ be an F_1-loop, with an order-preserving F_2-cover $\{1, \ldots, m\} = J \cup I_1 \cup \cdots \cup I_K$, where $J = \{i : \omega_i \in F_2(\bar{\omega}_i)\}$. Let τ_2 be an F_2-measurable experiment on a finite set S_2. Assume that the F_1-atoms $F_1(\bar{\omega}_t) = F_1(\omega_{t+1})$, $t = 1, \ldots, m$, are distinct. Then there is an F_1-measurable experiment τ_1 on S_2^K such that, for every i, the ordered pair $(\tau_1(\cdot|\omega_i), \tau_1(\cdot|\bar{\omega}_i))$ is a posterior-equivalent extension of $(\tau_2(\cdot|\omega_i), \tau_2(\cdot|\bar{\omega}_i))$.*

Proof. For each i, set $p_i := \tau_2(\cdot|\bar{\omega}_i) \in \Delta(S_2)$. If $i \in J$, then $\omega_i \in F_2(\bar{\omega}_i)$, so $\tau_2(\cdot|\omega_i) = p_i$. If $i \in I_k$, let $\mathrm{pred}(i)$ be the predecessor of i in the F_2-sub-loop indexed by I_k. Since the cover is order-preserving, $\tau_2(\cdot|\omega_i) = p_{\mathrm{pred}(i)}$ and $\tau_2(\cdot|\bar{\omega}_i) = p_i$.

For each block I_k, write its elements in the cyclic order inherited from L. For $t = 1, \ldots, m$, let $\ell_k(t)$ be the last element of I_k weakly before t in that cyclic order, and if there is no such element before t, let $\ell_k(t)$ be the last element of I_k. Define $q_t \in \Delta(S_2^K)$ by

$$q_t(x_1, \ldots, x_K) = \prod_{k=1}^{K} p_{\ell_k(t)}(x_k).$$

Now define τ_1 on the F_1-atoms in the loop by $\tau_1(\cdot|\bar{\omega}_t) = \tau_1(\cdot|\omega_{t+1}) = q_t$. Since these F_1-atoms are distinct, this assignment is unambiguous. Extend τ_1 arbitrarily, but F_1-measurably, on F_1-atoms not appearing in the loop.

Fix i. Under τ_1, the two laws in the CKC-pair $(\omega_i, \bar{\omega}_i)$ are q_{i-1} and q_i. If $i \in I_k$, then order preservation gives $\ell_k(i-1) = \text{pred}(i)$ and $\ell_k(i) = i$, while for every $h \neq k$, $\ell_h(i-1) = \ell_h(i)$. Therefore q_{i-1} and q_i differ only in coordinate k: on that coordinate the laws are $p_{\text{pred}(i)} = \tau_2(\cdot|\omega_i)$ and $p_i = \tau_2(\cdot|\bar{\omega}_i)$, and all other coordinates have the same law under both states. If $i \in J$, then $i \notin I_k$ for every k, so $\ell_k(i-1) = \ell_k(i)$ for every k. Hence $q_{i-1} = q_i$ as $\tau_2(\cdot|\omega_i) = \tau_2(\cdot|\bar{\omega}_i)$. Thus (q_{i-1}, q_i) is a posterior-equivalent extension of $(\tau_2(\cdot|\omega_i), \tau_2(\cdot|\bar{\omega}_i))$. □

A.6 Proof of Theorem 1

Proof. Assume first that F_1 and F_2 are equivalent. By Lemma 3 and Lemma 4, applied in both directions, for each i, F_i refines F_{-i} in every CKC and every F_i-loop is F_{-i}-covered. Since refinement holds in both directions, $F_1|_C = F_2|_C$ for every CKC C.

Fix an F_i-loop L_i. We show that L_i admits an order-preserving F_{-i}-cover. Consider first the case in which $L_i = (\omega_1, \bar{\omega}_1, \ldots, \omega_m, \bar{\omega}_m)$ is irreducible with at least six states. By Proposition 1, L_i is F_i-fully informative. Consider an F_{-i}-cover of L_i. No index can belong to the non-informative set $J = \{j : \omega_j \in F_{-i}(\bar{\omega}_j)\}$, because $F_i|_C = F_{-i}|_C$ in every CKC and L_i is F_i-fully informative. If the cover contained a proper F_{-i}-sub-loop, then Lemma 4, applied in the reverse direction, would imply that this sub-loop is F_i-covered. Since no pair is F_i-non-informative, this F_i-cover would contain a proper F_i-sub-loop of L_i, contradicting irreducibility. Hence the cover consists of a single F_{-i}-loop using exactly the pairs of L_i, and by Lemma 5, it is also order-preserving.

For irreducible L_i with four states, order preservation is vacuous. Now let L_i be reducible. Decompose L_i into irreducible F_i-sub-loops by repeatedly cutting either at repeated CKCs or at repeated F_i-atoms. Each cut produces shorter loops whose cyclic orders are inherited from L_i, and the process terminates. By the previous cases, each irreducible component admits an order-preserving F_{-i}-cover. Moreover, the cuts at repeated CKCs are valid for both partitions because $F_i|_C = F_{-i}|_C$, while cuts at repeated F_i-atoms correspond to the same original F_i-connection and therefore only determine how the components are concatenated back into the original cyclic order. Hence the component covers can be concatenated according to the same cyclic order as the components in L_i. The resulting family of F_{-i}-loops covers all pairs of L_i and preserves their cyclic order. Thus L_i admits an order-preserving F_{-i}-cover.

Conversely, assume the conditions in the theorem. We prove $F_1 \succeq_{\text{NE}} F_2$; the proof of $F_2 \succeq_{\text{NE}} F_1$ is symmetric. The two-sided local refinement condition gives $F_1|_C = F_2|_C$ for every CKC C. Call an F_i-loop *type-2 irreducible* if it is F_i-fully informative and no F_i-atom contains four states of the loop.

We first show that every type-2 irreducible F_1-loop is an F_2-loop. Let L be a type-2 irreducible F_1-loop. If L is irreducible in the ordinary sense, then the order-preserving F_2-cover of L guaranteed by the theorem's hypothesis cannot contain a non-informative pair, because $F_1|_C = F_2|_C$ and L is F_1-fully informative. It also cannot contain a proper F_2-sub-loop, because applying the hypothesis in the reverse direction to that sub-loop would generate a proper F_1-sub-loop of L. Thus the cover consists of a single F_2-loop, so L is an F_2-loop. If L is not irreducible, then by Proposition 1, it must intersect some CKC more than once. Splitting inside such repeated CKCs decomposes L into strict F_1-sub-loops. Iterating this procedure gives irreducible sub-loops (in the ordinary sense), hence F_2-loops by the previous case. All splitting and recombination takes place inside CKCs, where $F_1|_C = F_2|_C$. Therefore the same recombination turns these F_2-sub-loops back into the original loop L, so L is an F_2-loop.

We next focus on a decomposition process. Every F_i-fully informative and reducible loop can be decomposed into type-2 irreducible F_i-loops. Indeed, if an F_i-atom contains four loop states $\{\bar{\omega}_j, \omega_{j+1}, \bar{\omega}_\ell, \omega_{\ell+1}\}$, cut the loop at these two occurrences and reconnect inside that atom. This produces two shorter F_i-fully informative loops. Repeating this procedure terminates and yields type-2 irreducible F_i-loops. Moreover, if an F_i-connection uses only two states of the original loop, the cuts can be chosen so that this connection remains inside one of the resulting type-2 loops.

Now consider a connected family $\mathcal{L}$ of type-2 irreducible F_1-loops, where connected means that any two loops in the family are linked by a chain of type-2 irreducible loops such that consecutive loops share a CKC. We claim that if ω, ω' are states in CKCs visited by loops in $\mathcal{L}$ and $F_1(\omega) = F_1(\omega')$, then $F_2(\omega) = F_2(\omega')$. If ω and ω' lie in the same CKC, this follows from $F_1|_C = F_2|_C$. Otherwise, choose a chain of loops in $\mathcal{L}$ connecting a loop that visits the CKC of ω to a loop that visits the CKC of ω'. Follow the first loop from ω to a CKC shared with the next loop, switch inside that CKC, continue along the chain, and finally close the sequence using the connection $F_1(\omega) = F_1(\omega')$. After deleting any non-informative consecutive pair, this gives an F_1-informative loop preserving the closing connection between ω and ω'. By Proposition 1, this loop contains an F_1-fully informative sub-loop preserving that connection. By the decomposition step above, one of its type-2 irreducible components still preserves the connection, and that component is an F_2-loop. Hence $F_2(\omega) = F_2(\omega')$.

A *cluster* is the union of all CKCs visited by a maximal connected family of type-2 irreducible F_1-loops. The previous paragraph proves that every F_2-measurable experiment restricted to a cluster is also F_1-measurable. Let Ω^* be the partition of Ω into clusters and CKCs that are not contained in any cluster. Note that any two distinct elements of Ω^* are connected by at most one F_1-atom. If two distinct elements were connected by two different F_1-atoms, then we could embed these atoms in an informative loop, which would contain a type-2 irreducible loop (using the two atoms) and thus unify the two clusters into one, contradicting the definition of Ω^*.

We use the following gluing argument. Let A_1, A_2 be disjoint unions of elements of Ω^*. Assume that each A_i is solvable, meaning that every F_2-measurable experiment on A_i can be mimicked on A_i by an F_1-measurable experiment, generating the same posterior-profile distribution. We show that $A_1 \cup A_2$ is also solvable. Fix an F_2-measurable experiment τ_2 on $A_1 \cup A_2$. For $i = 1, 2$, let τ_1^i be an F_1-measurable experiment on A_i with signal space S_i reproducing the posterior-profile distribution induced by $\tau_2|_{A_i}$. Let $\widetilde{A}_i$ be the set of states in A_i whose F_1-atom intersects A_{-i}. By the one-connection condition, $\widetilde{A}_1 \cup \widetilde{A}_2$ is contained in a single F_1-atom. Let λ_i be the law of τ_1^i on $\widetilde{A}_i$ if $\widetilde{A}_i \neq \emptyset$, and choose any law on the signal space of τ_1^i otherwise. On the product signal space, define

$$\tau_1((s_1, s_2) \mid \omega) = \begin{cases} \tau_1^1(s_1 \mid \omega)\lambda_2(s_2), & \omega \in A_1, \\ \lambda_1(s_1)\tau_1^2(s_2 \mid \omega), & \omega \in A_2. \end{cases}$$

This experiment is F_1-measurable. Inside A_i, one coordinate reproduces τ_1^i and the other coordinate is common noise. On the possible connecting F_1-atom, both sides assign the same law $\lambda_1 \otimes \lambda_2$. Therefore the posterior-profile distribution is preserved on both A_1 and A_2, so $A_1 \cup A_2$ is solvable.

Form a graph G^* whose vertices are the elements of Ω^*, with an edge between two vertices whenever some F_1-atom intersects both. We already established that there is at most one edge between any two vertices of G^*, so it is a forest. Namely, if G^* had a cycle, following the corresponding F_1-atoms through the elements of Ω^* would produce an F_1-informative loop crossing distinct elements of Ω^*. As above, this loop would contain a type-2 irreducible F_1-loop crossing those elements, forcing them to belong to a common cluster, a contradiction.

Each vertex of G^*, whether a cluster or a CKC, is solvable. Since G^* is a forest, we can glue its vertices one leaf at a time using the gluing argument above, by induction on the number of vertices of G^*. The case of one and two vertices has already been established. For a forest with $k+1$ vertices, choose a leaf A_{k+1}, and let A be the union of the remaining vertices. Apply the induction hypothesis to A, and then glue A and A_{k+1} as before (the one-connector property ensures that at most one F_1-atom intersects both A and A_{k+1}). Thus all of Ω is solvable.

Hence, for every F_2-measurable experiment τ_2, there is an F_1-measurable experiment τ_1 inducing the same distribution over joint posterior profiles. By Lemma 1, $F_1 \succeq_{\text{NE}} F_2$. The symmetric argument gives $F_2 \succeq_{\text{NE}} F_1$. Therefore F_1 and F_2 are equivalent. □

A.7 Proof of Theorem 2

Proof. For necessity, local refinement in every CKC follows from Lemma 3. The existence of an F_2-cover for every F_1-loop follows from Lemma 4. If the cover is unique, order preservation follows from Lemma 5.

We prove sufficiency. Fix an arbitrary F_2-measurable experiment τ_2 on a finite signal set S_2. Let $\mathcal{C}$ be the CKC partition. We first construct an F_1-measurable experiment τ_1 that induces the same distribution of posteriors for a single decision maker whose information partition is $\mathcal{C}$. For each CKC C and each F_1-atom A with $A \cap C \neq \emptyset$, define $p_{A,C} \in \Delta(S_2)$ by $p_{A,C} := \tau_2(\cdot|\omega)$ for any $\omega \in A \cap C$. This is well-defined because F_1 refines F_2 inside C.

Construct the bipartite incidence graph Γ: its vertices are the CKCs and the F_1-atoms, and A is connected to C iff $A \cap C \neq \emptyset$. By the separated-loop assumption, each connected component of Γ is either a tree or contains one simple cycle, with trees attached to it.

We assign to each F_1-atom A a law q_A such that, for every CKC C, there is a law ν_C with $q_A = p_{A,C} \otimes \nu_C$ for every atom A adjacent to C. This condition means that, inside C, the first coordinate reproduces the experiment τ_2, while the remaining coordinates are common noise.

First consider a tree component. Denote its CKCs by T. Use one coordinate $S_D = S_2$ for each CKC $D \in T$, with signals $s_D \in S_D$. For an F_1-atom A and a CKC $D \in T$, let $\phi_D(A)$ be the unique F_1-atom adjacent to D that lies on the path from D to A. Define

$$q_A((s_D)_D) = \prod_{D \in T} p_{\phi_D(A),D}(s_D).$$

Fix a CKC C. If A is adjacent to C, then the C-coordinate of q_A is $p_{A,C}$. For every $D \neq C$, all F_1-atoms adjacent to C remain connected to C after deleting D. Hence the unique path from D to any atom adjacent to C, leaves D through the same neighboring F_1-atom. Therefore

$\phi_D(A)$ is the same for all F_1-atoms A adjacent to C. Thus all coordinates other than C are common across atoms adjacent to C, and the required property holds on tree components.

Now consider a component with one simple cycle, written as $A_0, C_1, A_1, \ldots, A_{m-1}, C_m, A_0$. This is an F_1-loop with distinct F_1-atoms (otherwise there are connected loops). By assumption, it has an order-preserving F_2-cover. Applying Lemma 7 to this cycle and τ_2 gives a signal set S_0 and laws h_{A_i} for the cycle atoms such that, for each cycle CKC C_i, there is a law η_i with $h_A = p_{A,C_i} \otimes \eta_i$ for the two cycle atoms A adjacent to C_i. For any off-cycle atom A adjacent to the same cycle CKC C_i, set $h_A := p_{A,C_i} \otimes \eta_i$. This is unambiguous because an off-cycle atom adjacent to two cycle CKCs would create a second cycle in Γ.

Delete the cycle CKC vertices and their incident edges. The remaining graph is a forest. Let $\mathcal{R}$ be the set of F_1-atoms R for which h_R has already been defined, i.e., $\mathcal{R}$ contains all cycle atoms and every off-cycle atom adjacent to a cycle CKC. Each tree in the remaining forest has a unique root $R \in \mathcal{R}$ (otherwise two such roots connected by a path in the forest and cycle would create a second cycle in Γ). For each $R \in \mathcal{R}$, let T_R be the tree rooted at R, including the root R itself, and let $\mathcal{D}_R$ be the set of off-cycle CKCs in T_R. Use one coordinate $S_D = S_2$ for each $D \in \mathcal{D}_R$. For every atom $A \in T_R$, define

$$r_A^R((s_D)_{D\in\mathcal{D}_R}) = \prod_{D\in\mathcal{D}_R} p_{\phi_D(A),D}(s_D),$$

where $\phi_D(A)$ is the unique F_1-atom adjacent to D that lies on the path from D to A inside T_R. If $\mathcal{D}_R = \emptyset$, take S_R to be a singleton and r_R^R to be the degenerate law. For the whole component use signal space $S_0 \times \prod_{R\in\mathcal{R}} S_R$, where $S_R := \prod_{D\in\mathcal{D}_R} S_D$. For every atom A, let $R(A)$ be the unique root such that $A \in T_{R(A)}$, and define

$$q_A = h_{R(A)} \otimes r_A^{R(A)} \otimes \bigotimes_{\substack{R\in\mathcal{R} \\ R\neq R(A)}} r_R^R.$$

The required property follows. If C_i is a cycle CKC, then every adjacent atom A has cycle-block law $p_{A,C_i} \otimes \eta_i$, and all tree-coordinate laws are the corresponding root laws, hence common across atoms adjacent to C_i. If C is off the cycle, all atoms adjacent to C belong to the same rooted tree; the cycle block and all other tree blocks are common, while the tree construction gives the C-coordinate $p_{A,C}$. Thus, for every CKC C in the component, there exists ν_C such that $q_A = p_{A,C} \otimes \nu_C$ for every adjacent atom A.

Doing this component by component and taking the product across components gives one finite signal space and one law q_A for every relevant F_1-atom A. Define $\tau_1(\cdot|\omega) := q_{F_1(\omega)}$. Then τ_1 is F_1-measurable. Moreover, for every CKC C, there exists a common-noise law

ν_C such that $q_A = p_{A,C} \otimes \nu_C$ for every F_1-atom A intersecting C. Therefore, in the single-DM environment with information partition $\mathcal{C}$, τ_1 and τ_2 induce the same distribution of posteriors.

Finally, return to the original multi-agent environment. Given the CKC-level posterior $P = \mu(\cdot|C(\omega), s)$ and the realized state ω, player i's posterior is $P(\cdot|\Pi_i(\omega))$. Hence equality of the distribution of CKC-level posteriors implies equality of the distribution of joint posterior profiles. By Lemma 1, $F_1 \succeq_{\text{NE}} F_2$. □

A.8 Proof of Corollary 2

Proof. Collapse, inside each CKC, states that are in the same F_2-atom. Since F_1 refines F_2 in every CKC, this projection preserves all posterior experiments that Oracle 2 can generate. If every F_1-loop is F_2-non-informative, then no projected F_1-loop remains: any projected loop would lift to an F_1-loop containing two states in some CKC that Oracle 2 separates, contradicting non-informativeness. Hence the projected problem has no F_1-loops. Apply the acyclic part of Corollary 1 and lift the mimicking experiment back to Ω. □

A.9 Proof of Proposition 2

Proof. Necessity follows from Lemmas 3 and 4, with Proposition 4. For sufficiency, assume local refinement and balancedness of every F_1-loop. With two CKCs, every irreducible loop has two CKC-pairs (the case of no F_1-loops follows from Corollary 1). By the cover criterion, a balanced two-pair loop is either an F_2-loop or F_2-non-informative. If all loops are non-informative, Corollary 2 applies. Otherwise there is an informative two-pair loop that is also an F_2-loop. Any other F_1-link between the two CKCs can be combined with this loop to form another two-pair loop; balancedness again forces that link to be F_2-measurable. Hence every cross-CKC F_1-constraint is also an F_2-constraint (if a loop exists), and every F_2-measurable experiment is already F_1-measurable across the two CKCs. Dominance follows. □

A.10 Proof of Proposition 3

Proof. If Oracle 1 is JMI than Oracle 2, then for every deterministic F_2-measurable experiment τ_2 there exists a deterministic F_1-measurable experiment τ_1 with $\Pi_i \vee \tau_1 = \Pi_i \vee \tau_2$ for all players i. The guided games are therefore isomorphic and have the same equilibrium outcome distributions.

Conversely, assume dominance and fix a deterministic F_2-measurable experiment τ_2. By Lemma 1, there is an F_1-measurable experiment τ_1 with $\mu_{\tau_1} = \mu_{\tau_2}$. Since τ_2 is deterministic, every posterior profile in its support corresponds to a cell of the partitions $\Pi_i \vee \tau_2$. Equality of the posterior-profile distribution forces τ_1 to generate the same posterior cells, and hence the same joined partitions $\Pi_i \vee \tau_1 = \Pi_i \vee \tau_2$ for all i. Thus Oracle 1 is jointly more informative than Oracle 2. □

A.11 Proof of Corollary 3

Proof. The if direction is immediate. For the converse, suppose the CKC is unique and F_i is JMI than F_{-i} for both i. If $F_1 \neq F_2$, choose $\omega, \bar{\omega}$ with $F_1(\omega) = F_1(\bar{\omega})$ and $F_2(\omega) \neq F_2(\bar{\omega})$. The binary partition that separates $F_2(\omega)$ from its complement must be reproducible by an F_1-measurable partition after joining with every player's partition. Since the CKC is unique, there is a path in the players' information graph from ω to $\bar{\omega}$. Along the first edge of this path at which the binary F_2-partition changes, the joined partitions disagree, contradicting JMI. Hence F_1 refines F_2. The symmetric argument gives the reverse refinement, so $F_1 = F_2$. □

B Online appendix

B.1 More than one CKC

As the benchmark corollaries below show, under a unique CKC, the partition-refinement condition ensures that Oracle 1 can produce the *exact* same experiment as Oracle 2. This conclusion, however, hinges on the existence of a unique CKC. In case there are several CKCs, Oracle 1 may need to follow a different experiment 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 3. *A mimicking experiment τ_1 may require more signals than τ_2.*

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 $F_2 = \{\{\omega_1\}, \{\omega_3\}, \{\omega_2, \omega_4\}\}$. 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.

Consider the experiment τ_2 given in Figure 7. 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 7: A stochastic F_2-measurable experiment of Oracle 2.

The set of τ_2-posteriors Post(τ_2) is

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

where the i^{th} coordinate of e_i is 1, and we can now try to mimic τ_2 using an F_1-measurable experiment. 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 $\left(\left(\frac{3}{7}, \frac{4}{7}, 0, 0\right), e_1\right)$ 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 $\left(e_3, (0, 0, \frac{1}{2}, \frac{1}{2})\right)$ necessitates that $\tau(s|\omega_4) = \alpha$ as well. These conditions are given in Table (a) within Figure 8.

$\tau_1(s\|\omega)$	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 8: An experiment 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 experiment. Thus, in order to mimic the stated experiment τ_2, Oracle 1 requires an additional signal as presented in Table (b), in Figure 8. To conclude, though the oracles' partitions refine one another in every CKC, they cannot always produce the exact same experiment when trying to mimic each other.

Remark 1. *Note that dominance does not imply refinement in general. To see this, consider the information structure* $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}, F_1 = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}\}$ *and* $F_2 =$

$\{\{\omega_1, \omega_2\}, \{\omega_3\}, \{\omega_4\}\}$. For every experiment τ_2, one can devise an experiment τ_1 that yields the same distribution over posterior beliefs. Evidently, Oracle 2 provides the players with no additional information regarding states ω_1 and ω_2, and this allows Oracle 1 to replicate τ_2 on ω_3 and ω_4, accordingly.

B.2 Balancedness and the likelihood-ratio obstruction

This subsection contains the longer motivation for the cover criterion used in Section 4.2. Consider a non-balanced F_1-loop as in Figure 6(b), and suppose that A and B are F_2-measurable sets with $\#(A \to B) \neq \#(B \to A)$. For example, if $A = \{\omega_1, \omega_2, \omega_3\}$ and $B = \{\overline{\omega}_1, \overline{\omega}_2, \overline{\omega}_3\}$, then $\#(A \to B) = 3$ and $\#(B \to A) = 0$. Let s be a signal and define an F_2-measurable experiment by

$$\tau_2(s \mid \omega) = \frac{1}{2} - \frac{1}{4}\mathbf{1}_{\{\omega \in A\}}.$$

Then each transition from A to B contributes the likelihood ratio 1/2, and each transition from B to A contributes the likelihood ratio 2. Hence the product of likelihood ratios around the loop is

$$\prod_j \frac{\tau_2(s \mid \omega_j)}{\tau_2(s \mid \overline{\omega}_j)} = 2^{\#(B\to A) - \#(A \to B)},$$

which is not equal to one whenever the loop is not balanced.

By contrast, any F_1-measurable experiment τ_1 must satisfy

$$\prod_j \frac{\tau_1(s \mid \omega_j)}{\tau_1(s \mid \overline{\omega}_j)} = 1,$$

because $\overline{\omega}_j$ and ω_{j+1} are in the same F_1-atom. Thus Oracle 1 cannot match the likelihood-ratio restrictions generated by such a τ_2. This is the likelihood-ratio obstruction behind Proposition 4: balancedness is exactly the algebraic condition that rules out this obstruction, and the proposition shows that it is equivalent to the geometric condition of being covered by F_2-loops.