8-17-2021

Learning Efficiency of Multi-Agent Information Structures

Mira Frick
Yale University

Ryota Iijima
Yale University

Yuhta Ishii

LEARNING EFFICIENCY OF MULTI-AGENT INFORMATION STRUCTURES

By

Mira Frick, Ryota Iijima, and Yuhta Ishii

August 2021

Learning Efficiency of Multi-Agent Information Structures*

Preliminary and incomplete

Mira Frick   Ryota Iijima   Yuhta Ishii

August 17, 2021

Abstract

We study settings in which, prior to playing an incomplete information game, players observe many draws of private signals about the state from some information structure. Signals are i.i.d. across draws, but may display arbitrary correlation across players. For each information structure, we define a simple *learning efficiency index*, which only considers the statistical distance between the worst-informed player's marginal signal distributions in different states. We show, first, that this index characterizes the speed of common learning (Cripps, Ely, Mailath, and Samuelson, 2008): In particular, the speed at which players achieve approximate common knowledge of the state coincides with the slowest player's speed of individual learning, and does not depend on the correlation across players' signals. Second, we build on this characterization to provide a ranking over information structures: We show that, with sufficiently many signal draws, information structures with a higher learning efficiency index lead to better equilibrium outcomes, robustly for a rich class of games and objective functions. We discuss implications of our results for constrained information design in games and for the question when information structures are complements vs. substitutes.

1 Introduction

1.1 Overview

Suppose a group of players (e.g., firms) are engaged in an incomplete information game (e.g., joint investment in a project of unknown profitability). Prior to choosing their actions in the game, players have access to many draws of private signals about the unknown state from some information structure (capturing, for instance, that data is "cheap" or abundant). While signals are assumed i.i.d. across draws, we allow them to be arbitrarily correlated across players.

This paper studies comparisons of information structures in such a setting, addressing two related questions. First, which information structures induce faster learning? In strategic settings, learning not only concerns each agent's beliefs about the state, but also agents' higher-order uncertainty about other agents' beliefs. Thus, for each information structure, we quantify the speed of *common learning* (Cripps, Ely, Mailath, and Samuelson, 2008), i.e., the speed at which repeated signal draws allow agents to achieve approximate common knowledge of the state. Second, when agents observe a large number of signal draws, which information structures induce "better" equilibrium outcomes? Based on our characterization of the speed of learning, we obtain a ranking over information structures that answers this question. The ranking applies for a rich class of games and objective functions, permitting a robust comparison of information structures that does not require an understanding of the full details of the strategic environment.

We answer both these questions by introducing a *learning efficiency index* for multi-agent information structures. An information structure maps each state to a joint distribution over all agents' private signals, where both states and signals are assumed finite. Our index reduces each information structure to a simple one-dimensional measure, which only quantifies how difficult the worst-informed agent finds it to distinguish the two states that are hardest to tell apart based on her private signal observations. Here, each agent i's difficulty of distinguishing any two states is measured by the (Chernoff) statistical distance between i's marginal signal distributions in each state. Notably, since the learning efficiency index is derived only from agents' marginal signal distributions, it does not depend on the correlation across agents' signals.

Our first main result is that this index characterizes agents' speed of common

learning. More precisely, for any information structure $\mathcal{I}$, we consider the probability that agents have approximate common knowledge of the true state after t i.i.d. signal draws from $\mathcal{I}$. Theorem 1 shows that, as t grows large, this probability converges to one at an exponential rate given by the learning efficiency index of $\mathcal{I}$. Approximate common knowledge is a much more demanding notion than individual knowledge, as it imposes confidence not only on agents' first-order beliefs about the state, but on their infinite hierarchy of higher-order beliefs. However, the fact that our learning efficiency index does not depend on the correlation across agents' signals has the following important implication: Common learning and individual learning occur at the *same* rate. Thus, with many signal observations, agents' higher-order belief uncertainty vanishes at least as fast as their first-order uncertainty. The proof of Theorem 1 relies on a key lemma that uses the "second law of thermodynamics" for Markov chains to relate players' observations and their higher-order beliefs via Kullback-Leibler divergence (Lemma 1) .

Second, building on Theorem 1, we use the learning efficiency index to provide a large-sample ranking over information structures in games. With any game, we associate an objective function over outcomes in each state, capturing, for instance, agents' welfare or a designer's preferences. Theorems 2–3 show that, for a rich class of games and objectives, information structures with a higher learning efficiency index induce better (Bayes Nash) equilibrium outcomes whenever agents observe sufficiently many signal draws. The only assumption imposed on the game and objective function is that, under common knowledge of the state, the first-best outcome (according to the objective) can be achieved by some strict Nash equilibrium of the game. As this assumption only requires the objective and agents' incentives to be aligned at certainty, it allows for rich strategic externalities. For instance, if the objective is to maximize utilitarian welfare, this assumption captures many important coordination games in the literature, such as the illustrative example below.

Based on the structure of the learning efficiency index, our ranking has implications for the design of information structures in games: In particular, if agents have access to many signal draws, then the way to achieve better equilibrium outcomes is by improving the worst-informed agent's information about the state. In contrast, providing signals about other agents' signals that do not contain additional information about the state is not effective. Thus, whereas a central insight in the literature is that higher-order belief uncertainty can be a significant source of ineffi-

3

ciency in incomplete information games, our results suggest that, when agents have access to large samples of signals, reducing higher-order belief uncertainty becomes a second-order concern.

Illustrative example: Joint investment. Consider two players $i = 1, 2$, with symmetric action sets $A_i = \{-1, 1, 0\}$, where action 1 (resp., -1) represents investing in project 1 (resp., project -1) and 0 represents no investment. The state of the world $\theta \in \{-1, 1\}$ captures which of the two projects will succeed and is drawn according to some non-degenerate prior p_0. Each player i's utility takes the form

$$u_i(a_1, a_2, \theta) = \mathbf{1}_{\{a_1 = a_2 = \theta\}} - c|a_i|;$$

that is, if i chooses to invest in either project, she incurs an investment cost of $c \in (0, 1)$, and she receives a payoff of 1 if and only if she invests in the correct project and her opponent also invests in this project. Under utilitarian welfare, $\frac{1}{2}(u_1(a, \theta) + u_2(a, \theta))$, the efficient outcome is to play (θ, θ) in state θ. This is a strict Nash equilibrium of the game under common knowledge of θ, but is not achievable as a Bayes-Nash equilibrium under incomplete information.

Now suppose that, prior to playing the game, players learn about state θ from repeated i.i.d. signal draws. Our learning efficiency index yields a generically complete ranking over information structures that makes it possible to compare how fast players achieve approximate common knowledge of θ, and hence how close the induced equilibrium play is to the efficient outcome after sufficiently many signal draws. For example, consider a simple class of binary information structures, where each player i's private signal realizations x_i are either -1 or 1, and the joint probabilities of players' signals conditional on state θ are

	$x_1 = \theta$	$x_1 \neq \theta$
$x_2 = \theta$	$\gamma\rho$	$\gamma(1 - \rho)$
$x_2 \neq \theta$	$\gamma(1 - \rho)$	$1 - \gamma(2 - \rho).$

Each information structure is summarized by an individual precision parameter $\gamma \in (1/2, 1)$, capturing the probability with which each player's signal matches the state, and a parameter $\rho \in (2 - \frac{1}{\gamma}, 1)$, capturing the extent of correlation across players' signals. Higher values of γ improve each player's individual learning about state θ, while higher values of ρ facilitate more accurate predictions of the opponent's signals (and hence their beliefs and actions). Thus, in comparing two information

structures parametrized by (γ, ρ) and $(\tilde{\gamma}, \tilde{\rho})$, it might not be clear how to trade off these two considerations. Indeed, if players observe only a small number of signal draws, whether (γ, ρ) or $(\tilde{\gamma}, \tilde{\rho})$ induces better equilibrium play can vary across different priors p_0 and investment costs c.

However, we will show that our learning efficiency index depends only on γ. Thus, for *any* p_0 and c, higher levels of individual precision γ guarantee better equilibrium welfare whenever players observe sufficiently many signal draws; in contrast, the effect of correlation ρ becomes negligible as the number of signals grows large. As we will see, this is due to the fact that the speed of common learning is the same as the speed of individual learning, because uncertainty about opponents' signals vanishes faster than uncertainty about the state.

1.2   Related Literature

Our paper bridges the literatures on higher-order beliefs and learning efficiency. Within the former, we relate most closely to Cripps, Ely, Mailath, and Samuelson (2008), henceforth CEMS. Their main result establishes that, in the current setting (with finite states and signals), every information structure leads to common learning as the number of signal observations goes to infinity.[1] We derive a simple learning efficiency index that characterizes the speed of common learning under each information structure. Characterizing the speed of learning is also essential for our second contribution of comparing how different information structures affect equilibrium outcomes after a large but finite number of signal draws. As we discuss (Remark 2), our proof uses Markov chain arguments that are related to CEMS' approach, but is based on a different construction.

Moscarini and Smith (2002) derive an efficiency index that characterizes the speed of single-agent learning.[2] As we discuss (Remark 1), our multi-agent index can be seen to reduce to theirs in the single-agent case. The main novelty of our analysis is to show that higher-order belief uncertainty vanishes at least as fast as first-order uncertainty; thus, the multi-agent index corresponds to the slowest individual agent's

---

[1]Several papers (e.g., Steiner and Stewart, 2011; Cripps, Ely, Mailath, and Samuelson, 2013) study common learning when signals are correlated across draws. Liang (2019) considers non-Bayesian agents who learn from public signals. Acemoglu, Chernozhukov, and Yildiz (2016) consider a setting that features identification problems due to uncertainty about the information structure.

[2]See also Frick, Iijima, and Ishii (2021) and Fudenberg, Lanzani, and Strack (2021) in the context of misspecified single-agent learning.

5

learning index, while the correlation across agents' signals plays no role.

Learning efficiency has also been analyzed in various social learning environments, but most work has not focused on the role of higher-order beliefs.[3] A notable exception is Harel, Mossel, Strack, and Tamuz (2021), who consider a setting in which long-lived agents repeatedly observe both private signals and other agents' actions, so that higher-order beliefs matter for agents' inferences. They derive an upper bound on the speed of first-order learning that holds uniformly across all population sizes. We study learning from exogenous signals rather than from others' actions, but provide an exact characterization of the convergence speed of both higher-order and first-order beliefs.

Theorems 2–3 relate to the literature on comparisons of information structures. Blackwell's (1951) order compares single-agent information structures in terms of their induced payoffs in all decision problems, assuming that a single signal draw is observed. Moscarini and Smith's (2002) aforementioned efficiency index extends this order to single-agent settings with a large number of i.i.d. signal draws. Mu, Pomatto, Strack, and Tamuz (2021) (see also Azrieli, 2014) consider a more demanding order than Moscarini and Smith (2002), by requiring the number of signal observations to be uniform across decision problems.

Several papers extend the Blackwell order to multi-agent settings, focusing on the single signal draw case. Gossner (2000) compares (Bayes Nash) equilibrium outcomes for general games and objective functions. He shows that this yields a very conservative order, where no two information structures that induce different (higher-order) beliefs can be compared. Thus, one needs to restrict the class of games and objectives to obtain less degenerate rankings.[4] In particular, Lehrer, Rosenberg, and Shmaya (2010) focus on common interest games with utilitarian welfare, and characterize the order based on a generalization of Blackwell's garbling condition. Analogously, Pęski (2008) compares min-max values in zero-sum games. Our exercise is most comparable to Lehrer, Rosenberg, and Shmaya (2010), in that we also impose a form of alignment on agents' incentives and the objective function. However, by assuming that agents observe many signal draws, we obtain a ranking that is a completion of

---

[3]See, e.g., Vives (1993); Duffie and Manso (2007); Hann-Caruthers, Martynov, and Tamuz (2018); Rosenberg and Vieille (2019); Liang and Mu (2020); Dasaratha and He (2019).

[4]Bergemann and Morris (2016) study general games by using a different approach. They consider Bayes correlated equilibria, which are equivalent to Bayes Nash equilibria in a setting with a mediator who commits to sending action recommendations after observing the state and signals.

Lehrer, Rosenberg, and Shmaya's (2010) order and that applies to a richer class of games and objective functions beyond the common-interest case.

2 Setting

Throughout the paper, we fix a finite set of agents I, a finite set of states Θ, and a full-support (common) prior belief $p_0 \in \Delta(\Theta)$.

An ***information structure*** $\mathcal{I}$ consists of a finite set of signals X_i for each agent i, with corresponding set of signal profiles $X := \prod_{i \in I} X_i$, as well as a distribution $\mu^\theta \in \Delta(X)$ over signal profiles conditional on each state $\theta \in \Theta$. Let $\mu_i^\theta \in \Delta(X_i)$ denote the marginal distribution over agent i's signals in state θ. We assume either that each distribution μ^θ has full support over X, or that signals are perfectly correlated with full-support marginals.[5] We also assume that $\mu_i^\theta \neq \mu_i^{\theta'}$ for all $i \in I$ and distinct $\theta, \theta' \in \Theta$.

A ***basic game*** $\mathcal{G}$ consists of a finite set of actions A_i for each agent i, with corresponding set of action profiles $A := \prod_{i \in I} A_i$, as well as a utility function $u_i : A \times \Theta \to \mathbb{R}$ over action profiles and states for each agent i.

We consider settings where prior to playing a basic game $\mathcal{G}$, agents observe repeated i.i.d. signal draws from an information structure $\mathcal{I}$. Formally, for each information structure $\mathcal{I}$ and $t \in \mathbb{N}$, let $\mathbb{P}_t^{\mathcal{I}} \in \Delta(\Theta \times X^t)$ denote the probability distribution over states and sequences of signal profiles that results when the state θ is drawn according to prior p_0 and, conditional on each state θ, a sequence $x^t = (x_\tau)_{\tau=1,\dots,t}$ of signal profiles is generated according to t independent draws from μ^θ. For each basic game $\mathcal{G}$, we consider the ***incomplete information game*** $\mathcal{G}_t(\mathcal{I})$, where states and signal sequences are drawn according to $\mathbb{P}_t^{\mathcal{I}}$ and a strategy for agent i is a map $\sigma_{it} : (X_i)^t \to \Delta(A_i)$ from i's observed signal sequences $x_i^t = (x_{i\tau})_{\tau=1,\dots,t}$ to (one-shot) mixed actions in A_i.

Let $\mathrm{BNE}_t(\mathcal{G}, \mathcal{I})$ denote the set of Bayes Nash equilibria (BNE) of $\mathcal{G}_t(\mathcal{I})$. That is, a strategy profile $\sigma_t = (\sigma_{it})_{i \in I}$ is in $\mathrm{BNE}_t(\mathcal{G}, \mathcal{I})$ if for each $i \in I$, $x_i^t \in X_i^t$, and a_i with

---

[5]Formally, signals are perfectly correlated if $X_i = X_j$ for all i, j, and for each $x \in X$ and θ,
$$\mu^\theta(x) = \begin{cases} \mu_i^\theta(x_i) \text{ if } x_i = x_j \text{ for all } i, j, \\ 0 \text{ otherwise} \end{cases} .$$

$$\sigma_{it}(a_i \mid x_i^t) > 0,$$

$$a_i \in \operatorname*{argmax}_{a_i' \in A_i} \sum_{\theta \in \Theta, x_{-i}^t \in X_{-i}^t} \mathbb{P}_t^{\mathcal{I}}(\theta, x_{-i}^t \mid x_i^t) \sum_{a_{-i} \in A_{-i}} \sigma_{-i}(a_{-i} \mid x_{-i}^t) u_i(a_i', a_{-i}, \theta).$$

3 Multi-Agent Learning Efficiency

3.1 Common Learning

Our first goal is to characterize the learning efficiency of each information structure $\mathcal{I}$. To formalize learning, we employ CEMS' notion of *common learning.* This captures that, in multi-agent settings, learning not only concerns agents' beliefs about the state θ, but also their higher-order uncertainty about other agents' beliefs.

Fix an information structure $\mathcal{I}$. For any $t \in \mathbb{N}$, $p \in (0, 1)$, and event $E \subseteq \Theta \times X^t$, let $B_t^p(E)$ denote the event that E is p-**believed** at t, i.e., that all agents assign probability at least p to E after t draws from $\mathcal{I}$. Formally,

$$B_t^p(E) := \bigcap_{i \in I} B_{it}^p(E), \quad \text{where} \quad B_{it}^p(E) := \Theta \times \{x_i^t \in X_i^t : \mathbb{P}_t^{\mathcal{I}}(E \mid x_i^t) \geq p\} \times \prod_{j \neq i} X_j^t.$$

Since $\mu_i^\theta \neq \mu_i^{\theta'}$ for all i and $\theta \neq \theta'$, standard arguments imply that all players ***individually learn*** the true state; that is, for all $p \in (0, 1)$ and $\theta \in \Theta$, we have

$$\lim_{t \to \infty} \mathbb{P}_t^{\mathcal{I}}\left(B_t^p(\theta) \mid \theta\right) = 1,$$

where, slightly abusing notation, we also use θ to denote the event $\{\theta\} \times X^t$.

While individual learning only requires all agents' first-order beliefs to eventually assign probability arbitrarily close to 1 to the true state, CEMS' notion of common learning additionally considers agents' higher-order beliefs. Let

$$C_t^p(E) := \bigcap_{k \in \mathbb{N}} (B_t^p)^k(E)$$

denote the event that E is ***commonly p-believed*** at t. Thus, at $C_t^p(E)$, the event E is p-believed, the event $B_t^p(E)$ is p-believed, and so on. ***Common learning*** obtains if the true state is eventually commonly p-believed for p arbitrarily close to 1; that

is, for all $p \in (0,1)$ and $\theta \in \Theta$,

$$\lim_{t \to \infty} \mathbb{P}_t^{\mathcal{I}} \left(C_t^p(\theta) \mid \theta \right) = 1. \tag{1}$$

The event that $C_t^p(\theta)$ for p close to 1 captures that players have approximate common knowledge of state θ. Conditional on this event, for any basic game $\mathcal{G}$, equilibria in $\mathrm{BNE}_t(\mathcal{G}, \mathcal{I})$ approximate equilibria of $\mathcal{G}$ under common knowledge of θ (Monderer and Samet, 1989).

The main result in CEMS is that when states and signals are finite, as in our setting, then common learning always obtains:[6]

Proposition 0 (CEMS). *For any information structure $\mathcal{I}$, common learning obtains.*

3.2 Characterization of Learning Efficiency

Proposition 0 shows that all information structures eventually lead to approximate common knowledge of the state. However, it says nothing about the *rate* at which the convergence in (1) is achieved, and hence about how each $\mathcal{I}$ affects equilibrium play in any game $\mathcal{G}_t(\mathcal{I})$. To measure the learning *efficiency* of each information structure, we derive a simple index that characterizes this rate for each $\mathcal{I}$.

First, define the ***Chernoff distance*** between any two distributions $\mu, \mu' \in \Delta(Y)$ over a finite set Y by

$$d(\mu, \mu') := \min_{\nu \in \Delta(Y)} \max\{\mathrm{KL}(\nu, \mu), \mathrm{KL}(\nu, \mu')\}. \tag{2}$$

Here, $\mathrm{KL}(\nu, \mu) := \sum_{y \in Y} \mu(y) \log \frac{\mu(y)}{\nu(y)}$ denotes the Kullback-Leibler (henceforth, KL) divergence of ν relative to μ.[7] Smaller values of $\mathrm{KL}(\nu, \mu)$ quantify that an empirical distribution ν is better explained by the theoretical distribution μ, in the sense that (a large number of) repeated i.i.d. draws from μ are more likely to generate empirical distributions ν with a smaller KL-divergence relative to μ. The Chernoff distance is a common statistical measure of the dissimilarity of distributions μ and μ' (e.g., Cover and Thomas, 1999). To understand this definition, observe that any minimizer ν of (2) must satisfy $\mathrm{KL}(\nu, \mu) = \mathrm{KL}(\nu, \mu')$, so $d(\mu, \mu')$ is the distance from μ and μ' to

---

[6]CEMS provide an example with countably infinite Θ, in which individual learning holds but common learning fails.

[7]We use the convention that $0 \log 0 = \frac{0}{0} = 0$ and $\log \frac{1}{0} = \infty$.

their KL-midpoint. Thus, the smaller $d(\mu, \mu')$, the more difficult it is to distinguish μ and μ', because repeated draws from either distribution are more likely to generate an empirical distribution ν that is explained equally well by μ and μ'. Note that $d(\mu, \mu') > 0$ whenever $\mu \neq \mu'$.

Using the Chernoff distance, we introduce the following learning efficiency index:

Definition 1. *For any information structure $\mathcal{I}$, define the **learning efficiency index in state θ** by*

$$\lambda^\theta(\mathcal{I}) := \min_{i \in I, \theta' \neq \theta} d(\mu_i^\theta, \mu_i^{\theta'}). \tag{3}$$

In each state θ, Definition 1 reduces each information structure $\mathcal{I}$ to a simple one-dimensional measure. For each agent i, the Chernoff distance $d(\mu_i^\theta, \mu_i^{\theta'})$ between i's marginal signal distribution in state θ and in any other state θ' captures how difficult i finds it to distinguish θ' from θ. The index $\lambda^\theta(\mathcal{I})$ is the minimum of $d(\mu_i^\theta, \mu_i^{\theta'})$ across all agents i and states $\theta' \neq \theta$. Thus, it focuses only on the worst-informed agent i and the state θ' that i finds most difficult to distinguish from the true state θ.

Notably, since the learning efficiency indices depend only on agents' marginal signal distributions, the correlation across different agents' signals plays no role. For instance, in the illustrative example (Section 1.1), where $\mathcal{I}$ is summarized by an individual precision parameter γ and a correlation parameter ρ, $\lambda^\theta(\mathcal{I})$ is strictly increasing in γ but does not depend on ρ. More generally, if each agent i's marginal signal distributions under $\mathcal{I}$ Blackwell-dominate those under $\tilde{\mathcal{I}}$, then $\lambda(\mathcal{I}) \geq \lambda(\tilde{\mathcal{I}})$.

Our first main result is that $\lambda^\theta(\mathcal{I})$ captures the rate of common learning under information structure $\mathcal{I}$. Moreover, this coincides with the rate of individual learning:

Theorem 1. *Fix any information structure $\mathcal{I}$, $\theta \in \Theta$, and $p \in (0, 1)$. As $t \to \infty$,*

$$\mathbb{P}_t^{\mathcal{I}}\left(B_t^p(\theta) \mid \theta\right) = 1 - \exp[-\lambda^\theta(\mathcal{I})t + o(t)]; \tag{4}$$

$$\mathbb{P}_t^{\mathcal{I}}\left(C_t^p(\theta) \mid \theta\right) = 1 - \exp[-\lambda^\theta(\mathcal{I})t + o(t)]. \tag{5}$$

As highlighted by a rich literature (going back to, e.g., Rubinstein, 1989), common p-belief is a much more demanding requirement than individual p-belief: $C_t^p(\theta)$ imposes confidence not only on agents' first-order beliefs about the state, but on their entire infinite hierarchy of higher-order beliefs. Based on this, it might be natural to expect common learning to occur more slowly than individual learning. However,

Theorem 1 shows that, as $t \to \infty$, the probability of common p-belief and the probability of individual p-belief of the true state θ both tend to 1 at the *same* exponential rate, which is given by the learning efficiency index $\lambda^\theta(\mathcal{I})$.

That is, when agents observe large samples of signals, higher-order belief uncertainty vanishes just as fast as first-order uncertainty. This point is also reflected by the fact that the learning efficiency index $\lambda^\theta(\mathcal{I}) = \min_{i \in I, \theta' \neq \theta} d(\mu_i^\theta, \mu_i^{\theta'})$ depends only on the worst-informed agent's marginal signal distributions, while correlation across agents' signals plays no role. When players observe a small sample of signals, increasing individual signal precision and increasing correlation of signals can both improve the probability of common p-belief of the correct state. However, under large samples, the effect of correlation becomes second-order.

The proof of Theorem 1 is in Appendix B. We sketch the argument in the next section. The key step is a lemma relating higher-order beliefs to KL-divergence (Lemma 1), which we use to show that higher-order uncertainty vanishes at least as fast as first-order uncertainty.

Remark 1 (Single-agent learning efficiency)**.** Applied to the single-agent case, $I = \{i\}$, Theorem 1 yields that each agent i's individual rate of learning (i.e., the rate at which $\mathbb{P}_t^{\mathcal{I}}\left(B_{it}^p(\theta) \mid \theta\right) \to 1$) is given by $\lambda_i^\theta(\mathcal{I}) := \min_{\theta' \in \Theta \setminus \{\theta\}} d(\mu_i^\theta, \mu_i^{\theta'})$. This is equivalent to the single-agent learning efficiency index introduced by Moscarini and Smith (2002), which is based on the Hellinger transform:

$$\lambda_{i,\mathrm{MS}}^\theta(\mathcal{I}) = \min_{\theta' \in \Theta \setminus \{\theta\}} \max_{\kappa \in [0,1]} - \log \sum_{x_i \in X_i} \mu_i^\theta(x_i)^\kappa \mu_i^{\theta'}(x_i)^{1-\kappa}. \tag{6}$$

Indeed, the variational formula (e.g., Dupuis and Ellis, 2011, Lemma 6.2.3.f) ensures that $d(\mu_i^\theta, \mu_i^{\theta'}) = \max_{\kappa \in [0,1]} - \log \sum_{x_i \in X_i} \mu_i^\theta(x_i)^\kappa \mu_i^{\theta'}(x_i)^{1-\kappa}$ for any distinct θ, θ'.

Thus, our efficiency index can be viewed as a multi-player generalization of Moscarini and Smith (2002), and Theorem 1 shows that the rate of common learning, $\lambda^\theta(\mathcal{I}) = \min_{i \in I} \lambda_i^\theta(\mathcal{I})$, corresponds to the slowest agent's rate of individual learning. ▲

3.3   Proof Sketch of Theorem 1

Speed of individual learning. We first show that each agent i's rate of individual

learning in state θ is $\lambda_i^\theta(\mathcal{I}) = \min_{\theta' \neq \theta} d(\mu_i^\theta, \mu_i^{\theta'})$, i.e., as $t \to \infty$,

$$\mathbb{P}_t^{\mathcal{I}}(B_{it}^p(\theta) \mid \theta) = 1 - \exp[-\lambda_i^\theta(\mathcal{I})t + o(t)]. \tag{7}$$

This can be seen by showing that $\lambda_i^\theta(\mathcal{I})$ is equivalent to Moscarini and Smith's (2002) single-agent efficiency index (see Remark 1). However, for clarity, we sketch a direct proof on which we will build below to characterize the speed of common learning.

Let $\nu_{it} \in \Delta(X_i)$ denote the empirical distribution of i's signals up to t, which is a sufficient statistic for i's beliefs. By standard arguments, as $t \to \infty$, i's beliefs concentrate on the state whose signal distribution minimizes KL-divergence relative to ν_{it}. Thus, for any $\varepsilon > 0$, we have that for all large enough t,

$$\left\{ \mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq \min_{\theta' \neq \theta} \mathrm{KL}(\nu_{it}, \mu_i^{\theta'}) - \varepsilon \right\} \subseteq B_{it}^p(\theta) \subseteq \left\{ \mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq \min_{\theta' \neq \theta} \mathrm{KL}(\nu_{it}, \mu_i^{\theta'}) + \varepsilon \right\}. \tag{8}$$

Moreover, by Sanov's theorem from large deviation theory, for any closed $D_i \subseteq \Delta(X_i)$ with non-empty interior,

$$\mathbb{P}_t^{\mathcal{I}}(\nu_{it} \in D_i \mid \theta) = 1 - \exp[-\inf_{\nu \notin D_i} \mathrm{KL}(\nu, \mu_i^\theta)t + o(t)], \quad \text{as } t \to \infty.$$

For $D_i := \{\nu_i \in \Delta(X_i) : \mathrm{KL}(\nu_i, \mu_i^\theta) \leq \min_{\theta' \neq \theta} \mathrm{KL}(\nu_i, \mu_i^{\theta'})\}$, it can be seen that $\inf_{\nu_i \notin D_i} \mathrm{KL}(\nu_i, \mu_i^\theta) = \lambda_i^\theta(\mathcal{I}).$[8] Thus, applying Sanov's theorem to the upper and lower bounds in (8) and letting $\varepsilon \to 0$ yields the desired conclusion.

Finally, (7) implies (4): Since $B_t^p(\theta) = \bigcap_{i \in I} B_{it}^p(\theta)$, the speed of convergence of $\mathbb{P}_t^{\mathcal{I}}(B_t^p(\theta) \mid \theta)$ is governed by the slowest individual learning rate, $\lambda^\theta(\mathcal{I}) = \min_{i \in I} \lambda_i^\theta(\mathcal{I})$.

Speed of common learning. Since $C_t^p(\theta) \subseteq B_t^p(\theta)$, the speed of common learning at θ cannot exceed the speed of individual learning and thus, by the first part, is at most $\lambda^\theta(\mathcal{I})$. The main step of the proof establishes that the speed of

---

[8]Indeed, note that

$$\inf_{\nu_i \notin D_i} \mathrm{KL}(\nu_i, \mu_i^\theta) = \inf \left\{ \mathrm{KL}(\nu_i, \mu_i^\theta) : \mathrm{KL}(\nu_i, \mu_i^\theta) > \mathrm{KL}(\nu_i, \mu_i^{\theta'}) \text{ for some } \theta' \neq \theta \right\}$$

$$= \min_{\theta' \neq \theta} \left\{ \mathrm{KL}(\nu_i, \mu_i^\theta) : \mathrm{KL}(\nu_i, \mu_i^\theta) = \mathrm{KL}(\nu_i, \mu_i^{\theta'}) \right\} = \min_{\theta' \neq \theta} d(\mu_i^\theta, \mu_i^{\theta'}) = \lambda_i^\theta(\mathcal{I}).$$

common learning at θ is at least $\lambda^\theta(\mathcal{I})$, i.e., as $t \to \infty$,

$$\mathbb{P}_t^\mathcal{I}(C_t^p(\theta) \mid \theta) \geq 1 - \exp[-\lambda^\theta(\mathcal{I})t + o(t)].$$

To see this, fix any $d < \lambda^\theta(\mathcal{I})$. For each t, we consider the event

$$F_t(\theta, d) := \bigcap_{i \in I} F_{it}(\theta, d), \quad \text{where } F_{it}(\theta, d) := \left\{ \mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d \right\}.$$

Observe that $d < \lambda^\theta(\mathcal{I})$ together with (8) implies that, for all large enough t,

$$F_t(\theta, d) \subseteq B_t^p(\theta). \tag{9}$$

We now show more strongly that, for all large enough t,

$$F_t(\theta, d) \subseteq C_t^p(\theta). \tag{10}$$

Given this, Sanov's theorem implies that

$$\mathbb{P}_t^\mathcal{I}(C_t^p(\theta) \mid \theta) \geq \mathbb{P}_t^\mathcal{I}(F_t(\theta, d) \mid \theta) = 1 - \exp[-dt + o(t)], \quad \text{as } t \to \infty.$$

This yields the desired conclusion since d can be chosen arbitrarily close to $\lambda^\theta(\mathcal{I})$.

By Monderer and Samet (1989), to show (10) it is enough to prove that event $F_t(\theta, d)$ is *p-evident*. That is, we want to show that, for all large enough t,

$$F_t(\theta, d) \subseteq B_t^p(F_t(\theta, d)). \tag{11}$$

For this, we establish the following key lemma that uses KL-divergence to relate i's own observations ν_{it} to i's beliefs about others' observations. For any two agents i and j, let $M_{ij}^\theta \in \mathbb{R}^{X_i \times X_j}$ denote the matrix whose (x_i, x_j)-th element is $M_{ij}^\theta(x_i, x_j) = \mu^\theta(x_j \mid x_i)$. As CEMS observed, if agent i's empirical signal distribution at t is ν_{it}, then conditional on state θ, i's expectation of j's empirical distribution is given by $\nu_{it} M_{ij}^\theta \in \Delta(X_j)$. Moreover, $\mu_i^\theta M_{ij}^\theta = \mu_j^\theta$.

Lemma 1. *For each $\theta \in \Theta$, distinct $i, j \in I$, and $\nu_i \in \Delta(X_i)$, we have $\mathrm{KL}(\nu_i, \mu_i^\theta) \geq \mathrm{KL}(\nu_i M_{ij}^\theta, \mu_j^\theta)$. Moreover, the inequality is strict whenever $\nu_i \neq \mu_i^\theta$ and signals are not perfectly correlated.*

13

To understand the result, suppose that i observes an empirical signal distribution ν_i. Then $\mathrm{KL}(\nu_i, \mu_i^\theta)$ quantifies how much i's observations deviate from i's average signal distribution μ_i^θ in state θ. Likewise, $\mathrm{KL}(\nu_i M_{ij}^\theta, \mu_j^\theta)$ quantifies how much i's *expectation* of j's observations deviates from j's average signal distribution μ_j^θ. Thus, Lemma 1 says that when i forms an estimate of j's signal observations based on i's signal observations, this estimate cannot be any more "atypical" than i's own signal observations. For example, if i and j's signals are independent, then regardless of her own observations, i's estimate of j's observations is always the average distribution (i.e., $\mathrm{KL}(\nu_i M_{ij}^\theta, \mu_j^\theta) = 0$). At the opposite extreme, if i and j's signals are perfectly correlated, then i expects j to observe the same signals as herself, so her estimate of j's observations is just as atypical as her own observations (i.e., $\mathrm{KL}(\nu_i M_{ij}^\theta, \mu_j^\theta) = \mathrm{KL}(\nu_i, \mu_i^\theta)$).

Finally, to see how Lemma 1 implies (11), consider agent i's reasoning conditional on the event that $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d$ when t is large enough. By (9), i assigns high probability to state θ. Hence, by a law of large numbers argument, i assigns high probability to every other agent j's realized empirical distribution ν_{jt} being close to i's expectation $\nu_{it} M_{ij}^\theta$ conditional on state θ. But then, Lemma 1 together with the fact that $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d$ implies that i also assigns high probability to the event that $\mathrm{KL}(\nu_{jt}, \mu_j^\theta) \leq d$ for all j. Thus, $F_t(\theta, d)$ is p-evident at all large t. $\qquad\square$

Remark 2. *Second law of thermodynamics.* The weak inequality in Lemma 1 is an implication of the "second law of thermodynamics" for Markov chains, which says that the KL-divergence between any two initial distributions shrinks under iterated application of the transition matrix (see, e.g., Section 4.4 in Cover and Thomas, 1999). Indeed, consider the Markov chain defined on the state space $X_i \cup X_j$, whose transition matrix is given by M_{ij}^θ if the current state is in X_i and by M_{ji}^θ if the current state is in X_j. Then the second law applied to the initial distributions ν_i and μ_i^θ implies that $\mathrm{KL}(\nu_i, \mu_i^\theta) \geq \mathrm{KL}(\nu_i M_{ij}^\theta, \mu_i^\theta M_{ij}^\theta)$, which yields the desired inequality as $\mu_i^\theta M_{ij}^\theta = \mu_j^\theta$.

Relationship with CEMS. In proving Proposition 0, CEMS consider a different Markov chain, which is defined on the space X_i and has transition matrix $M_{ij}^\theta M_{ji}^\theta$. They show that this transition matrix is a sup-norm contraction on $\Delta(X_i)$ (see their Lemma 4), and based on this construct a different sequence of p-evident events F_t (that are defined using the sup-norm rather than KL-divergence). While the probability of these events F_t also converges to 1, the rate of convergence is less than $\lambda^\theta(\mathcal{I})$. Thus, this construction cannot be used to provide a tight bound on the speed

of common learning.

Convergence of belief hierarchies. Theorem 1 characterizes the speed at which players achieve approximate common knowledge in the sense of common p-belief of the true state. Analogous results hold if proximity to common knowledge is instead formalized in terms of commonly used topologies over belief hierarchies. See Appendix ?? for details. ▲

4 Ranking Information Structures in Games

So far, we have analyzed learning efficiency under each information structure $\mathcal{I}$. We now return to the setting where, following t draws of signals from $\mathcal{I}$, agents play a game. We show that, when t is sufficiently large, the learning efficiency index can be used to rank information structures in terms of their equilibrium outcomes: Information structures with a higher learning efficiency index induce better equilibrium outcomes, robustly for a rich class of games and objective functions.

4.1 Objective Functions

Given any basic game $\mathcal{G}$, we introduce an **objective function** $W : A \times \Theta \to \mathbb{R}$, which assigns a value to each action profile and state. We assume that in each state θ, W is maximized by a unique action profile, $\{a^{\theta,W}\} = \mathrm{argmax}_{a \in A} W(a, \theta)$. The objective function can be interpreted as capturing a designer's preferences over outcomes in the game. A benevolent designer might seek to maximize agents' welfare, for example, via utilitarian aggregation, $W = \frac{1}{I} \sum_{i \in I} u_i$. However, we also allow for objective functions that do not relate to agents' utilities in any particular way.

For any information structure $\mathcal{I}$ and number t of signal draws, we use W to evaluate expected equilibrium outcomes in the incomplete-information game $\mathcal{G}_t(\mathcal{I})$. Specifically, for any strategy profile $\sigma_t = (\sigma_{it})_{i \in I}$ in game $\mathcal{G}_t(\mathcal{I})$, let

$$W_t(\sigma_t, \mathcal{I}) := \sum_{\theta \in \Theta, x^t \in X^t} \mathbb{P}^{\mathcal{I}}(\theta, x^t) \sum_{a \in A} \sigma_t(a \mid x^t) W(a, \theta)$$

denote the ex-ante expected value of the objective when signal sequences x^t are drawn from information structure $\mathcal{I}$ in each state and each agent then chooses their action

according to strategy $\sigma_{it}(\cdot \mid x_i^t)$. Define the **objective value**

$$W_t(\mathcal{G}, \mathcal{I}) := \sup_{\sigma_t \in \text{BNE}_t(\mathcal{G}, \mathcal{I})} W_t(\sigma_t, \mathcal{I}) \tag{12}$$

to be the ex-ante expected value of the objective under the best BNE of $\mathcal{G}_t(\mathcal{I})$.

For any two information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$, we seek to compare their objective values $W_t(\mathcal{G}, \mathcal{I})$ and $W_t(\mathcal{G}, \tilde{\mathcal{I}})$ when the number t of signal draws is large. We will see that this comparison can be carried out robustly for a rich class of games $\mathcal{G}$ and objective functions W.

The one substantive restriction we impose is the following joint assumption on $\mathcal{G}$ and W. Let $\text{SNE}(\mathcal{G}, \theta) \subseteq A$ denote the set of strict Nash equilibria of $\mathcal{G}$ under common knowledge of θ.

Assumption 1 (Alignment at certainty). *For each $\theta \in \Theta$, $a^{\theta, W} \in \text{SNE}(\mathcal{G}, \theta)$.*

Assumption 1 requires that under common knowledge of each state θ, the W-first best outcome $a^{\theta, W}$ is achievable as a strict Nash equilibrium of $\mathcal{G}$. Note that the condition does not require $a^{\theta, W}$ to be the only strict Nash of $\mathcal{G}$ at θ.

One simple environment that satisfies Assumption 1 is when $\mathcal{G}$ is a common-interest game and W represents utilitarian welfare, i.e., $u_i = u_j = W$ for all i, j.[9] In this case, agents' incentives in $\mathcal{G}$ are *fully* aligned with W: Indeed, for any $\mathcal{I}$ and t, *any* strategy profile σ_t that maximizes the expected objective $W_t(\sigma_t, \mathcal{I})$ is a BNE of $\mathcal{G}_t(\mathcal{I})$.

However, Assumption 1 is substantially more permissive than imposing full alignment on $\mathcal{G}$ and W: We only require maximization of W to be an equilibrium of $\mathcal{G}$ under *certainty*, i.e., when players have common knowledge of the state. Except for this requirement, there is no restriction on players' incentives in game $\mathcal{G}$ or the relationship with W. Beyond common interest games, this allows for rich patterns of strategic externalities. For instance, under utilitarian welfare, Assumption 1 is satisfied by many important coordination games (e.g., bank runs, currency attack games, etc.): For example, in the joint investment game in Section 1.1, the efficient action profile at θ is (θ, θ), which is a strict Nash equilibrium under common knowledge of θ (another strict Nash is $(0,0)$). At the same time, Assumption 1 rules out set-

---

[9] Generically, any common interest game $\mathcal{G}$ admits a strict Nash equilibrium that uniquely maximizes utilitarian welfare.

tings where agents' incentives and the objective are misaligned even under complete information (e.g., a prisoner's dilemma game under utilitarian welfare).

Finally, an objective function W might also serve to quantify how close play after t signal draws comes to any particular common knowledge equilibrium. Indeed, given any basic game $\mathcal{G}$ and any selection $a^\theta \in \mathrm{SNE}(\mathcal{G}, \theta)$ of a common knowledge equilibrium at each state θ, define W by

$$W(a, \theta) = \begin{cases} 1 \text{ if } a = a^\theta \\ 0 \text{ otherwise.} \end{cases}$$

Then $\mathcal{G}$ and W trivially satisfy Assumption 1. In this case, the objective value $W_t(\mathcal{G}, \mathcal{I})$ measures the ex-ante probability that, after t draws of signals from $\mathcal{I}$, agents are able to play the common knowledge equilibrium a^θ in each state θ.

4.2 Ranking under Full Separation

Under Assumption 1, we now proceed to rank information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$ in terms of their objective values $W_t(\mathcal{I}, \mathcal{G})$ and $W_t(\tilde{\mathcal{I}}, \mathcal{G})$ at large t. In this section, we additionally assume that all agents must distinguish all states in order to maximize W, as is the case, for instance, in the joint investment game in Section 1.1 (where $a_i^{\theta, W} = \theta$ for all i, θ):

Assumption 2 (Full separation). *For all $i \in I$ and distinct $\theta, \theta' \in \Theta$, $a_i^{\theta, W} \neq a_i^{\theta', W}$.*

Define the ***(ex-ante) learning efficiency index*** by

$$\lambda(\mathcal{I}) := \min_{\theta \in \Theta} \lambda^\theta(\mathcal{I}) = \min_{i \in I, \theta, \theta' \in \Theta, \theta' \neq \theta} d(\mu_i^\theta, \mu_i^{\theta'}). \tag{13}$$

That is, $\lambda(\mathcal{I})$ considers the worst-case across all states of the conditional learning efficiency indices $\lambda^\theta(\mathcal{I})$.

Theorem 2. *Take any information structures $\mathcal{I}, \tilde{\mathcal{I}}$ with $\lambda(\mathcal{I}) \neq \lambda(\tilde{\mathcal{I}})$. The following are equivalent:*

1. *$\lambda(\mathcal{I}) > \lambda(\tilde{\mathcal{I}})$.*

2. *For every basic game $\mathcal{G}$ and objective function W satisfying Assumptions 1–2, there exists T such that $W_t(\mathcal{G}, \mathcal{I}) > W_t(\mathcal{G}, \tilde{\mathcal{I}})$ for all $t > T$.*

17

Theorem 2 shows that for all games $\mathcal{G}$ and objectives W satisfying Assumptions 1–2, the learning efficiency index eventually permits a generically complete ranking over information structures: Except when the efficiency indices $\lambda(\mathcal{I})$ and $\lambda(\tilde{\mathcal{I}})$ are exactly tied, $\mathcal{I}$ and $\tilde{\mathcal{I}}$ can be ranked, and the information structure with the higher efficiency index strictly outperforms that with the lower index whenever agents observe sufficiently many signals.

The structure of $\lambda(\mathcal{I})$ suggests some general implications for the design of information structures in games. In particular, recall that $\lambda(\mathcal{I})$ depends only on the worst-informed agent's marginal signal distributions, while the correlation across agents' signals is irrelevant. Thus, under Assumptions 1–2, Theorem 2 implies that, if agents have access to many signal draws, then a designer should be "egalitarian" and focus on improving the worst-informed agent's information about the state. On the other hand, providing signals about other agents' signals that do not convey any additional information about the state is not effective under large samples. This contrasts with the central insight (e.g., Rubinstein, 1989; Carlsson and Van Damme, 1993; Weinstein and Yildiz, 2007) that (even small amounts of) uncertainty about other agents' signals can be a significant source of inefficiency in incomplete information games (including environments satisfying Assumptions 1–2). The reason for this difference is that, as captured by Theorem 1, higher-order belief uncertainty vanishes at least as fast as first-order uncertainty as $t \to \infty$. Thus, when agents have access to sufficiently many signal draws, interventions that reduce uncertainty about other agents' signals have a negligible effect relative to ones that directly improve agents' information about the state.

Theorem 2 can also be contrasted with Lehrer, Rosenberg, and Shmaya (2010). They consider the case in which agents observe a single signal draw from each information structure and show that a generalization of Blackwell's single-agent garbling condition characterizes when $W_1(\mathcal{G}, \mathcal{I})$ exceeds $W_1(\mathcal{G}, \tilde{\mathcal{I}})$ for any common-interest game $\mathcal{G}$ and utilitarian welfare criterion W. When agents observe many signal draws, Theorem 2 yields a ranking that (i) is a completion of Lehrer, Rosenberg, and Shmaya's (2010) order, and (ii) applies to a richer class of environments that allows for misalignment between agents' incentives and the objective under incomplete information.[10]

---

[10] The former can be seen by noting that when $\mathcal{I} \succsim \tilde{\mathcal{I}}$ in the sense of Lehrer, Rosenberg, and Shmaya (2010), then each agent i's marginal signal distributions under $\mathcal{I}$ Blackwell-dominate those under $\tilde{\mathcal{I}}$, which implies that $\lambda(\mathcal{I}) \geq \lambda(\tilde{\mathcal{I}})$.

Both (i) and (ii) rely on the assumption that agents observe sufficiently many signal draws: When $t = 1$, then even if information structure $\mathcal{I}$ yields higher utilitarian welfare than $\tilde{\mathcal{I}}$ in common-interest games, $\mathcal{I}$ can be strictly worse than $\tilde{\mathcal{I}}$ in some other environments that satisfy Assumption 1–2 but feature misaligned incentives. Moreover, when $t = 1$, many information structures are incomparable even when restricting attention to common-interest games.

To illustrate the proof of Theorem 2 (Appendix C), suppose $\mathcal{G}$ and W satisfy Assumptions 1–2. We show that, for any sequence of equilibria $\sigma_t \in \text{BNE}_t(\mathcal{G}, \mathcal{I})$,

$$\sum_{\theta \in \Theta, x^t \in X^t} \mathbb{P}_t^{\mathcal{I}}(\theta, x^t) \sigma_t(a^{\theta, W} \mid x^t) \leq 1 - \exp[-t\lambda(\mathcal{I}) + o(t)] \text{ as } t \to \infty, \qquad (14)$$

and that (14) holds with equality for some BNE sequence (σ_t). That is, under information structure $\mathcal{I}$, $\lambda(\mathcal{I})$ is the maximal rate at which ex-post inefficient behavior (i.e., not choosing $a^{\theta, W}$ at θ) vanishes in some equilibrium. Thus, if $\lambda(\mathcal{I}) > \lambda(\tilde{\mathcal{I}})$, then $W_t(\mathcal{G}, \mathcal{I}) > W_t(\mathcal{G}, \tilde{\mathcal{I}})$ for all large enough t, because $W_t(\mathcal{G}, \mathcal{I})$ approaches the first-best payoff $\sum_\theta p_0(\theta) W(a^{\theta, W}, \theta)$ faster than does $W_t(\mathcal{G}, \tilde{\mathcal{I}})$.

The argument for inequality (14) is purely statistical and does not consider agents' incentives. Indeed, in Lemma C.1, we show that (14) holds for any sequence of *strategy profiles* (σ_t), regardless of whether or not (σ_t) are equilibria. The basic idea is that, for each agent i, the question whether i's action under σ_{it} matches the correct efficient action $a_i^{\theta, W}$ in each state θ can be recast as a randomized hypothesis test. Given this, the Neyman-Pearson lemma implies that no σ_{it} can achieve a lower ex-ante error probability than a "likelihood ratio test," where agent i chooses action $a_i^{\theta, W}$ whenever her empirical signal frequency ν_{it} is best explained by μ_i^θ (i.e., $\text{KL}(\nu_{it}, \mu_i^\theta) < \text{KL}(\nu_{it}, \mu_i^{\theta'})$ for all $\theta' \neq \theta$). By Sanov's theorem and Assumption 2, the error probability of the latter test decays at rate $\min_{\theta \neq \theta'} d(\mu_i^\theta, \mu_i^{\theta'})$ as $t \to \infty$. Taking the minimum over all agents yields (14).

Finally, the existence of a sequence of *equilibria* for which (14) holds with equality follows from the characterization of the speed of common learning in Theorem 1. By Assumption 1, each $a^{\theta, W}$ is a strict Nash equilibrium under common knowledge of θ. Given this, for any sufficiently large $p \in (0, 1)$ and any t, there exists a BNE σ_t^* under which each agent i plays action $a_i^{\theta, W}$ in the event that θ is common p-belief at t.[11]

---

[11] The reason that $a^{\theta, W}$ is required to be a *strict* Nash equilibrium in Assumption 1 is to ensure that it can be played in a BNE even when players only have *approximate* common knowledge of θ.

Thus, conditional on state θ, the probability that $a^{\theta,W}$ is played under σ_t^* is at least $\mathbb{P}_t^{\mathcal{I}}\left(C_t^p(\theta) \mid \theta\right)$. By Theorem 1, the latter probability goes to 1 at rate $\lambda^\theta(\mathcal{I})$ as $t \to \infty$. Thus, the ex-ante probability of efficient play under sequence (σ_t^*) approaches 1 *at least* at rate $\lambda(\mathcal{I})$. Since, by the previous paragraph, the rate of convergence cannot exceed $\lambda(\mathcal{I})$, (14) must hold with equality under (σ_t^*).

Remark 3. *Comparison across different sample sizes.* The same arguments as in Theorem 2 can be used to obtain a ranking of information structures under different sample sizes: Suppose $\lambda(\mathcal{I}) > k\lambda(\tilde{\mathcal{I}})$ for some $k > 0$. Then for any basic game $\mathcal{G}$ and objective W satisfying Assumptions 1–2, there exists T such that $W_t(\mathcal{G},\mathcal{I}) > W_{kt}(\mathcal{G},\tilde{\mathcal{I}})$ for all $t > T$ with $kt \in \mathbb{N}$.

Beyond best-case equilibrium. In defining the objective value $W_t(\mathcal{G},\mathcal{I})$, (12) considered the best-case BNE. If one focuses instead on the worst-case objective value and replaces Assumption 1 with the assumption that each $W(\cdot, \theta)$ is strictly *minimized* by some action profile in $\mathrm{SNE}(\mathcal{G},\theta)$, then Theorem 2 (applied to the objective $-W$) implies that information structures with a higher learning efficiency index induce a *lower* worst-case objective value at all large t, because equilibrium play can approximate the worst-case common knowledge equilibrium faster. Relatedly, in Appendix **??**, we use the learning efficiency index to characterize the speed at which the entire equilibrium set $\mathrm{BNE}_t(\mathcal{G},\mathcal{I})$ approaches the set of common knowledge equilibria in each state.   ▲

4.3 General Ranking

In Theorem 2, the ranking over information structures reduces to comparing their speed of common learning, because Assumption 2 requires all agents to distinguish all states in order to play the efficient action profile. We now drop Assumption 2, so that some players need not distinguish some pairs of states in order to maximize W. We generalize Theorem 2 to this setting by constructing learning efficiency indices that account for the presence of "equivalent" states for some players.

Formally, given any objective function W, define a partition Π_i^W over Θ for each agent i, whose cells are given by

$$\Pi_i^W(\theta) := \{\theta' \in \Theta : a_i^{\theta,W} = a_i^{\theta',W}\} \text{ for each } \theta,$$

and let $\Pi^W := (\Pi_i^W)_{i \in I}$ denote the collection of all agents' partitions. That is, Π_i^W

divides Θ into equivalence classes of states in which the W-optimal action profile features the same action for agent i.

Given any collection of partitions $\Pi = (\Pi_i)_{i \in I}$ over Θ, we define the learning efficiency index

$$\lambda(\mathcal{I}, \Pi) := \min_{i \in I, \theta, \theta' \in \Theta, \theta' \notin \Pi_i(\theta)} d(\mu_i^\theta, \mu_i^{\theta'}).^{12}$$

That is, in identifying the worst-informed agent and hardest to distinguish states, we do not consider all agents and pairs of states as in (13). Instead, for each agent i, we restrict attention to pairs of states at which i's W-optimal actions are different.

Theorem 3. *Fix any collection $\Pi = (\Pi_i)_{i \in I}$ of partitions over Θ. Take any information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$ with $\lambda(\mathcal{I}, \Pi) \neq \lambda(\tilde{\mathcal{I}}, \Pi)$. The following are equivalent:*

1. $\lambda(\mathcal{I}, \Pi) > \lambda(\tilde{\mathcal{I}}, \Pi)$.

2. *For every $(\mathcal{G}, W)$ satisfying Assumption 1 and $\Pi^W = \Pi$, there exists T such that $W_t(\mathcal{I}, \mathcal{G}) > W_t(\tilde{\mathcal{I}}, \mathcal{G})$ for all $t > T$.*

Theorem 3 extends Theorem 2 by dropping Assumption 2. Based on the generalized learning efficiency indices $\lambda(\cdot, \Pi)$, we again obtain a (generically complete) ranking over the equilibrium outcomes induced by different information structures at large enough t: This ranking applies for all games and objective functions that are aligned at certainty and give rise to the same partitions Π of equivalent states.

Theorem 3 also implies the following partial order over information structures that applies in *all* environments $(\mathcal{G}, W)$ satisfying Assumption 1:

Corollary 1. *Take any information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$ such that $\lambda(\mathcal{I}, \Pi) \neq \lambda(\tilde{\mathcal{I}}, \Pi)$ for all non-degenerate collections of partitions Π. The following are equivalent:*

1. $\lambda(\mathcal{I}, \Pi) > \lambda(\tilde{\mathcal{I}}, \Pi)$ *for all non-degenerate Π.*

2. *For every $(\mathcal{G}, W)$ satisfying Assumption 1, there exists T such that $W_t(\mathcal{I}, \mathcal{G}) > W_t(\tilde{\mathcal{I}}, \mathcal{G})$ for all $t > T$.*

The proof of Theorem 3 generalizes the argument in Theorem 2. That is, as in (14), we show that, for any sequence of strategy profiles (σ_t),

$$\sum_{\theta \in \Theta, x^t \in X^t} \mathbb{P}_t^{\mathcal{I}}(\theta, x^t) \sigma_t(a^{\theta, W} \mid x^t) \leq 1 - \exp[-t\lambda(\mathcal{I}, \Pi^W) + o(t)], \tag{15}$$

---

[12]Slightly abusing notation, we set the index to be ∞ when Π is degenerate (i.e., $\Pi_i(\theta) = \Theta$ for all i).

with equality for some BNE sequence (σ_t). Note that in general $\lambda(\mathcal{I}, \Pi^W) \geq \lambda(\mathcal{I})$. Thus, unlike in the full-separation case, to show that (15) holds with equality for some BNE sequence (σ_t), it is not enough to invoke the fact that the speed of common learning in each state is $\lambda^\theta(\mathcal{I})$. Nevertheless, we show based on Lemma 1 that a similar equilibrium construction as in Theorem 2 remains valid.

Remark 4 (Monotone information structures.). Many economic environments involve information structures that satisfy the monotone-likelihood ratio property with respect to some linear order over states and signals. Appendix **??** considers such environments. We show that, in this case, the condition in Corollary 1 (i.e., $\lambda(\mathcal{I}, \Pi) > \lambda(\tilde{\mathcal{I}}, \Pi)$ for all Π) can be relaxed to one that is easier to verify. This exercise can be viewed as an analog of the relaxation of the Blackwell order considered by Lehmann (1988); Persico (2000); Athey and Levin (2018) in settings with a single agent and single signal draw. ▲

5   Concluding Discussions

5.1   Information Structures as Complements vs. Substitutes

So far, we have considered repeated draws from a single information structure $\mathcal{I}$. However, our learning efficiency index can also be used to formalize whether two information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are complements or substitutes, by considering the effect of combining signal observations from $\mathcal{I}$ and $\tilde{\mathcal{I}}$.

Specifically, given two information structures $\mathcal{I} = (X, (\mu^\theta)_{\theta \in \Theta})$ and $\tilde{\mathcal{I}} = (\tilde{X}, (\tilde{\mu}^\theta)_{\theta \in \Theta})$, consider the combined information structure $\mathcal{I} \times \tilde{\mathcal{I}} := (X \times \tilde{X}, (\mu^\theta \times \tilde{\mu}^\theta)_{\theta \in \Theta})$ under which the signal distribution at each state θ is the product of μ^θ and $\tilde{\mu}^\theta$.

Definition 2. *We say that information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are **complements** if $\lambda(\mathcal{I} \times \tilde{\mathcal{I}}) \geq \lambda(\mathcal{I}) + \lambda(\tilde{\mathcal{I}})$ and **substitutes** if $\lambda(\mathcal{I} \times \tilde{\mathcal{I}}) \leq \lambda(\mathcal{I}) + \lambda(\tilde{\mathcal{I}})$.*

To interpret this definition, consider the case in which $\lambda(\mathcal{I}) = \lambda(\tilde{\mathcal{I}})$ and $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are strict complements, i.e., $\lambda(\mathcal{I} \times \tilde{\mathcal{I}}) > \lambda(\mathcal{I}) + \lambda(\tilde{\mathcal{I}}) = 2\lambda(\mathcal{I})$. Then, by Theorem 1, the speed of common learning under the combined information structure $\mathcal{I} \times \tilde{\mathcal{I}}$ is more than twice as fast as the speed of common learning under $\mathcal{I}$ or $\tilde{\mathcal{I}}$ alone.[13]

---

[13]That is, for all $p \in (0, 1)$ and large enough t, the (ex-ante) probability of common p-belief of the true state is strictly greater if agents observe t signal draws from $\mathcal{I} \times \tilde{\mathcal{I}}$ than if agents observe $2t$

Equivalently, Theorem 2 implies that for any basic game $\mathcal{G}$ and objective function W satisfying Assumptions 1–2 and any large enough t,

$$W_t(\mathcal{I} \times \tilde{\mathcal{I}}, \mathcal{G}) > \max\{W_{2t}(\mathcal{I}, \mathcal{G}), W_{2t}(\tilde{\mathcal{I}}, \mathcal{G})\}.$$

That is, holding fixed any (large enough) total number of signal observations, better equilibrium outcomes are achieved if players observe a mix of signals from $\mathcal{I}$ and $\tilde{\mathcal{I}}$ than if they specialize in only $\mathcal{I}$ or $\tilde{\mathcal{I}}$.

The structure of our efficiency index suggests two conflicting channels that determine whether $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are complements or substitutes. On the one hand, a "force for substitutes" is that the Chernoff distance is subadditive, i.e., for all agents i and states θ, θ',

$$d(\mu_i^\theta \times \tilde{\mu}_i^\theta, \mu_i^{\theta'} \times \tilde{\mu}_i^{\theta'}) \leq d(\mu_i^\theta, \mu_i^{\theta'}) + d(\tilde{\mu}_i^\theta, \tilde{\mu}_i^{\theta'}). \tag{16}$$

Intuitively, this captures that combining multiple information sources creates more scope for "confusing" signal realizations that do not allow an agent to distinguish some states. For example, if observed in isolation, a particular sequence of signal realizations from $\mathcal{I}$ might be indicative of state θ and a sequence of signal realizations from $\tilde{\mathcal{I}}$ might be indicative of state θ', but if the two sequences are observed jointly, these two effects might cancel out and render θ and θ' indistinguishable.[14]

On the other hand, the efficiency index is defined by considering the *worst-case* Chernoff distance across all agents and states. When the worst agent or pair of states differ across $\mathcal{I}$ and $\tilde{\mathcal{I}}$ this creates a hedging value to combining $\mathcal{I}$ and $\tilde{\mathcal{I}}$, which acts as a "force for complements." The following example illustrates both possibilities:

Example 1. Suppose states are binary, $\Theta = \{\theta, \theta'\}$.

Suppose first that signals under either $\mathcal{I}$ or $\tilde{\mathcal{I}}$ are perfectly correlated. Then the worst-informed agent is the same across $\mathcal{I}$ and $\tilde{\mathcal{I}}$. Thus, only the first channel is

---

signal draws from $\mathcal{I}$ or $\tilde{\mathcal{I}}$ alone. An analogous result holds for the speed of learning conditional on any state θ if complementarity is defined using the conditional learning efficiency index λ^θ.

[14]Formally, observe that $d(\mu_i^\theta, \mu_i^{\theta'}) = \min_{\nu_i \in \Delta(X_i)} \mathrm{KL}(\nu_i, \mu_i^\theta)$ s.t. $\mathrm{KL}(\nu_i, \mu_i^\theta) = \mathrm{KL}(\nu_i, \mu_i^{\theta'})$. Combined with the fact that KL-divergence is additive across independent distributions, this yields

$$d(\mu_i^\theta \times \tilde{\mu}_i^\theta, \mu_i^{\theta'} \times \tilde{\mu}_i^{\theta'}) = \min_{\nu_i \in \Delta(X_i), \tilde{\nu}_i \in \Delta(\tilde{X}_i)} \mathrm{KL}(\nu_i, \mu_i^\theta) + \mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta)$$

$$\text{s.t. } \mathrm{KL}(\nu_i, \mu_i^\theta) + \mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta) = \mathrm{KL}(\nu_i, \mu_i^{\theta'}) + \mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^{\theta'}).$$

This implies (16), because $\mathrm{KL}(\nu_i, \mu_i^\theta) + \mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta) = \mathrm{KL}(\nu_i, \mu_i^{\theta'}) + \mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^{\theta'})$ is possible even if $\mathrm{KL}(\nu_i, \mu_i^\theta) \neq \mathrm{KL}(\nu_i, \mu_i^{\theta'})$ and $\mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta) \neq \mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^{\theta'})$.

relevant and $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are substitutes. In particular, (under binary states) this is always the case if there is a single agent.

Suppose next that signals are binary, $X_i = \{x_i, x_i'\}$ and each i's signal distributions are symmetric, i.e., $\mu_i^\theta(x_i) = \mu_i^{\theta'}(x_i')$, $\tilde{\mu}_i^\theta(x_i) = \tilde{\mu}_i^{\theta'}(x_i')$. Then (16) holds with equality. Thus, only the second channel is relevant and $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are complements. ▲

Existing work has studied the complementarity/substitutability of information structures in other settings. Börgers, Hernando-Veciana, and Krähmer (2013) formalize notions of complements/substitutes for single-agent information structures with a single signal observation. Under Gaussian priors and signal distributions, Liang and Mu (2020) study a form of complementarity, where combining multiple information structures allows for identification of the state while each information structure alone leads to non-identification. Complementing these papers, our approach applies to multi-agent information structures and is based on the speed of learning.

5.2   Information Design in Games

The analysis in Section 4 has implications for the design of information structures in games. Beyond the general design implications highlighted following Theorem 2, the learning efficiency index can be used to solve constrained design problems where information is relatively "cheap."

Concretely, given any game $\mathcal{G}$ and objective W, consider the optimal choice of an information structure from some set $\mathbb{I}$ subject to a budget constraint:

$$\max_{\mathcal{I} \in \mathbb{I}, t \in \mathbb{N}} W_t(\mathcal{I}, \mathcal{G}) \ \text{ s.t. } \ tc(\mathcal{I}) \leq \kappa.$$

That is, the designer optimally selects both an information structure $\mathcal{I} \in \mathbb{I}$ and the number t of signal draws from $\mathcal{I}$, subject to a marginal cost of $c(\mathcal{I}) > 0$ per draw from $\mathcal{I}$ and an overall budget of $\kappa > 0$.

For any $\mathcal{G}$ and W satisfying Assumptions 1–2 and any finite set $\mathbb{I}$, our analysis implies that whenever κ is sufficiently large (i.e., information is sufficiently cheap), the designer's problem simplifies to

$$\max_{\mathcal{I} \in \mathbb{I}} \frac{\lambda(\mathcal{I})}{c(\mathcal{I})}.$$

Thus, the optimal information structure can be determined solely based on the learning efficiency index and per-sample cost, and the solution is robust across all games and objectives satisfying Assumptions 1–2.

5.3 Higher-Order Expectations

Beyond its use in the proofs of Theorems 1–3, Lemma 1 can shed light on agents' higher-order beliefs more broadly. To see this, consider a finite set of types T_i for each agent i, with $T := \prod_{i \in I} T_i$. Let $\pi \in \Delta(T)$ be a (full-support) common prior over type profiles, with marginals $\pi_i \in \Delta(T_i)$. Each type $t_i \in T_i$ of player i induces a conditional distribution $\pi(\cdot \mid t_i) \in \Delta(T)$ over type profiles. By identifying each $t_j \in T_j$ with the point-mass distribution $\delta_{t_j} \in \Delta(T_j)$, we can associate with $\pi(\cdot \mid t_i)$ a sequence of higher-order expectations about other agents' types. In particular, $\mathbb{E}_{t_i}[t_j] := \sum_{t_j \in T_j} \pi(t_j \mid t_i) \delta_{t_j} \in \Delta(T_j)$ is t_i's expectation of j's type, $\mathbb{E}_{t_i} \mathbb{E}_{t_j}[t_k] := \sum_{t_j \in T_j, t_k \in T_k} \pi(t_j \mid t_i) \pi(t_k \mid t_j) \delta_{t_k} \in \Delta(T_k)$ is t_i's expectation of j's expectation of k's type, and so on.

A seminal result due to Samet (1998) is that any such sequence of higher-order expectations converges to the prior distribution as the number of iterations grows large. Formally, consider any sequence of agents $i_0, i_1, \ldots \in I$ in which all $i \in I$ appear infinitely often and any initial type $t_{i_0} \in T_{i_0}$. Then

$$\left\| \mathbb{E}_{t_{i_0}} \mathbb{E}_{t_{i_1}} \cdots \mathbb{E}_{t_{i_{k-1}}}[t_{i_k}] - \pi_{i_k} \right\| \to 0 \text{ as } k \to \infty.$$

By applying Lemma 1 to this setting, we can formalize a sense in which agents' higher-order expectations grow closer to the prior distribution at *each step* of the iteration. In particular, Lemma 1 implies that

$$\mathrm{KL}(\mathbb{E}_{t_{i_0}}[t_{i_1}], \pi_{i_1}) \geq \mathrm{KL}(\mathbb{E}_{t_{i_0}} \mathbb{E}_{t_{i_1}}[t_{i_2}], \pi_{i_2}),$$

and iteratively, for each k,

$$\mathrm{KL}(\mathbb{E}_{t_{i_0}} \mathbb{E}_{t_{i_1}} \cdots \mathbb{E}_{t_{i_{k-1}}}[t_{i_k}], \pi_{i_k}) \geq \mathrm{KL}(\mathbb{E}_{t_{i_0}} \mathbb{E}_{t_{i_1}} \cdots \mathbb{E}_{t_{i_k}}[t_{i_{k+1}}], \pi_{i_{k+1}}).$$

Thus, complementing Samet's asymptotic result, this clarifies that the informativeness of agents' higher-order expectations, as measured by their KL-divergence relative

to the prior distribution, decreases monotonically along any sequence.

Appendix: Proofs

A   Preliminaries

Let the transition matrix M_{ij}^θ and events $F_{it}(\theta, d)$, $F_t(\theta, d)$ be as defined in Section 3.3.

A.1   Proof of Lemma 1

We prove the following more general claim. Lemma 1 follows since $\mu_i^\theta M_{ij}^\theta = \mu_j^\theta$.

Lemma A.1. *For each $\theta \in \Theta$, distinct $i, j \in I$, and $\nu_i, \nu_i' \in \Delta(X_i)$ with $\operatorname{supp}(\nu_i) \subseteq \operatorname{supp}(\nu_i')$, we have $\operatorname{KL}(\nu_i, \nu_i') \geq \operatorname{KL}(\nu_i M_{ij}^\theta, \nu_i' M_{ij}^\theta)$. The inequality is strict whenever $\nu_i \neq \nu_i'$ and signals are not perfectly correlated.*

Proof. Consider $m, m' \in \Delta(X_i \times X_j)$ defined by

$$m(x_i, x_j) = \nu_i(x_i) M_{ij}^\theta(x_i, x_j), \quad m'(x_i, x_j) = \nu_i'(x_i) M_{ij}^\theta(x_i, x_j)$$

for each x_i, x_j. Note that $\operatorname{supp}(m) \subseteq \operatorname{supp}(m')$ and that the marginals of m, m' on X_i are ν_i, ν_i', and the marginals on X_j are $\nu_i M_{ij}^\theta, \nu_i' M_{ij}^\theta$, respectively.

Let $m(\cdot \mid x_i)$, $m(\cdot \mid x_j)$, $m'(\cdot \mid x_i)$, $m'(\cdot \mid x_i)$ denote the corresponding conditional distributions; conditional on a zero-probability signal, we specify these distributions arbitrarily. By the chain rule for KL-divergence we have

$$
\begin{aligned}
\operatorname{KL}(m, m') &= \operatorname{KL}(\nu_i, \nu_i') + \sum_{x_i \in \operatorname{supp}(\nu_i)} \nu_i(x_i) \operatorname{KL}(m(\cdot \mid x_i), m'(\cdot \mid x_i)) \\
&= \operatorname{KL}(\nu_i M_{ij}^\theta, \nu_i' M_{ij}^\theta) + \sum_{x_j \in \operatorname{supp}(\nu_i M_{ij}^\theta)} (\nu_i M_{ij}^\theta)(x_j) \operatorname{KL}(m(\cdot \mid x_j), m'(\cdot \mid x_j)).
\end{aligned}
$$

Since $m(\cdot \mid x_i) = m'(\cdot \mid x_i) = M_{ij}^\theta(x_i, \cdot)$ for every $x_i \in \operatorname{supp}(\nu_i)$, we have

$$\sum_{x_i \in \operatorname{supp}(\nu_i)} \nu_i(x_i) \operatorname{KL}(m(\cdot \mid x_i), m'(\cdot \mid x_i)) = 0,$$

which implies the weak inequality $\operatorname{KL}(\nu_i, \nu_i') \geq \operatorname{KL}(\nu_i M_{ij}^\theta, \nu_i' M_{ij}^\theta)$.

To show the strict inequality, suppose that $\nu_i \neq \nu_i'$ and signals are not perfectly correlated. Then there exist x_i, x_i' such that $\nu_i(x_i) > \nu_i'(x_i)$ and $\nu_i(x_i') < \nu_i'(x_i')$. For

any $x_j \in \text{supp}(\nu_i M_{ij}^\theta)$,

$$\frac{m(x_i \mid x_j)}{m(x_i' \mid x_j)} = \frac{\nu_i(x_i) M_{ij}^\theta(x_i, x_j)}{\nu_i(x_i') M_{ij}^\theta(x_i', x_j)} \neq \frac{\nu_i'(x_i) M_{ij}^\theta(x_i, x_j)}{\nu_i'(x_i') M_{ij}^\theta(x_i', x_j)} = \frac{m'(x_i \mid x_j)}{m'(x_i' \mid x_j)},$$

where the inequality holds since $M_{ij}^\theta(x_i, x_j), M_{ij}^\theta(x_i', x_j) > 0$ by the full-support assumption on μ^θ. This guarantees

$$\sum_{x_j \in \text{supp}(\nu_i' M_{ij}^\theta)} (\nu_i M_{ij}^\theta)(x_j) \text{KL}(m(\cdot \mid x_j), m'(\cdot \mid x_j)) > 0$$

by Gibbs inequality. $\qquad\square$

A.2 Preliminary lemmas

Let $\|\cdot\|$ denote the sup norm for finite-dimensional real vectors. The following result is proved by CEMS (Lemma 3) based on a concentration inequality:

Lemma A.2. *For any $\varepsilon > 0$ and $q < 1$, there is T such that for all $t \geq T$, $\theta \in \Theta$, $i \in I$,*

$$\mathbb{P}_t^{\mathcal{I}}(\{\|\nu_{it} M_{ij}^\theta - \nu_{jt}\| < \varepsilon, \forall j \neq i\} \mid x_i^t, \theta) > q.$$

Let $F_{-it}(\theta, d) := \bigcap_{j \neq i} F_{jt}(\theta, d)$. The following result follows from Lemma 1 and Lemma A.2 and plays a key role in the proofs of Theorems 1–3:

Lemma A.3. *Consider any $i \in I$ and partition Π_i over Θ. Fix any $\theta \in \Theta$, $d \in (0, \min_{i \in I, \theta' \notin \Pi_i(\theta)} d(\mu_i^\theta, \mu_i^{\theta'}))$ and $p \in (0, 1)$. There exists T such that for all $t \geq T$,*

$$\text{KL}(\nu_{it}, \mu_i^\theta) \leq d \quad \Longrightarrow \quad \mathbb{P}_t^{\mathcal{I}}\left(\bigcup_{\theta' \in \Pi_i(\theta)} (\{\theta'\} \cap F_{-it}(\theta', d)) \mid x_i^t\right) \geq p. \qquad (17)$$

Proof. We only consider the case in which signals are not perfectly correlated. Under perfect correlation, the argument is straightforward.

Claim 1: *There exist $\kappa \in \left(0, \min_{i \in I, \theta' \notin \Pi_i(\theta)} d(\mu_i^\theta, \mu_i^{\theta'}) - d\right)$ and $T' > 0$ such that for all $t \geq T'$ and $\theta' \in \Theta$,*

$$\text{KL}(\nu_{it}, \mu_i^{\theta'}) \leq d + \kappa \quad \Longrightarrow \quad \mathbb{P}_t^{\mathcal{I}}(F_{-it}(\theta', d) \mid x_i^t, \theta') \geq \sqrt{p}.$$

Proof of Claim 1. Lemma 1 implies that for all $j \neq i$, $\nu_i \in \Delta(X_i)$, and $\theta' \in \Theta$,

$$\mathrm{KL}(\nu_i, \mu_i^{\theta'}) \leq d \implies \mathrm{KL}(\nu_i M_{ij}^{\theta'}, \mu_j^{\theta'}) \leq \mathrm{KL}(\nu_i, \mu_i^{\theta'}) \leq d.$$

Moreover, the first inequality on the RHS is strict when $\nu_i \neq \mu_i^{\theta'}$ (by Lemma 1), and the second inequality on the RHS is strict when $\nu_i = \mu_i^{\theta'}$. Note that $\mathrm{KL}(\cdot, \mu_i)$ is continuous for each full-support $\mu_i \in \Delta(X_i)$. Thus, since $\Delta(X_i)$ is compact, there exists $\eta > 0$ such that for all $j \neq i$, $\nu_i \in \Delta(X_i)$, and $\theta' \in \Theta$,

$$\mathrm{KL}(\nu_i, \mu_i^{\theta'}) \leq d \implies \mathrm{KL}(\nu_i M_i^{\theta'}, \mu_j^{\theta'}) \leq d - \eta.$$

Then there exists $\kappa \in (0, \min_{i \in I, \theta' \notin S_i(\theta)} d(\mu_i^\theta, \mu_i^{\theta'}) - d)$ such that for all $j \neq i$, $\nu_i \in \Delta(X_i)$, and $\theta' \in \Theta$,

$$\mathrm{KL}(\nu_i, \mu_i^{\theta'}) \leq d + \kappa \implies \mathrm{KL}(\nu_i M_i^{\theta'}, \mu_j^{\theta'}) \leq d - \eta/2.$$

Moreover, there exists $\varepsilon > 0$ such that for all $j \neq i$, $\nu_i \in \Delta(X_i)$, and $\theta' \in \Theta$,

$$\left[\mathrm{KL}(\nu_i, \mu_i^{\theta'}) \leq d + \kappa \text{ and } \|\nu_i M_{ij}^{\theta'} - \nu_j\| \leq \varepsilon \right] \implies \mathrm{KL}(\nu_j, \mu_j^{\theta'}) \leq d.$$

Thus, by Lemma A.2, there exists T' such that $\mathbb{P}_t^{\mathcal{I}}(F_t(\theta', d) \mid x_i^t, \theta') \geq \sqrt{p}$ holds for all $t \geq T'$ and $\theta' \in \Theta$. $\qquad \square$

Claim 2: *Consider any κ as found in Claim 1. There exists T'' such that for all $t \geq T''$,*

$$\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d \implies \mathbb{P}_t^{\mathcal{I}}(\{\theta' \in \Pi_i(\theta) : \mathrm{KL}(\nu_{it}, \mu^{\theta'}) \leq d + \kappa\} \mid x_i^t) \geq \sqrt{p}.$$

Proof of Claim 2. Take any $t \geq 1$ and x_i^t such that $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d$. Then for each $\theta' \notin \Pi_i(\theta)$, we have $\mathrm{KL}(\nu_{it}, \mu_i^{\theta'}) > d + \kappa$. Indeed, otherwise $\mathrm{KL}(\nu_i', \mu_i^\theta), \mathrm{KL}(\nu_i', \mu_i^{\theta'}) < d + \kappa \leq d(\mu_i^\theta, \mu_i^{\theta'})$ holds for some $\nu_i' = (1 - \varepsilon)\nu_{it} + \varepsilon\mu_i^{\theta'}$ with $\varepsilon > 0$ small enough, contradicting the definition of $d(\mu_i^\theta, \mu_i^{\theta'})$.

Thus, whenever $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d$, then for any θ' such that either $\theta' \notin \Pi_i(\theta)$ or

$\mathrm{KL}(\nu_{it}, \mu^{\theta'}) > d + \kappa$, we have

$$\log \mathbb{P}_t^{\mathcal{I}}(\theta' \mid x_i^t) \leq \log \frac{\mathbb{P}_t^{\mathcal{I}}(\theta'|x_i^t)}{\mathbb{P}_t^{\mathcal{I}}(\theta|x_i^t)} \quad = \quad \log \frac{p_0(\theta')}{p_0(\theta)} + t \sum_{x_i \in X_i} \nu_{it}(x_i) \log \frac{\mu_i^{\theta'}(x_i)}{\mu_i^\theta(x_i)}$$

$$= \quad \log \frac{p_0(\theta')}{p_0(\theta)} + t(\mathrm{KL}(\nu_{it}, \mu_i^\theta) - \mathrm{KL}(\nu_{it}, \mu_i^{\theta'}))$$

$$\leq \quad \log \frac{p_0(\theta')}{p_0(\theta)} - t\kappa.$$

Hence, by choosing $T'' > 0$ large enough, we have that for all $t \geq T''$ and all θ' such that either $\theta' \notin \Pi_i(\theta)$ or $\mathrm{KL}(\nu_{it}, \mu^{\theta'}) > d + \kappa$,

$$\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d \quad \Longrightarrow \quad \mathbb{P}_t^{\mathcal{I}}(\theta'|x_i^t) < \frac{1 - \sqrt{p}}{|\Theta|},$$

proving Claim 2. $\qquad \square$

Finally, to prove Lemma A.3, let $T = \max\{T', T''\}$, with T' and T'' as found in Claims 1–2. Then, whenever $t \geq T$ and $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d$, we have

$$\mathbb{P}_t^{\mathcal{I}}\Big(\bigcup_{\theta' \in \Pi_i(\theta)} (\{\theta'\} \cap F_{-it}(\theta', d)) \mid x_i^t \Big) \geq \sum_{\theta' \in \Pi_i(\theta) \text{ s.t. } \mathrm{KL}(\nu_{it}, \mu^{\theta'}) \leq d + \kappa} \mathbb{P}_t^{\mathcal{I}}(\{\theta'\} \cap F_{-it}(\theta', d) \mid x_i^t)$$

$$= \sum_{\theta' \in \Pi_i(\theta) \text{ s.t. } \mathrm{KL}(\nu_{it}, \mu^{\theta'}) \leq d + \kappa} \mathbb{P}_t^{\mathcal{I}}(F_{-it}(\theta', d) \mid x_i^t, \theta')\mathbb{P}_t^{\mathcal{I}}(\theta' \mid x_i^t)$$

$$\geq \sum_{\theta' \in \Pi_i(\theta) \text{ s.t. } \mathrm{KL}(\nu_{it}, \mu^{\theta'}) \leq d + \kappa} \sqrt{p} \times \mathbb{P}_t^{\mathcal{I}}(\theta' \mid x_i^t) \geq p,$$

where the second inequality uses Claim 1 and the last inequality uses Claim 2. $\qquad \square$

B   Proof of Theorem 1

Fix any information structure $\mathcal{I}$, $\theta \in \Theta$ and $p \in (0, 1)$. We first establish that

$$\limsup_{t \to \infty} \frac{1}{t} \log \big(1 - \mathbb{P}_t^{\mathcal{I}}(C_t^p(\theta) \mid \theta) \big) \leq -\lambda^\theta(\mathcal{I}). \tag{18}$$

Take any $d \in (0, \lambda^\theta(\mathcal{I}))$. Applying Lemma A.3 to the case with $\Pi_i(\theta) = \{\theta\}$ for each $i \in I$, there exists $T > 0$ such that (i) $F_t(\theta, d) \subseteq B_t^p(\theta)$, and (ii) $F_t(\theta, d) \subseteq B_t^p(F(\theta, d))$

for each $i \in I$ and $t \geq T$. Thus, by Monderer and Samet (1989), we have $F_t(\theta, d) \subseteq C_t^p(\theta)$ for all $t \geq T$. Therefore,

$$
\begin{aligned}
\limsup_{t \to \infty} \frac{1}{t} \log \left(1 - \mathbb{P}_t^{\mathcal{I}}(C_t^p(\theta) \mid \theta) \right) &\leq \limsup_{t \to \infty} \frac{1}{t} \log \left(1 - \mathbb{P}_t^{\mathcal{I}}(F_t(\theta, d) \mid \theta) \right) \\
&\leq \limsup_{t \to \infty} \frac{1}{t} \log \left(\sum_i \mathbb{P}_t^{\mathcal{I}}(\{\mathrm{KL}(\nu_{it}, \mu_i^\theta) > d\} \mid \theta) \right) \\
&= \max_i \limsup_{t \to \infty} \frac{1}{t} \log \mathbb{P}_t^{\mathcal{I}}(\{\mathrm{KL}(\nu_{it}, \mu_i^\theta) > d\} \mid \theta) \\
&= -d,
\end{aligned}
$$

where the last equality follows from Sanov's theorem. Since this holds for all $d < \lambda^\theta(\mathcal{I})$, this establishes (18).

We next establish that

$$
\liminf_{t \to \infty} \frac{1}{t} \log \left(1 - \mathbb{P}_t^{\mathcal{I}}(B_t^q(\theta) \mid \theta) \right) \geq -\lambda^\theta(\mathcal{I}). \tag{19}
$$

Take $i \in I$ and $\theta' \neq \theta$ such that $d(\mu_i^\theta, \mu_i^{\theta'}) = \lambda^\theta(\mathcal{I})$. Take any $d > d(\mu_i^\theta, \mu_i^{\theta'})$. Then there is $\nu_i \in \Delta(X_i)$ with $\mathrm{KL}(\nu_i, \mu_i^\theta) = \mathrm{KL}(\nu_i, \mu_i^{\theta'}) < d$. Hence for some ν_i' nearby ν_i,

$$
\mathrm{KL}(\nu_i', \mu_i^{\theta'}) < \mathrm{KL}(\nu_i', \mu_i^\theta) < d.
$$

Thus, there exist $\varepsilon > 0$ and an open set $K_i \ni \nu_i'$ of signal distributions such that for all $\nu_i'' \in K_i$,

$$
\mathrm{KL}(\nu_i'', \mu_i^{\theta'}) + \varepsilon < \mathrm{KL}(\nu_i'', \mu_i^\theta) < d.
$$

Then, for all large enough t, $B_{it}^p(\theta) \cap \{\nu_{it} \in K_i\} = \emptyset$, because by standard arguments, i's beliefs at large t concentrate on states whose signal distributions minimize KL-divergence relative to ν_{it}. Thus,

$$
\liminf_{t \to \infty} \frac{1}{t} \log \left(1 - \mathbb{P}_t^{\mathcal{I}}(B_{it}^p(\theta) \mid \theta) \right) \geq \liminf_{t \to \infty} \frac{1}{t} \log \mathbb{P}_t^{\mathcal{I}}(\{\nu_{it} \in K_i\} \mid \theta) \geq -d,
$$

where the final inequality holds by Sanov's theorem. Since this is true for all $d > \lambda^\theta(\mathcal{I})$, this establishes (19). $\qquad \square$

C Proof of Theorems 2–3

Below we prove Theorem 3. Theorem 2 corresponds to the special case in which $\Pi_i(\theta) = \{\theta\}$ for all θ and i. To simplify notation, we drop the superscript W from $a^{\theta,W}$ when there is no risk of confusion.

C.1 Bounds on inefficiency

For any $\mathcal{I}$, $\mathcal{G}$, and W, we first derive bounds on the probability of inefficient play (i.e., not playing a^θ in state θ) as t grows large. The following result provides a *lower* bound on this probability for arbitrary sequences of strategy profiles (σ_t):

Lemma C.1. *Fix any $\mathcal{I}$, $\mathcal{G}$, and W. For any sequence of strategy profiles (σ_t) of $\mathcal{G}_t(\mathcal{I})$,*

$$\liminf_{t\to\infty} \max_\theta \frac{1}{t} \log \left(1 - \sum_{x^t \in X^t} \mathbb{P}_t^{\mathcal{I}}(x^t \mid \theta) \sigma_t(a^\theta \mid x^t) \right) \geq -\lambda(\mathcal{I}, \Pi^W).$$

Proof. Pick i, θ, and $\theta' \notin \Pi_i^W(\theta)$ such that $\lambda(\mathcal{I}, \Pi^W) = d(\mu_i^\theta, \mu_i^{\theta'})$. Consider any sequence of strategy profiles (σ_t) of $\mathcal{G}_t(\mathcal{I})$. Consider modified strategies $(\tilde{\sigma}_{it})$ for player i such that, for each x_i^t,

1. $\tilde{\sigma}_{it}(a_i^\theta \mid x_i^t) \geq \sigma_{it}(a_i^\theta \mid x_i^t)$ and $\tilde{\sigma}_{it}(a_i^{\theta'} \mid x_i^t) \geq \sigma_{it}(a_i^{\theta'} \mid x_i^t)$

2. $\tilde{\sigma}_{it}(a_i^\theta \mid x_i^t) + \tilde{\sigma}_{it}(a_i^{\theta'} \mid x_i^t) = 1$.

That is, $(\tilde{\sigma}_{it})$ is obtained by shifting all weight (σ_{it}) puts on actions other than $a_i^\theta, a_i^{\theta'}$ to $a_i^\theta, a_i^{\theta'}$ at all signal realizations.

We also consider the sequence of strategies (σ_{it}^*) given by

$$\begin{cases} \sigma_{it}^*(a_i^\theta \mid x_i^t) &= 1 \text{ if } \mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq \mathrm{KL}(\nu_{it}, \mu_i^{\theta'}) \\ \sigma_{it}^*(a_i^{\theta'} \mid x_i^t) &= 1 \text{ if } \mathrm{KL}(\nu_{it}, \mu_i^\theta) > \mathrm{KL}(\nu_{it}, \mu_i^{\theta'}), \end{cases}$$

where ν_{it} is the empirical signal frequency associated with x_i^t. Note that σ_{it}^* can be seen as a likelihood ratio test (with threshold 1). Thus, the Neyman-Pearson lemma for randomized tests (Theorem 3.2.1 in Lehmann and Romano, 2006) implies that for

32

each t,

$$\sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta)\tilde{\sigma}_{it}(a_i^\theta | x_i^t) \leq \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta)\sigma_{it}^*(a_i^\theta | x_i^t)$$

$$\text{or } \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta')\tilde{\sigma}_{it}(a_i^{\theta'} | x_i^t) \leq \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta')\sigma_{it}^*(a_i^{\theta'} | x_i^t). \tag{20}$$

Hence,

$$\liminf_{t \to \infty} \frac{1}{t} \log \left(\max \left\{ 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta)\sigma_{it}(a_i^\theta | x_i^t), 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta')\sigma_{it}(a_i^{\theta'} | x_i^t) \right\} \right)$$

$$\geq \liminf_{t \to \infty} \frac{1}{t} \log \left(\max \left\{ 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta)\tilde{\sigma}_{it}(a_i^\theta | x_i^t), 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta')\tilde{\sigma}_{it}(a_i^{\theta'} | x_i^t) \right\} \right)$$

$$\geq \liminf_{t \to \infty} \frac{1}{t} \log \left(\min \left\{ 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta)\sigma_{it}^*(a_i^\theta | x_i^t), 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta')\sigma_{it}^*(a_i^{\theta'} | x_i^t) \right\} \right)$$

$$= \min_{\theta'' \in \{\theta, \theta'\}} \liminf_{t \to \infty} \frac{1}{t} \log \left(1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta'')\sigma_{it}^*(a_i^{\theta''} | x_i^t) \right),$$

where the first inequality follows from the construction of $(\tilde{\sigma}_{it})$ and the second inequality uses (20). The last line is equal to $-d(\mu_i^\theta, \mu_i^{\theta'}) = -\lambda(\mathcal{I}, \Pi^W)$, because the asymptotic error rate under a likelihood-ratio test with threshold 1 is given by Chernoff information (Theorem 3.4.3 in Dembo and Zeitouni, 2010),[15] i.e.,

$$\lim_{t \to \infty} \frac{1}{t} \log \left(1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta)\sigma_{it}^*(a_i^\theta | x_i^t) \right) = \lim_{t \to \infty} \frac{1}{t} \log \left(1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta')\sigma_{it}^*(a_i^{\theta'} | x_i^t) \right)$$

$$= -d(\mu_i^\theta, \mu_i^{\theta'}).$$

This implies that

$$\liminf_{t \to \infty} \max_{\theta'' \in \Theta} \frac{1}{t} \log \left(1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta'')\sigma_{it}(a_i^{\theta''} | x_i^t) \right) \geq -\lambda(\mathcal{I}, \Pi^W),$$

---

[15]This in turn follows from a simple application of Sanov's theorem.

as claimed. □

Under Assumption 1, the following result provides an *upper* bound on the probability of inefficient play under some *equilibrium* sequence (σ_t):

Lemma C.2. *Fix any $\mathcal{I}$ and any $(\mathcal{G}, W)$ satisfying Assumption 1. There exists a sequence of BNE strategy profiles $(\sigma_t) \in \mathrm{BNE}_t(\mathcal{G}, \mathcal{I})$ such that, for all $\theta \in \Theta$,*

$$\limsup_{t \to \infty} \frac{1}{t} \log \left(1 - \sum_{x^t \in X^t} \mathbb{P}_t^{\mathcal{I}}(x^t \mid \theta) \sigma_t(a^\theta \mid x^t) \right) \leq -\lambda(\mathcal{I}, \Pi^W).$$

Proof. Take $p \in (0, 1)$ sufficiently close to 1 that for any i and θ, choosing a_i^θ is optimal whenever i's belief about the state and opponents' actions assigns probability at least p to $\{(\theta', a_{-i}^{\theta'}) : \theta' \in \Pi_i^W(\theta)\}$. Such a p exists because, by Assumption 1, a_i^θ is the unique maximizer of $u_i(\cdot, a_{-i}^{\theta'}, \theta')$ for each $\theta' \in \Pi_i^W(\theta)$.

Fix any $d < \lambda(\mathcal{I}, \Pi^W) := \min_{i \in I, \theta \in \Theta, \theta' \notin \Pi_i(\theta)} d(\mu_i^\theta, \mu_i^{\theta'})$. Let $\Sigma_{it}(d)$ denote the set of i's strategies at t such that $\sigma_{it}(a^\theta \mid x_i^t) = 1$ whenever $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d$. This set is well-defined by the choice of d, i.e., there is no $\nu_i \in \Delta(X_i)$ such that $\mathrm{KL}(\nu_i, \mu_i^\theta), \mathrm{KL}(\nu_i, \mu_i^{\theta'}) \leq d$ for some θ and $\theta' \notin \Pi_i^W(\theta)$.

Given such a p and d, use Lemma A.3 to construct a large enough T such that (17) holds for all i and θ and $t \geq T$. Then for all $t \geq T$, each i's best response against any strategy profile in $\prod_{j \neq i} \Sigma_{jt}(d)$ must be in $\Sigma_{it}(d)$. This is because each agent i with $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq d$ assigns probability at least p to $\{(\theta', a_{-i}^{\theta'}) : \theta' \in \Pi_i^W(\theta)\}$. Thus, applying Kakutani's fixed point theorem to the best-response correspondences defined on the restricted strategy space $\prod_i \Sigma_{it}(d)$ for all $t \geq T$, there exists a BNE sequence (σ_t) such that $\sigma_t(a^\theta \mid x^t) = 1$ at $F_t(\theta, d)$ for every θ.

For this sequence of BNEs (σ_t) and for all θ, we have that as $t \to \infty$,

$$1 - \sum_{x^t \in X^t} \mathbb{P}_t^{\mathcal{I}}(\theta, x^t) \sigma_t(a^\theta \mid x^t) \leq \sum_i \mathbb{P}_t^{\mathcal{I}}(\{\mathrm{KL}(\nu_{it}, \mu_i^\theta) > d\}) = \exp[-td + o(t)],$$

where the equality follows from Sanov's theorem. Since this holds for all $d < \lambda(\mathcal{I}, \Pi^W)$, this yields the desired conclusion. □

C.2 Proof of Theorem 3

We prove that 1. implies 2. The converse is then immediate from the assumption that $\lambda(\mathcal{I}, \Pi) \neq \lambda(\tilde{\mathcal{I}}, \Pi)$.

Fix any information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$ with $\lambda(\mathcal{I}, \Pi) > \lambda(\tilde{\mathcal{I}}, \Pi)$, and any $(\mathcal{G}, W)$ satisfying Assumption 1 and $\Pi^W = \Pi$. Since $\{a^\theta\} = \arg\max_a W(a, \theta)$ for each $\theta \in \Theta$, there exist constants $c \geq \tilde{c} > 0$ such that for all t, strategy profiles σ_t of $\mathcal{G}_t(\mathcal{I})$ and $\tilde{\sigma}_t$ and $\mathcal{G}_t(\tilde{\mathcal{I}})$, and all $\theta \in \Theta$,

$$W(a^\theta, \theta) - \sum_{x^t, a} \mathbb{P}_t^{\mathcal{I}}(x^t \mid \theta)\sigma_t(a \mid x^t)W(a, \theta) \leq c \left(1 - \sum_{x^t} \mathbb{P}_t^{\mathcal{I}}(x^t \mid \theta)\sigma_t(a^\theta \mid x^t)\right), \quad (21)$$

$$W(a^\theta, \theta) - \sum_{\tilde{x}^t, a} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}^t \mid \theta)\tilde{\sigma}_t(a \mid \tilde{x}^t)W(a, \theta) \geq \tilde{c} \left(1 - \sum_{\tilde{x}^t} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}^t \mid \theta)\tilde{\sigma}_t(a^\theta \mid \tilde{x}^t)\right). \quad (22)$$

By Lemma C.2, there exists a sequence of BNE $\sigma_t \in \mathrm{BNE}_t(\mathcal{G}, \mathcal{I})$ such that

$$
\begin{aligned}
-\lambda(\mathcal{I}, \Pi) &\geq \max_\theta \limsup_{t\to\infty} \frac{1}{t} \log\left(1 - \sum_{x^t} \mathbb{P}_t^{\mathcal{I}}(x^t \mid \theta)\sigma_t(a^\theta \mid x^t)\right) \\
&= \limsup_{t\to\infty} \frac{1}{t} \log \sum_\theta p_0(\theta)\left(1 - \sum_{x^t} \mathbb{P}_t^{\mathcal{I}}(x^t \mid \theta)\sigma_t(a^\theta \mid x^t)\right),
\end{aligned}
$$

which by (21) implies

$$\limsup_{t\to\infty} \frac{1}{t} \log \sum_\theta p_0(\theta)\left(W(a^\theta, \theta) - \sum_{x^t} \mathbb{P}_t^{\mathcal{I}}(x^t \mid \theta)\sigma_t(a^\theta \mid x^t)\right) \leq -\lambda(\mathcal{I}, \Pi). \quad (23)$$

Let $\tilde{\sigma}_t$ denote a strategy profile that maximizes $W_t(\cdot, \tilde{\mathcal{I}})$. By Lemma C.1,

$$
\begin{aligned}
-\lambda(\tilde{\mathcal{I}}, \Pi) &\leq \liminf_{t\to\infty} \max_\theta \frac{1}{t} \log\left(1 - \sum_{\tilde{x}^t} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}^t \mid \theta)\tilde{\sigma}_t(a^\theta \mid \tilde{x}^t)\right) \\
&\leq \liminf_{t\to\infty} \frac{1}{t} \log \sum_\theta p_0(\theta)\left(1 - \sum_{\tilde{x}^t} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}^t \mid \theta)\tilde{\sigma}_t(a^\theta \mid \tilde{x}^t)\right),
\end{aligned}
$$

which by (22) implies

$$\liminf_{t \to \infty} \frac{1}{t} \log \sum_{\theta} p_0(\theta) \left(W(a^\theta, \theta) - \sum_{\tilde{x}^t} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}^t \mid \theta) \tilde{\sigma}_t(a^\theta \mid \tilde{x}^t) \right) \geq -\lambda(\tilde{\mathcal{I}}, \Pi). \quad (24)$$

Thus, for all large enough t, we have $W_t(\mathcal{G}, \mathcal{I}) \geq W_t(\sigma_t, \mathcal{I}) > W_t(\tilde{\sigma}_t, \tilde{\mathcal{I}}) \geq W_t(\mathcal{G}, \tilde{\mathcal{I}})$, where the strict inequality follows from (23) and (24) and the assumption that $\lambda(\mathcal{I}, \Pi) > \lambda(\tilde{\mathcal{I}}, \Pi)$. □

References

ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2016): "Fragility of asymptotic agreement under Bayesian learning," *Theoretical Economics*, 11(1), 187–225.

ATHEY, S., AND J. LEVIN (2018): "The value of information in monotone decision problems," *Research in Economics*, 72(1), 101–116.

AZRIELI, Y. (2014): "Comment on "the law of large demand for information"," *Econometrica*, 82(1), 415–423.

BERGEMANN, D., AND S. MORRIS (2016): "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 11(2), 487–522.

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

BÖRGERS, T., A. HERNANDO-VECIANA, AND D. KRÄHMER (2013): "When are signals complements or substitutes?," *Journal of Economic Theory*, 148(1), 165–195.

CARLSSON, H., AND E. VAN DAMME (1993): "Global games and equilibrium selection," *Econometrica*, pp. 989–1018.

COVER, T. M., AND J. A. THOMAS (1999): *Elements of information theory*. John Wiley & Sons.

CRIPPS, M. W., J. C. ELY, G. J. MAILATH, AND L. SAMUELSON (2008): "Common learning," *Econometrica*, 76(4), 909–933.

———— (2013): "Common learning with intertemporal dependence," *International Journal of Game Theory*, 42(1), 55–98.

DASARATHA, K., AND K. HE (2019): "Aggregative Efficiency of Bayesian Learning in Networks," working paper.

DEMBO, A., AND O. ZEITOUNI (2010): *Large Deviations Techniques and Applications*. Springer.

DUFFIE, D., AND G. MANSO (2007): "Information percolation in large markets," *American Economic Review*, 97(2), 203–209.

DUPUIS, P., AND R. S. ELLIS (2011): *A weak convergence approach to the theory of large deviations*, vol. 902. John Wiley & Sons.

FRICK, M., R. IIJIMA, AND Y. ISHII (2021): "Welfare Comparisons for Biased Learning," working paper.

FUDENBERG, D., G. LANZANI, AND P. STRACK (2021): "Pathwise Concentration Bounds for Misspecified Bayesian Beliefs," *Available at SSRN 3805083*.

GOSSNER, O. (2000): "Comparison of information structures," *Games and Economic Behavior*, 30(1), 44–63.

HANN-CARUTHERS, W., V. V. MARTYNOV, AND O. TAMUZ (2018): "The speed of sequential asymptotic learning," *Journal of Economic Theory*, 173, 383–409.

HAREL, M., E. MOSSEL, P. STRACK, AND O. TAMUZ (2021): "Rational groupthink," *The Quarterly Journal of Economics*, 136(1), 621–668.

LEHMANN, E. L. (1988): "Comparing Location Experiments," *The Annals of Statistics*, 16(2), 521–533.

LEHMANN, E. L., AND J. P. ROMANO (2006): *Testing statistical hypotheses*. Springer Science & Business Media.

LEHRER, E., D. ROSENBERG, AND E. SHMAYA (2010): "Signaling and mediation in games with common interests," *Games and Economic Behavior*, 68(2), 670–682.

LIANG, A. (2019): "Games of incomplete information played by statisticians," working paper.

LIANG, A., AND X. MU (2020): "Complementary information and learning traps," *The Quarterly Journal of Economics*, 135(1), 389–448.

MONDERER, D., AND D. SAMET (1989): "Approximating common knowledge with common beliefs," *Games and Economic Behavior*, 1(2), 170–190.

MOSCARINI, G., AND L. SMITH (2002): "The law of large demand for information," *Econometrica*, 70(6), 2351–2366.

MU, X., L. POMATTO, P. STRACK, AND O. TAMUZ (2021): "From Blackwell Dominance in Large Samples to Rényi Divergences and Back Again," *Econometrica*, 89(1), 475–506.

PERSICO, N. (2000): "Information acquisition in auctions," *Econometrica*, 68(1), 135–148.

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

ROSENBERG, D., AND N. VIEILLE (2019): "On the efficiency of social learning," *Econometrica*, 87(6), 2141–2168.

RUBINSTEIN, A. (1989): "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"," *The American Economic Review*, pp. 385–391.

SAMET, D. (1998): "Iterated expectations and common priors," *Games and Economic Behavior*, 24(1-2), 131–141.

STEINER, J., AND C. STEWART (2011): "Communication, timing, and common learning," *Journal of Economic Theory*, 146(1), 230–247.

VIVES, X. (1993): "How fast do rational agents learn?," *The Review of Economic Studies*, 60(2), 329–347.

WEINSTEIN, J., AND M. YILDIZ (2007): "A structure theorem for rationalizability with application to robust predictions of refinements," *Econometrica*, 75(2), 365–400.