Learning Efficiency of Multi-Agent Information Structures*

Mira Frick Ryota Iijima Yuhta Ishii

First posted version: August 17, 2021
This version: December 21, 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 that are "aligned at certainty." We discuss implications of our results for constrained information design in games and for the question when information structures are complements vs. substitutes.

Keywords: common learning, speed of learning, higher-order beliefs, comparison of information structures.

*Frick: Yale University (mira.frick@yale.edu); Iijima: Yale University (ryota.iijima@yale.edu); Ishii: Pennsylvania State University (yxi5014@psu.edu). We thank numerous seminar audiences, as well as Nageeb Ali, Dirk Bergemann, Roberto Corrao, Drew Fudenberg, Olivier Gossner, Johannes Hörner, Andreas Kleiner, Vijay Krishna, Giacomo Lanzani, Bart Lipman, Antonio Penta, Larry Samuelson, Eddie Schlee, Ron Siegel, Jakub Steiner, Colin Stewart, Philipp Strack, and Nick Wu.

1 Introduction

1.1 Overview

Suppose a group of players (e.g., firms) are about to engage 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 prior to playing a game, 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 that are "aligned at certainty," permitting a robust comparison of information structures that does not require an understanding of the full details of the strategic environment.

To address both these questions, we introduce 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 (in the sense of common p-belief) 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 agents' 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 identify a class of games and objectives for which information structures with a higher learning efficiency index induce better (Bayes-Nash) equilibrium outcomes whenever agents observe sufficiently many signal draws. The substantive 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, this 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 on incomplete information games is that higher-order belief uncertainty can be a significant source of inefficiency, our results suggest that, when agents have access to large samples of signals, reducing higher-order belief uncertainty becomes a second-order concern.

Section 5 discusses further implications of our results for constrained information design in games, the question when information structures are complements vs. substitutes, and the informativeness of agents' higher-order expectations.

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 successful 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 (best-case) 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 as follows:

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

Each information structure is summarized by two parameters: The individual precision parameter $\gamma \in (1/2, 1)$ captures the probability with which each player's signal matches the state; the parameter $\rho \in [0, 1]$ captures the extent of correlation across players' signals (signals are independent if $\rho = 0$ and perfectly correlated if $\rho = 1$). 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.[1]

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 the speed of learning. 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.[2] 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 of Theorem 1 uses Markov chain arguments that are related to CEMS' approach, but

[1] For example, suppose players observe only one signal draw. Then the best equilibrium (under utilitarian welfare) varies across different parameters (p_0, c) (e.g., it might involve each player i always investing in project x_i, or i investing in project 1 when $x_i = 1$ and not investing when $x_i = -1$), and each of these possibilities induces a different ranking over (γ, ρ).

[2] 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.

is based on a different construction.

Moscarini and Smith (2002) derive an efficiency index that characterizes the speed of single-agent learning.[3] 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 learning index, while the correlation across agents' signals plays no role.

The speed of learning has also been analyzed in various social learning environments, but most work has not focused on the role of higher-order beliefs.[4] 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 the agent observes a single signal draw. 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.[5] In particular, Lehrer, Rosenberg, and Shmaya

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

[4]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).

[5]Bergemann and Morris (2016) study general games using a different approach. They consider

(2010) focus on common interest games with utilitarian welfare, and characterize the order based on a generalization of Blackwell's garbling condition, while 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, as we discuss in Section 4.2, by assuming that agents observe many signal draws, we obtain a ranking that is a completion of Lehrer, Rosenberg, and Shmaya's (2010) order and applies to a richer class of games and objective functions beyond the common-interest (i.e., full-alignment) 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 private 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 private signals in state θ. We assume that, for all agents i and states θ, μ_i^θ has full support and $\mu_i^\theta \neq \mu_i^{\theta'}$ for all $\theta' \neq \theta$. Note that the joint distribution μ^θ may display arbitrary correlation.

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 $\sigma_{it} : (X_i)^t \to \Delta(A_i)$ for agent i maps i's observed sequence of private signals $x_i^t = (x_{i\tau})_{\tau=1,\dots,t}$ to a mixed

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.

action in A_i. Note that each agent i only observes her own sequence of private signals x_i^t and that the game is played once at the end of the entire sequence of t signal draws.

Let $\text{BNE}_t(\mathcal{G}, \mathcal{I})$ 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 $\text{BNE}_t(\mathcal{G}, \mathcal{I})$ if for each $i \in I$, $x_i^t \in X_i^t$, and a_i with $\sigma_{it}(a_i \mid x_i^t) > 0$,

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

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 $C_t^p(\theta)$ 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 (as made precise by 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 different information structures $\mathcal{I}$ affect equilibrium play in game $\mathcal{G}_t(\mathcal{I})$ at finite t. 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, $\mathrm{KL}(\nu,\mu) := \sum_{y\in Y} \nu(y) \log \frac{\nu(y)}{\mu(y)}$ 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

[6]See Section 5.4 for a discussion of more general settings.

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

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 their KL-midpoint. Thus, the smaller $d(\mu, \mu')$, the more difficult it is to statistically 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 $d(\mu, \mu')$ is symmetric and 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*

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

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 $\lambda^\theta(\mathcal{I}) \geq \lambda^\theta(\tilde{\mathcal{I}})$ for all θ.

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 $\lambda^\theta(\mathcal{I})$.[8]

That is, when agents observe large samples of signals, higher-order belief uncertainty vanishes at least 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 Appendices B and D. 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 $\mathbb{P}_t^{\mathcal{I}}(B_{it}^p(\theta) \mid \theta) \to 1$) 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. ▲

[8]The $o(t)$ terms can differ across (4) and (5) and can depend on p_0, p, and features of $\mathcal{I}$ other than $\lambda^\theta(\mathcal{I})$, but these terms become negligible as $t \to \infty$.

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 $\lambda_i^\theta(\mathcal{I}) = \min_{\theta' \in \Theta \setminus \{\theta\}} d(\mu_i^\theta, \mu_i^{\theta'})$, 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,

$$\begin{aligned} &\left\{ \mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq \min_{\theta' \in \Theta \setminus \{\theta\}} \mathrm{KL}(\nu_{it}, \mu_i^{\theta'}) - \varepsilon \right\} \subseteq B_{it}^p(\theta) \\ \subseteq &\left\{ \mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq \min_{\theta' \in \Theta \setminus \{\theta\}} \mathrm{KL}(\nu_{it}, \mu_i^{\theta'}) + \varepsilon \right\}. \end{aligned} \tag{8}$$

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 $D_i := \{\nu_i \in \Delta(X_i) : \mathrm{KL}(\nu_i, \mu_i^\theta) \leq \min_{\theta' \in \Theta \setminus \{\theta\}} \mathrm{KL}(\nu_i, \mu_i^{\theta'})\}$, it can be seen that $\inf_{\nu_i \notin D_i} \mathrm{KL}(\nu_i, \mu_i^\theta) = \lambda_i^\theta(\mathcal{I})$.[9] 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

[9] Indeed, note that

$$\begin{aligned} \inf_{\nu_i \notin D_i} \mathrm{KL}(\nu_i, \mu_i^\theta) &= \inf \left\{ \mathrm{KL}(\nu_i, \mu_i^\theta) : \mathrm{KL}(\nu_i, \mu_i^\theta) > \mathrm{KL}(\nu_i, \mu_i^{\theta'}) \text{ for some } \theta' \neq \theta \right\} \\ &= \min_{\theta' \in \Theta \setminus \{\theta\}} \left\{ \mathrm{KL}(\nu_i, \mu_i^\theta) : \mathrm{KL}(\nu_i, \mu_i^\theta) = \mathrm{KL}(\nu_i, \mu_i^{\theta'}) \right\} = \min_{\theta' \in \Theta \setminus \{\theta\}} d(\mu_i^\theta, \mu_i^{\theta'}) = \lambda_i^\theta(\mathcal{I}). \end{aligned}$$

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

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

Below we illustrate the argument, assuming for simplicity that each joint distribution $\mu^\theta \in \Delta(X)$ has full support. Fix any $d < \lambda^\theta(\mathcal{I})$. For each t, we consider the event

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

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

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

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 (given (9)) 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 $\nu_{it} M_{ij}^\theta \in \Delta(X_j)$ (treating $\nu_{it} \in \Delta(X_i) \subseteq \mathbb{R}^{1 \times X_i}$ as a vector). 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 μ^θ has full support and $\nu_i \neq \mu_i^\theta$.*

To understand the result, suppose that i observes an empirical signal distribution ν_{it} at any t. Then $\mathrm{KL}(\nu_{it}, \mu_i^\theta)$ quantifies how much i's observations deviate from i's theoretical signal distribution μ_i^θ in state θ. Likewise, $\mathrm{KL}(\nu_{it} M_{ij}^\theta, \mu_j^\theta)$ quantifies how much i's *expectation* of j's observations deviates from j's theoretical signal distribution μ_j^θ. Thus, Lemma 1 says that when i forms an estimate of j's signal observations based on i's own signal observations, then (conditional on any state θ) this estimate cannot be 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 theoretical distribution (i.e., $\mathrm{KL}(\nu_{it} M_{ij}^\theta, \mu_j^\theta) = 0$). 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., $\mathrm{KL}(\nu_{it} M_{ij}^\theta, \mu_j^\theta) = \mathrm{KL}(\nu_{it}, \mu_i^\theta)$).

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, event $F_t(\theta, d)$ is p-evident at all large t. □

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 (an iteration of) 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.[10]

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 Online Appendix E 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, and we apply our learning efficiency index to rank information structures in terms of their induced equilibrium outcomes.

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, $\{a^{\theta,W}\} = \operatorname{argmax}_{a\in A} W(a,\theta)$. 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

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

[10] Sugaya (2021); Fong, Gossner, Hörner, and Sannikov (2010); Chan and Zhang (2018) use arguments related to CEMS' contraction result in the context of repeated games with private monitoring.

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***

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

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, using our learning efficiency index, 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 $\mathrm{SNE}(\mathcal{G},\theta) \subseteq A$ denote the set of strict Nash equilibria of $\mathcal{G}$ under common knowledge of θ.

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

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.[11] 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.[12] 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,

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

[12] In particular, Assumption 1 allows for environments where, away from the common knowledge limit, improving players' information can lead to worse equilibrium outcomes; see the discussion of Lehrer, Rosenberg, and Shmaya (2010) on p. 18-19.

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 settings 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 $a^\theta \in \text{SNE}(\mathcal{G}, \theta)$ 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:

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

Assumption 2 is satisfied, for instance, in the joint investment game in Section 1.1, where $a_i^{\theta,W} = \theta$ for all i, θ. Section 4.3 will drop this assumption.

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 $t > T$.*

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.

Based on the structure of $\lambda(\mathcal{I})$, Theorem 2 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. At the same time, 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; Kajii and Morris, 1997; 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.[13] Both (i) and (ii) rely on the assumption that agents observe sufficiently many signal draws: When $t = 1$, many information structures are incomparable even when restricting attention to common-interest games. Moreover, when $t = 1$, then even if $\mathcal{I}$ is more informative than $\tilde{\mathcal{I}}$ in the sense of Lehrer, Rosenberg, and Shmaya (2010), $\mathcal{I}$ can be strictly worse than $\tilde{\mathcal{I}}$ in some environments that satisfy Assumptions 1–2 but feature misaligned incentives.

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

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

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 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., $\mathrm{KL}(\nu_{it}, \mu_i^\theta) < \mathrm{KL}(\nu_{it}, \mu_i^{\theta'})$ 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

[13]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 $\lambda(\mathcal{I}) \geq \lambda(\tilde{\mathcal{I}})$.

which each agent i plays action $a_i^{\theta,W}$ in the event that θ is common p-belief at t.[14] 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 every 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 $\text{SNE}(\mathcal{G}, \theta)$, 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 Online Appendix G, we use the learning efficiency index to characterize the speed at which the entire equilibrium set $\text{BNE}_t(\mathcal{G}, \mathcal{I})$ approaches the set of common knowledge equilibria in each state. ▲

4.3 Ranking without Full Separation

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. Maintaining Assumption 1, we generalize Theorem 2 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

[14] 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.

agent i, whose cells are given by

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

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. Let $\Pi^W := (\Pi_i^W)_{i\in I}$ denote the collection of all agents' partitions.

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

$$\lambda(\mathcal{I}, \Pi) := \min_{i\in I, \theta, \theta' \in \Theta, \theta' \notin \Pi_i(\theta)} d(\mu_i^\theta, \mu_i^{\theta'}).^{15}$$

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.

In the following result, we restrict attention to information structures that are either ***fully private***, in the sense that each joint distribution μ^θ has full support on X, or ***public***, in the sense that signals are perfectly correlated across agents.[16]

Theorem 3. *Fix any collection $\Pi = (\Pi_i)_{i\in I}$ of partitions over Θ. Take any information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$, each of which is either fully private or public, and suppose that $\lambda(\mathcal{I}, \Pi) \neq \lambda(\tilde{\mathcal{I}}, \Pi)$. 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.

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

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

Theorem 3 also immediately 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 $\tilde{\mathcal{I}}$, each of which is either fully private or public, and suppose 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),

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

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. Online Appendix F 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 Discussion

5.1 Information Design in Games

The preceding analysis 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 $\kappa > 0$.

Then, for any $\mathcal{G}$ and W satisfying Assumptions 1–2 and any finite set $\mathbb{I}$, our analysis implies that, whenever the budget κ is sufficiently large (i.e., information is sufficiently cheap), the designer's problem simplifies to[17]

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

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. Based on this observation, future work might explore properties of the optimal information structure under multi-agent generalizations of information cost functions c in the literature (e.g., Pomatto, Strack, and Tamuz, 2020).

5.2 Information Structures as Complements vs. Substitutes

Our learning efficiency index also suggests a novel formalization of when two information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are complements or substitutes.[18] To this end, we extend our baseline setting with repeated draws from a single information structure $\mathcal{I}$ by

[17]Indeed, as $\kappa \to \infty$, the analysis in Section 4.2 implies that it is optimal to exhaust the budget and that the difference between the first-best payoff $\sum_\theta p_0(\theta)W(a^{\theta,W},\theta)$ and the best equilibrium payoff under information structure $\mathcal{I}$ takes the form $\exp[-\kappa\frac{\lambda(\mathcal{I})}{c(\mathcal{I})} + o(\kappa)]$.

[18]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.

considering the effect of combining signal observations from $\mathcal{I} = (X, (\mu^\theta)_{\theta\in\Theta})$ and $\tilde{\mathcal{I}} = (\tilde{X}, (\tilde{\mu}^\theta)_{\theta\in\Theta})$. Let $\mathcal{I} \times \tilde{\mathcal{I}} := (X \times \tilde{X}, (\mu^\theta \times \tilde{\mu}^\theta)_{\theta\in\Theta})$ denote the combined information structure under which the signal distribution in 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.[19] Likewise, 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.[20]

[19] 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 $2t$ 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.

[20] 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. $\mathrm{KL}(\nu_i, \mu_i^\theta) = \mathrm{KL}(\nu_i, \mu_i^{\theta'})$. Com-

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 $\mathcal{I}$ and $\tilde{\mathcal{I}}$ have in common a worst-informed agent. Thus, only the first channel is relevant and $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are substitutes. In particular, (under binary states) this is always the case if there is only one 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. ▲

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. In particular, it can be used to analyze the "informativeness" of agents' higher-order expectations, which plays an important role, for instance, in beauty-contest games (e.g., Morris and Shin, 2002; Golub and Morris, 2017).

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.

---

bined 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 $\mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta) \neq \mathrm{KL}(\tilde{\nu}_i, \tilde{\mu}_i^{\theta'})$.

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 his result adapted to the current setting implies that[21]

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

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.

5.4 More General Information Structures

This paper has characterized the speed of common learning in settings where signal spaces are finite and signals are generated i.i.d. across draws. One challenge in moving beyond these settings is that it is not known under which general conditions common learning obtains.

With infinite signals, CEMS exhibit a setting in which common learning fails even though individual learning is successful (see their Section 4); at the same time, there are other natural infinite-signal settings, in particular Gaussian signal structures, that do give rise to common learning.[22] In Online Appendix H, we analyze the latter Gaussian environment. We again show that common learning and individual learning

[21] See the proof of his Proposition 6.

[22] In contrast, Dogan (2018) shows that with (uncountably) infinite states, common learning fails under mild conditions (even if signals are finite).

occur at the same exponential rate, which depends only on the worst-informed agent's signal precision.

When signals are correlated across draws, it is also known that common learning can fail even when individual learning is successful (e.g., Steiner and Stewart, 2011), while Cripps, Ely, Mailath, and Samuelson (2013) exhibit some settings with intertemporally correlated signals that give rise to common learning. We leave the analysis of such settings for future work, in particular, the question whether there are environments in which common learning is successful but occurs at a slower rate than individual learning.

Finally, farther afield, one might consider settings in which players engage in basic game $\mathcal{G}$ not only once, at t, but repeatedly following each signal draw. In this case, players' past actions can reveal information about their private signals. Basu, Chatterjee, Hoshino, and Tamuz (2020) and Sugaya and Yamamoto (2020) study such settings and construct equilibria that lead to common learning. An interesting open question is to analyze the speed of common learning and how this is affected by players' strategic incentives.

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

Fix $\theta \in \Theta$, distinct $i, j \in I$, and $\nu_i \in \Delta(X_i)$. Define $m, m' \in \Delta(X_i \times X_j)$ by

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

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

Let $m(\cdot \mid x_i)$, $m(\cdot \mid x_j)$, $m'(\cdot \mid x_i)$, $m'(\cdot \mid x_j)$ 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

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

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

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

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

To show the strict inequality, suppose that $\nu_i \neq \mu_i^\theta$ and μ^θ has full support on X. Then there exist x_i, x_i' such that $\nu_i(x_i) > \mu_i^\theta(x_i)$ and $\nu_i(x_i') < \mu_i^\theta(x_i')$. For any $x_j \in \mathrm{supp}(\nu_i M_{ij}^\theta)$,

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

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. By Gibbs' inequality, this guarantees

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

and hence $\mathrm{KL}(\nu_i, \mu_i^\theta) > \mathrm{KL}(\nu_i M_{ij}^\theta, \mu_j^\theta)$. □

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

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

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

Lemma A.2. *Take any collection of partitions $(\Pi_i)_{i\in I}$ over Θ, $\theta\in\Theta$, $p\in(0,1)$, and $d\in(0,\min_{i\in I,\theta'\notin\Pi_i(\theta)} d(\mu_i^\theta,\mu_i^{\theta'}))$. Assume that μ^θ has full support. There exists T such that for all $i\in I$ and $t\geq T$,*

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

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

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

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 $\theta'\in\Theta$,

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

Given this, there exists $\kappa\in(0,\min_{i\in I,\theta'\notin\Pi_i(\theta)} d(\mu_i^\theta,\mu_i^{\theta'})-d)$ 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 $\theta'\in\Theta$,

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

Combined with Lemma A.1, this yields the desired conclusion. □

Claim 2. *Consider any κ as found in Claim 1. There exists T'' such that for all*

$t \geq T''$ and $i \in I$,

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

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

$$\begin{aligned}
\log \mathbb{P}_t^{\mathcal{I}}(\theta' \mid x_i^t) \leq \log \frac{\mathbb{P}_t^{\mathcal{I}}(\theta' | x_i^t)}{\mathbb{P}_t^{\mathcal{I}}(\theta | x_i^t)} &= \log \frac{p_0(\theta')}{p_0(\theta)} + t \sum_{x_i \in X_i} \nu_{it}(x_i) \log \frac{\mu_i^{\theta'}(x_i)}{\mu_i^\theta(x_i)} \\
&= \log \frac{p_0(\theta')}{p_0(\theta)} + t(\mathrm{KL}(\nu_{it}, \mu_i^\theta) - \mathrm{KL}(\nu_{it}, \mu_i^{\theta'})) \\
&\leq \log \frac{p_0(\theta')}{p_0(\theta)} - t\kappa.
\end{aligned}$$

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. □

Finally, to prove Lemma A.2, 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

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

where the second inequality uses Claim 1 and the last inequality uses Claim 2. □

B Proof of Theorem 1 (Fully Private Case)

This appendix proves Theorem 1, assuming for ease of exposition that information structure $\mathcal{I}$ is fully private, i.e., the joint distribution μ^θ has full support on X for each θ. Appendix D extends the proof to general information structures.

Fix any $\theta \in \Theta$ and $p \in (0,1)$. We first establish that

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

Take any $d \in (0, \lambda^\theta(\mathcal{I}))$. Applying Lemma A.2 to the case with $\Pi_i(\theta) = \{\theta\}$ for each $i \in I$, there exists $T > 0$ such that, for all $t \geq T$, (i) $F_t(\theta, d) \subseteq B_t^p(\theta)$, and (ii) $F_t(\theta, d) \subseteq B_t^p(F(\theta, d))$. 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' close to 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, $B^p_{it}(\theta) \cap \{\nu_{it} \in K_i\} = \emptyset$, because by standard arguments, i's beliefs at large t concentrate on states whose signal distributions minimize KL-divergence relative to ν_{it}. Thus,

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

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

C Proof of Theorem 2 (Fully Private Case) and Theorem 3

Below we prove Theorem 3. When $\mathcal{I}$ and $\tilde{\mathcal{I}}$ are either fully private or public, Theorem 2 then follows as the special case in which $\Pi_i(\theta) = \{\theta\}$ for all θ and i. Appendix D proves Theorem 2 for general information structures. 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 $\mathcal{G}_t(\mathcal{I})$,*

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

Proof. Pick i, θ, and $\theta' \notin \Pi^W_i(\theta)$ such that $\lambda(\mathcal{I}, \Pi^W) = d(\mu^\theta_i, \mu^{\theta'}_i)$. 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^t_i,

1. $\tilde{\sigma}_{it}(a^\theta_i \mid x^t_i) \geq \sigma_{it}(a^\theta_i \mid x^t_i)$ and $\tilde{\sigma}_{it}(a^{\theta'}_i \mid x^t_i) \geq \sigma_{it}(a^{\theta'}_i \mid x^t_i)$
2. $\tilde{\sigma}_{it}(a^\theta_i \mid x^t_i) + \tilde{\sigma}_{it}(a^{\theta'}_i \mid x^t_i) = 1$.

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 distribution 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 each t,

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

Hence,

$$\begin{aligned} & \liminf_{t\to\infty} \frac{1}{t} \log \left(\max \left\{ 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta) \sigma_{it}(a_i^\theta | x_i^t), 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta') \sigma_{it}(a_i^{\theta'} | x_i^t) \right\} \right) \\ \geq \; & \liminf_{t\to\infty} \frac{1}{t} \log \left(\max \left\{ 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta) \tilde{\sigma}_{it}(a_i^\theta | x_i^t), 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta') \tilde{\sigma}_{it}(a_i^{\theta'} | x_i^t) \right\} \right) \\ \geq \; & \liminf_{t\to\infty} \frac{1}{t} \log \left(\min \left\{ 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta) \sigma_{it}^*(a_i^\theta | x_i^t), 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta') \sigma_{it}^*(a_i^{\theta'} | x_i^t) \right\} \right) \\ = \; & \min_{\theta'' \in \{\theta, \theta'\}} \liminf_{t\to\infty} \frac{1}{t} \log \left(1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta'') \sigma_{it}^*(a_i^{\theta''} | x_i^t) \right), \end{aligned}$$

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),[23] 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

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

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 $\mathcal{I}$ that is either fully private or public 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 such that, for all i and θ, choosing a_i^θ is u_i-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 $\sigma_{it}(a_i^\theta\mid x_i^t)=1$ 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 $\theta'\notin\Pi_i^W(\theta)$.

We show that there exists T such that for any $t>T$, there is a BNE σ_t of $\mathcal{G}_t(\mathcal{I})$ with $\sigma_{it}\in\Sigma_{it}(d)$ for every i. To see this, first consider the case in which $\mathcal{I}$ is fully private. Then, by Lemma A.2 with p as chosen above, there is T such that (17) holds for all i, θ, and $t\geq T$. Thus, for all $t\geq T$, each agent i's best response against any

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

strategy profile in $\prod_{j\neq i}\Sigma_{jt}(d)$ must be in $\Sigma_{it}(d)$, because whenever $\mathrm{KL}(\nu_{it},\mu_i^\theta)\leq d$, then i assigns probability at least p to $\{(\theta', a_{-i}^{\theta'}) : \theta' \in \Pi_i^W(\theta)\}$. Thus, for every $t\geq T$, applying Kakutani's fixed point theorem to the best-response correspondences defined on the restricted strategy space $\prod_i \Sigma_{it}(d)$, we obtain a BNE σ_t of $\mathcal{G}_t(\mathcal{I})$ such that $\sigma_{it}\in\Sigma_{it}(d)$ for every i. Next, suppose $\mathcal{I}$ is public. In this case, all players' posteriors coincide, i.e., $\mathbb{P}_t^{\mathcal{I}}(\cdot|x_t^i) = \mathbb{P}_t^{\mathcal{I}}(\cdot|x_t^j)$ for all i, j, and t. Moreover, $\mathrm{KL}(\nu_{it},\mu_i^\theta)\leq d \iff \mathrm{KL}(\nu_{jt},\mu_j^\theta)\leq d$ for all i,j,t. Thus, if we choose T large enough, the same argument as in Claim 2 in the proof of Lemma A.2 ensures that

$$\mathrm{KL}(\nu_{it},\mu_i^\theta)\leq d \quad\Longrightarrow\quad \mathbb{P}_t^{\mathcal{I}}(\{\theta'\in\bigcap_j \Pi_j(\theta)\}\mid x_i^t)\geq p$$

for all $t\geq T$. Based on this observation, the same argument as in the fully private case yields a sequence of BNE $\sigma_t\in\prod_i\Sigma_{it}(d)$ for all $t\geq T$.

The above implies that there is a sequence of BNEs (σ_t) such that for all θ, we have that, as $t\to\infty$,

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

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 Remaining Proof

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 $\tilde{\mathcal{I}}$, each of which is either fully private or public, and any $(\mathcal{G},W)$ satisfying Assumption 1 and $\Pi^W=\Pi$. Suppose $\lambda(\mathcal{I},\Pi) > \lambda(\tilde{\mathcal{I}},\Pi)$. Since A is finite and $\{a^\theta\}=\arg\max_a W(a,\theta)$ 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 $\tilde{\sigma}_t$ of $\mathcal{G}_t(\tilde{\mathcal{I}})$, and all $\theta\in\Theta$,

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

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

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

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

which by (21) implies

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

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

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

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

D Proofs of Theorems 1–2 (General Case)

In this section, we extend the proofs of Theorems 1–2 to general information structures that need not be fully private. The main complication stems from the fact that the strict inequality part of Lemma 1 need not hold when μ^θ does not have full support.

We handle this issue by modifying the events $F_t(\theta, d)$ appropriately.

Fix any information structure $\mathcal{I}$ and state θ. Let $X^\theta \subseteq X$ denote the support of μ^θ. Conditional on state θ, define $H_i^\theta = (h_i^\theta(x))_{x \in X^\theta}$ to be agent i's information partition of X^θ based on observing her own private signal; that is

$$h_i^\theta(x) := \{x' \in X^\theta : x'_i = x_i\}, \quad \text{for all } x \in X^\theta.$$

For any distribution $\nu \in \Delta(X^\theta)$ and any partition H of X^θ, let $\nu_H \in \Delta(H)$ denote the induced distribution over the cells in H; that is, $\nu_H(h) := \sum_{x \in h} \nu(x)$ for all $h \in H$. Letting $\nu_t \in \Delta(X^\theta)$ denote the joint empirical distribution of signals up to t, note that $(\nu_t)_{H_i^\theta}$ can be identified with i's empirical distribution ν_{it}. For each subset of agents $S \subseteq I$, define $H_S^\theta := \bigwedge_{i \in S} H_i^\theta$ to be the finest common coarsening of all the partitions H_i^θ with $i \in S$. For any joint empirical signal distribution ν_t, distribution $(\nu_t)_{H_S^\theta}$ is commonly known among all agents in S.

Finally, for any $d > 0$ and $\varepsilon_1, \ldots, \varepsilon_{|I|} \in [0, d)$, define the following event:

$$F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|}) := \left\{ x^t \in (X^\theta)^t : \mathrm{KL}\left((\nu_t)_{H_S^\theta}, \mu^\theta_{H_S^\theta} \right) \leq d - \varepsilon_{|S|}, \ \forall S \subseteq I \right\}.$$

Note that, for any $i \in S$, $\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \geq \mathrm{KL}\left((\nu_t)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right)$. Thus, $F_t(\theta, d, 0, \ldots, 0) = F_t(\theta, d)$. Observe also that if μ^θ has full support, then $H_S^\theta = \{X\}$ for all non-singleton S, so $F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|}) = F_t(\theta, d - \varepsilon_1)$.

The main step in extending the proofs of Theorems 1–2 is the following result, which we prove in Appendix D.1.

Proposition D.1. *Take any $d \in (0, \lambda^\theta(\mathcal{I}))$ and $\varepsilon \in (0, d)$. There exists a sequence $\varepsilon = \varepsilon_n > \cdots > \varepsilon_2 > \varepsilon_1 = 0$ such that, for all $p \in (0,1)$, there exists T such that*

$$\mathbb{P}_t^{\mathcal{I}}\left(\{\theta\} \cap F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|}) \mid x_i^t \right) \geq p$$

holds for every $i \in I$, $t \geq T$, and signal sequence $x^t \in F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|})$.

Using Proposition D.1, the proof of Theorem 1 extends as follows. It suffices to prove (18) for general $\mathcal{I}$, as the argument for (19) in Appendix B did not rely on the full-support assumption. To prove (18), take any $d \in (0, \lambda^\theta(\mathcal{I}))$ and $\varepsilon \in (0, d)$. Then for all $p \in (0,1)$ and large enough t, the events $F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|})$ constructed in Proposition D.1 satisfy

$$F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|}) \subseteq C_t^p(\theta),$$

since Proposition D.1 ensures that these events are p-evident and $F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|}) \subseteq B_t^p(\theta)$ at large t by the usual argument. Moreover, by Sanov's theorem and the fact that $F_t(\theta, d, 0, \ldots, 0) = F_t(\theta, d)$,

$$\begin{aligned}
&\lim_{\varepsilon_k \to 0 \forall k} \lim_{t\to\infty} \frac{1}{t} \log\left(1 - \mathbb{P}_t^{\mathcal{I}}\left(F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|}) \mid \theta\right)\right) \\
&= \lim_{t\to\infty} \frac{1}{t} \log\left(1 - \mathbb{P}_t^{\mathcal{I}}\left(F_t(\theta, d) \mid \theta\right)\right) = -d.
\end{aligned}$$

Since this holds for all $d < \lambda^\theta(\mathcal{I})$, (18) follows.

To extend the proof of Theorem 2, it is sufficient to establish Lemma C.2 for general $\mathcal{I}$ under Assumption 2, as the remaining steps of the proof in Appendix C did not rely on the full-support assumption. To this end, fix $p \in (0, 1)$ and $d \in (0, \lambda(\mathcal{I}))$ as in the original proof of Lemma C.2, and take any $\varepsilon \in (0, d)$. Applying Proposition D.1 and following the same steps as in the original proof of Lemma C.2, we construct a BNE sequence (σ_t) such that for all large enough t and each θ, we have $\sigma_t(a^\theta|x^t) = 1$ at all signal sequences $x^t \in F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|})$. Thus,

$$\lim_{t\to\infty} \frac{1}{t} \log\left(1 - \sum_{x^t \in X^t} \mathbb{P}_t^{\mathcal{I}}(x^t|\theta)\sigma_t(a^\theta \mid x^t)\right) \leq \lim_{t\to\infty} \frac{1}{t} \log\left(1 - \mathbb{P}_t^{\mathcal{I}}\left(F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|}) \mid \theta\right)\right).$$

As above, the right-hand side tends to $-d$ as $\varepsilon_k \to 0$ for each k. Since this holds for all $d < \lambda(\mathcal{I})$, we obtain the desired conclusion.

D.1 Proof of Proposition D.1

D.1.1 Generalization of Lemma 1

The key step in proving Proposition D.1 is the following generalization of Lemma 1. For each $i \in I$ and $\nu \in \Delta(X^\theta)$ with $\nu_i = \mathrm{marg}_{X_i}\nu$, define distribution $\nu M_i^\theta \in \Delta(X^\theta)$ by

$$(\nu M_i^\theta)(x_i, x_{-i}) := \nu_i(x_i)\mu^\theta(x_{-i}|x_i), \quad \text{for all } (x_i, x_{-i}) \in X^\theta. \tag{25}$$

When the joint empirical signal distribution is ν_t, then $\nu_t M_i^\theta$ is i's expectation of this joint distribution conditional on state θ and on observing ν_{it}.

Lemma D.1. *Take any $\nu \in \Delta(X^\theta)$, $i \in I$, and $S \subseteq I$. Then* $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) \leq \mathrm{KL}\left(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}\right)$*. Moreover, the inequality is an equality only if $\nu_{H_i^\theta}(\cdot|h) = \mu^\theta_{H_i^\theta}(\cdot|h)$ for every $h \in H_i^\theta \wedge H_S^\theta$ with $\nu_{H_i^\theta \wedge H_S^\theta}(h) > 0$.*

Proof. To show the inequality, first note that

$$\begin{aligned}\mathrm{KL}\left(\nu M_i^\theta, \mu^\theta\right) &= \mathrm{KL}\left((\nu M_i^\theta)_{H_i^\theta}, \mu^\theta_{H_i^\theta}\right) + \sum_{h \in H_i^\theta} (\nu M_i^\theta)_{H_i^\theta}(h) \mathrm{KL}((\nu M_i^\theta)(\cdot|h), \mu^\theta(\cdot|h)) \\ &= \mathrm{KL}\left(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}\right), \end{aligned} \tag{26}$$

where the first equality uses the chain rule for KL-divergence and the second one holds because $\nu_{H_i^\theta} = (\nu M_i^\theta)_{H_i^\theta}$ and $(\nu M_i^\theta)(\cdot|h) = \mu^\theta(\cdot|h)$ for each $h \in H_i^\theta$ by (25). The chain rule also implies that

$$\begin{aligned}\mathrm{KL}\left(\nu M_i^\theta, \mu^\theta\right) &= \mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) + \sum_{h \in H_S^\theta} (\nu M_i^\theta)_{H_S^\theta}(h) \mathrm{KL}\left((\nu M_i^\theta)(\cdot|h) \mid \mu^\theta(\cdot|h)\right) \\ &\geq \mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right). \end{aligned} \tag{27}$$

Combining (26)–(27) yields $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) \leq \mathrm{KL}\left(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}\right)$.

For the "moreover" part, suppose that $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) = \mathrm{KL}\left(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}\right)$. Then, by (26)-(27), for every $h \in H_S^\theta$ such that $(\nu M_i^\theta)_{H_S^\theta}(h) > 0$, we have $(\nu M_i^\theta)(\cdot|h) = \mu^\theta(\cdot|h)$. In addition, for any $h \in H_i^\theta$ such that $(\nu M_i^\theta)_{H_i^\theta}(h) > 0$, (25) implies $(\nu M_i^\theta)(\cdot|h) = \mu^\theta(\cdot|h)$. These two observations yield that for any $h \in H_i^\theta \wedge H_S^\theta$ with $(\nu M_i^\theta)_{H_i^\theta \wedge H_S^\theta}(h) > 0$, we have $(\nu M_i^\theta)(\cdot|h) = \mu^\theta(\cdot|h)$, and hence $(\nu M_i^\theta)_{H_i^\theta}(\cdot|h) = \mu^\theta_{H_i^\theta}(\cdot|h)$. But by (25), $(\nu M_i^\theta)_{H_i^\theta} = \nu_{H_i^\theta}$ and $(\nu M_i^\theta)_{H_i^\theta \wedge H_S^\theta} = \nu_{H_i^\theta \wedge H_S^\theta}$. Thus, $\nu_{H_i^\theta}(\cdot|h) = \mu^\theta_{H_i^\theta}(\cdot|h)$ for all $h \in H_i^\theta \wedge H_S^\theta$ with $\nu_{H_i^\theta \wedge H_S^\theta}(h) > 0$. □

Lemma D.1 yields the following corollary:

Corollary D.1. *Take any $d > 0$ and $\varepsilon \in (0, d)$. There exists $\rho \in (0, \varepsilon)$ such that for all $S \subseteq I$, $i \notin S$, and $\nu \in \Delta(X^\theta)$ with*

$$\mathrm{KL}(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}) \leq d \quad \textit{and} \quad \max_{|S'|=|S|+1} \mathrm{KL}(\nu_{H_{S'}^\theta}, \mu^\theta_{H_{S'}^\theta}) \leq d - \varepsilon,$$

we have $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) < d - \rho$.

Proof. Consider any $S \subseteq I$, $i \notin S$, and $\nu \in \Delta(X^\theta)$ with $\mathrm{KL}(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}) \leq d$ and $\max_{|S'|=|S|+1} \mathrm{KL}(\nu_{H^\theta_{S'}}, \mu^\theta_{H^\theta_{S'}}) \leq d - \varepsilon$. It suffices to prove that $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) < d$, as the left-hand side of this inequality is continuous in ν and $\Delta(X^\theta)$ is compact.

To show the latter inequality, note that Lemma D.1 implies $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) \leq \mathrm{KL}(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}) \leq d$. Thus, we can focus on the case in which $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) = \mathrm{KL}(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta})$. In this case,

$$\begin{aligned}
\mathrm{KL}(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}) &= \mathrm{KL}\left(\nu_{H_i^\theta \wedge H_S^\theta}, \mu^\theta_{H_i^\theta \wedge H_S^\theta}\right) + \sum_{h \in H_i^\theta \wedge H_S^\theta} \nu_{H_i^\theta \wedge H_S^\theta}(h) \mathrm{KL}\left(\nu_{H_i^\theta}(\cdot|h), \mu^\theta_{H_i^\theta}(\cdot|h)\right) \\
&= \mathrm{KL}\left(\nu_{H_i^\theta \wedge H_S^\theta}, \mu^\theta_{H_i^\theta \wedge H_S^\theta}\right) \leq d - \varepsilon,
\end{aligned}$$

where the first equality uses the chain rule and the second one holds by the "moreover" part of Lemma D.1. □

D.1.2 Completing the Proof

To prove Proposition D.1, we first set $\varepsilon_{|I|} = \varepsilon$. By Corollary D.1, there exists $\rho_{|I|-1} \in (0, \varepsilon_{|I|})$ such that for all $i \in I$ and $S = I \setminus \{i\}$, whenever

$$\mathrm{KL}(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}) \leq d \quad \text{and} \quad \mathrm{KL}(\nu_{H_I^\theta}, \mu^\theta_{H_I^\theta}) \leq d - \varepsilon,$$

we have $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) < d - \rho_{|I|-1}$.

Next, choose some $\varepsilon_{|I|-1} \in (0, \rho_{|I|-1})$, and proceed inductively in the same manner. In particular, once we have constructed ε_{k+1}, use Corollary D.1 to find $\rho_k \in (0, \varepsilon_{k+1})$ such that for all $i \in I$ and $S \subseteq I$ with $|S| = k$ and $i \notin S$, whenever

$$\mathrm{KL}(\nu_{H_i^\theta}, \mu^\theta_{H_i^\theta}) \leq d \quad \text{and} \quad \max_{|S'|=k+1} \mathrm{KL}(\nu_{H^\theta_{S'}}, \mu^\theta_{H^\theta_{S'}}) \leq d - \varepsilon_{k+1},$$

we have $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) < d - \rho_k$.

This yields a sequence

$$\varepsilon = \varepsilon_{|I|} > \rho_{|I|-1} > \varepsilon_{|I|-1} > \cdots > \varepsilon_2 > \rho_1 > \varepsilon_1 = 0$$

with the property that whenever

$$\mathrm{KL}(\nu_{H_S^\theta}, \mu^\theta_{H_S^\theta}) \leq d - \varepsilon_{|S|} \text{ for all } S \subseteq I,$$

we have $\mathrm{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) < d - \rho_{|S|}$ for all $S \subseteq I$ and $i \notin S$.

We now show that this sequence is as required by Proposition D.1. As noted, for any $p \in (0,1)$ and sufficiently large t, we have $F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|}) \subseteq B_t^p(\theta)$. Thus, it suffices to show that, for any $p \in (0,1)$, $i \in I$, and $S \subseteq I$, there exists T such that

$$\mathbb{P}_t^{\mathcal{I}}\left(\left\{\mathrm{KL}\left((\nu_t)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) \leq d - \varepsilon_{|S|}\right\} \mid x_i^t, \theta\right) \geq p \tag{28}$$

holds for every $t \geq T$ and signal sequence $x^t \in F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|})$.

To show (28), fix any $p \in (0,1)$. First, consider $i \in I$ and $S \subseteq I$ with $i \in S$. Then H_S^θ is coarser than H_i^θ. Hence, for any $t \geq 1$ and signal sequence x^t with corresponding empirical distribution $\tilde{\nu}_t \in \Delta(X^\theta)$, we have

$$\mathbb{P}_t^{\mathcal{I}}\left(\left\{\nu_t \in \Delta(X^\theta) : (\nu_t)_{H_S^\theta} = (\tilde{\nu}_t)_{H_S^\theta}\right\} \mid x_i^t, \theta\right) = 1.$$

Thus, if $x^t \in F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|})$, then

$$\mathbb{P}_t^{\mathcal{I}}\left(\left\{\mathrm{KL}\left((\nu_t)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) \leq d - \varepsilon_{|S|}\right\} | x_i^t, \theta\right) = 1 > p,$$

as required.

Next, consider $i \in I$ and $S \subseteq I$ with $i \notin S$. Then the way in which sequence $(\varepsilon_k, \rho_k)_{k=1,\ldots,|I|}$ was constructed ensures that, for any $t \geq 1$ and $x^t \in F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|})$ with corresponding empirical frequency $\tilde{\nu}_t$, we have

$$\mathrm{KL}\left((\tilde{\nu}_t M_i^\theta)_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) \leq d - \rho_{|S|}. \tag{29}$$

Since $\rho_{|S|} > \varepsilon_{|S|}$ and $\Delta(X^\theta)$ is compact, there exists $\kappa > 0$ such that, for all $\nu, \nu' \in \Delta(X^\theta)$,

$$\mathrm{KL}\left(\nu'_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) \leq d - \rho_{|S|} \text{ and } \|\nu' - \nu\| < \kappa \implies \mathrm{KL}\left(\nu_{H_S^\theta}, \mu^\theta_{H_S^\theta}\right) \leq d - \varepsilon_{|S|}. \tag{30}$$

By the same law of large numbers argument as in the full-support case, there exists T such that, for all $t \geq T$ and signal sequences x^t with empirical distribution

$\tilde{\nu}_t$, we have

$$\mathbb{P}_t^{\mathcal{I}}\left(\left\{\left\|\nu_t - \tilde{\nu}_t M_i^\theta\right\| < \kappa\right\} \mid x_i^t, \theta\right) \geq p.$$

Combined with (29)–(30), this implies that (28) holds for every $t \geq T$ and signal sequence $x^t \in F_t(\theta, d, \varepsilon_1, \ldots, \varepsilon_{|I|})$. □

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.

Basu, P., K. Chatterjee, T. Hoshino, and O. Tamuz (2020): "Repeated coordination with private learning," *Journal of Economic Theory*, 190, 105106.

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.

Brandenburger, A., and E. Dekel (1993): "Hierarchies of beliefs and common knowledge," *Journal of Economic Theory*, 59(1), 189–198.

Carlsson, H., and E. Van Damme (1993): "Global games and equilibrium selection," *Econometrica*, pp. 989–1018.

Chan, J. H., and W. Zhang (2018): "Self-evident Events and the Value of Linking," working paper.

Chen, Y.-C., A. Di Tillio, E. Faingold, and S. Xiong (2010): "Uniform topologies on types," *Theoretical Economics*, 5(3), 445–478.

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.

Dekel, E., D. Fudenberg, and S. Morris (2006): "Topologies on types," *Theoretical Economics*, 1, 275.

Dembo, A., and O. Zeitouni (2010): *Large Deviations Techniques and Applications*. Springer.

Dogan, E. (2018): “Common Learning on a Continuum State Space,” working paper.

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.

Fong, K., O. Gossner, J. Hörner, and Y. Sannikov (2010): “Efficiency in the Repeated Prisoner’s Dilemma with Imperfect Private Monitoring,” working paper.

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*.

Golub, B., and S. Morris (2017): “Expectations, networks, and conventions,” working paper.

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.

Kajii, A., and S. Morris (1997): “The robustness of equilibria to incomplete information,” *Econometrica*, pp. 1283–1309.

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.

Lipman, B. L. (2003): “Finite order implications of common priors,” *Econometrica*, 71(4), 1255–1267.

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

Morris, S., and H. S. Shin (2002): “Social value of public information,” *American Economic Review*, 92(5), 1521–1534.

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.

Pomatto, L., P. Strack, and O. Tamuz (2020): "The cost of information," working paper.

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.

Shaked, M., and G. Shanthikumar (2007): *Stochastic orders*. Springer.

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

Sugaya, T. (2021): "Folk Theorem in Repeated Games with Private Monitoring," *Review of Economic Studies*, forthcoming.

Sugaya, T., and Y. Yamamoto (2020): "Common learning and cooperation in repeated games," *Theoretical Economics*, 15(3), 1175–1219.

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.

Online Appendix to "Learning Efficiency of Multi-Agent Information Structures"

Mira Frick, Ryota Iijima, and Yuhta Ishii

E Convergence of Belief Hierarchies

Theorem 1 showed that the learning efficiency index $\lambda^\theta(\mathcal{I})$ characterizes the speed at which players achieve approximate common knowledge of the true state in the sense of common p-belief. An analogous result holds if proximity to common knowledge is instead formalized in terms of commonly used topologies over belief hierarchies.

Recall that a belief hierarchy for player i is a sequence $\tau_i := (\tau_i^1, \tau_i^2, \ldots) \in Z_i = (Z_i^1, Z_i^2, \ldots)$, where $Z_i^1 := \Delta(\Theta)$ and $Z_i^k := \Delta(\Theta \times \prod_{j \neq i} Z_j^{k-1})$ denotes the space of player i's kth order beliefs, subject to standard coherency requirements across the kth order beliefs τ_i^k for different k (e.g., Brandenburger and Dekel, 1993).[24] Fix an information structure $\mathcal{I}$. Each observation of a signal sequence x_i^t induces a belief hierarchy $\tau_i(x_i^t) \in Z_i$ for player i. We let $\tau_i(\theta) \in Z_i$ denote the belief hierarchy for player i when there is common certainty of state θ.

Let ρ_i^{product} denote a metric on Z_i that induces the ***product topology*** over player i's belief hierarchies. For example, define $\rho_i^{\text{product}}(\tau_i, \tilde{\tau}_i) := \sum_k \beta^k \rho^k(\tau_i^k, \tilde{\tau}_i^k)$, where $\beta \in (0,1)$ and ρ^k denotes the Prokhorov metric over kth order beliefs. Since the product topology may in general be too coarse (e.g., Lipman, 2003; Weinstein and Yildiz, 2007), the literature has proposed several alternative metrics that refine this topology. In particular, consider the metric for the ***uniform-weak topology*** (Chen, Di Tillio, Faingold, and Xiong, 2010), which is given by $\rho_i^{\text{uniform}}(\tau_i, \tilde{\tau}_i) := \sup_k \rho^k(\tau_i^k, \tilde{\tau}_i^k)$. Then for all $\theta \in \Theta$ and sufficiently small $\varepsilon > 0$, Theorem 1 implies that, as $t \to \infty$,[25]

$$\mathbb{P}_t^{\mathcal{I}}(\{\max_i \rho_i^{\text{product}}(\tau_i(x_i^t), \tau_i(\theta)) < \varepsilon\} \mid \theta) = 1 - \exp[-\lambda^\theta(\mathcal{I})t + o(t)],$$

$$\mathbb{P}_t^{\mathcal{I}}(\{\max_i \rho_i^{\text{uniform}}(\tau_i(x_i^t), \tau_i(\theta)) < \varepsilon\} \mid \theta) = 1 - \exp[-\lambda^\theta(\mathcal{I})t + o(t)].$$

That is, the speed of convergence to common certainty is the *same* under both these topologies and is again given by the exponential rate $\lambda^\theta(\mathcal{I})$.[26] While differences

[24]For any topological space Y, we let $\Delta(Y)$ denote the space of Borel probability measures over Y and endow it with the topology of weak convergence.

[25]To see the latter equality, note that the proof of Proposition 6 in Chen, Di Tillio, Faingold, and Xiong (2010) implies that the ε-ball around $\tau_i(\theta)$ consists of all belief hierarchies for player i that have common $(1-\varepsilon)$-belief on θ.

[26]Analogous results hold under other metrics considered in the literature, for example, the metric for the strategic topology (Dekel, Fudenberg, and Morris, 2006), which is coarser than the product topology but finer than the uniform-weak topology.

between these topologies play a significant role in one-shot signal observation settings, our finding shows that these differences do not matter for the speed of convergence in the current large-sample setting.

F Monotone Information Structures

Suppose states are linearly ordered; without loss of generality, write $\Theta = \{1, \dots, n\}$. An information structure $\mathcal{I}$ is ***monotone*** if, for each agent $i \in I$, X_i is linearly ordered in such a way that the signal distributions $(\mu_i^\theta)_{\theta \in \Theta}$ satisfy the monotone likelihood-ratio property with respect to the orders over states and signals.

The following result shows that, for monotone information structures, the condition in Corollary 1 can be relaxed to one that is easier to verify:

Corollary F.1. *Let $\Theta = \{1, \dots, n\}$. Take any monotone information structures $\mathcal{I}$ and $\tilde{\mathcal{I}}$, each of which is either fully private or public, and suppose that $\min_{i \in I} d(\mu_i^\theta, \mu_i^{\theta+1}) \neq \min_{i \in I} d(\tilde{\mu}_i^\theta, \tilde{\mu}_i^{\theta+1})$ for all $\theta = 1, \dots, n-1$. The following are equivalent:*

(i). For any basic game $\mathcal{G}$ and objective W satisfying Assumption 1, there exists T such that $W_t(\mathcal{I}, \mathcal{G}) > W_t(\tilde{\mathcal{I}}, \mathcal{G})$ for all $t > T$.

(ii). We have $\min_{i \in I} d(\mu_i^\theta, \mu_i^{\theta+1}) > \min_{i \in I} d(\tilde{\mu}_i^\theta, \tilde{\mu}_i^{\theta+1})$ for all $\theta = 1, \dots, n-1$.

In Corollary 1, one has to calculate index $\lambda(\mathcal{I}, \Pi)$ for every partition Π of Θ, where the number of partitions grows exponentially as $|\Theta|$ becomes large. In contrast, for monotone structures, Corollary F.1 shows that it suffices to consider the Chernoff distance between each consecutive pair of signal distributions $\mu_i^\theta, \mu_i^{\theta+1}$, whose number grows only linearly as $|\Theta|$ becomes large.

The proof of Corollary 1 uses the following property of Chernoff distances under monotone information structures:

Lemma F.1. *Let $\Theta = \{1, ..., n\}$. Take any monotone information structure $\mathcal{I}$. For any $\theta < \theta' < \theta''$ and $i \in I$, we have $\max\{d(\mu_i^\theta, \mu_i^{\theta'}), d(\mu_i^{\theta'}, \mu_i^{\theta''})\} \leq d(\mu_i^\theta, \mu_i^{\theta''})$.*

Proof. We establish that $d(\mu_i^\theta, \mu_i^{\theta'}) \leq d(\mu_i^\theta, \mu_i^{\theta''})$. The proof of the remaining inequality is analogous. By the equivalent expression for Chernoff distance based on the Hellinger transform (Remark 1), we have that $d(\mu_i^\theta, \mu_i^{\theta'}) \leq d(\mu_i^\theta, \mu_i^{\theta''})$ if and only if

$$\min_{p \in [0,1]} \sum_{x_i} \mu_i^\theta(x_i) \left(\frac{\mu_i^{\theta'}(x_i)}{\mu_i^\theta(x_i)} \right)^p \geq \min_{p \in [0,1]} \sum_{x_i} \mu_i^\theta(x_i) \left(\frac{\mu_i^{\theta''}(x_i)}{\mu_i^\theta(x_i)} \right)^p.$$

By the concavity of $z \mapsto z^p$ for each $p \in [0,1]$ and because $\mu_i^{\theta'} \neq \mu_i^{\theta''}$, this inequality is satisfied if the distribution of $\frac{\mu_i^{\theta''}(x_i)}{\mu_i^\theta(x_i)}$ is a mean-preserving spread of the distribution of $\frac{\mu_i^{\theta'}(x_i)}{\mu_i^\theta(x_i)}$, when x_i is drawn according to μ_i^θ.

To show the latter, let $\geq_i$ denote the linear order on X_i. By the monotone likelihood-ratio property, both $\frac{\mu_i^{\theta'}(x_i)}{\mu_i^{\theta}(x_i)}$ and $\frac{\mu_i^{\theta''}(x_i)}{\mu_i^{\theta}(x_i)}$ are increasing in x_i. Moreover, $\mu_i^{\theta''}$ first-order stochastically dominates $\mu_i^{\theta'}$. Thus, for all $\overline{x}_i \in X_i$,

$$\begin{aligned} \sum_{x_i \in X_i \text{ s.t. } \overline{x}_i \geq_i x_i} \mu_i^{\theta}(x_i)\frac{\mu_i^{\theta'}(x_i)}{\mu_i^{\theta}(x_i)} = \sum_{x_i \in X_i \text{ s.t. } \overline{x}_i \geq_i x_i} \mu_i^{\theta'}(x_i) &\geq \sum_{x_i \in X_i \text{ s.t. } \overline{x}_i \geq_i x_i} \mu_i^{\theta''}(x_i) \\ &= \sum_{x_i \in X_i \text{ s.t. } \overline{x}_i \geq_i x_i} \mu_i^{\theta}(x_i)\frac{\mu_i^{\theta''}(x_i)}{\mu_i^{\theta}(x_i)}, \end{aligned}$$

with equality if $\overline{x}_i$ is $\geq_i$-maximal. This implies the desired mean-preserving spread relationship (e.g., Theorem 3.A.5. in Shaked and Shanthikumar, 2007). □

Proof of Corollary F.1. First, take any non-degenerate collection of partitions Π and any $\theta \in \Theta$, $i \in I$, and $\theta' \notin \Pi_i(\theta)$ such that $d(\mu_i^{\theta}, \mu_i^{\theta'}) = \lambda(\mathcal{I}, \Pi)$. Note that we can assume that $|\theta - \theta'| = 1$. Indeed, otherwise, Lemma F.1 yields a state θ'' in between θ and θ' such that $\max\{d(\mu_i^{\theta}, \mu_i^{\theta''}), d(\mu_i^{\theta''}, \mu_i^{\theta'})\} \leq d(\mu_i^{\theta}, \mu_i^{\theta'})$, where either $\theta'' \notin \Pi_i(\theta)$ or $\theta'' \notin \Pi_i(\theta')$ holds. The same argument applies to $\tilde{\mathcal{I}}$.

To show that (ii) implies (i), note that if $\min_{i\in I} d(\mu_i^{\theta}, \mu_i^{\theta+1}) > \min_{i\in I} d(\tilde{\mu}_i^{\theta}, \tilde{\mu}_i^{\theta+1})$ holds for all $\theta = 1, \ldots, n-1$, then $\lambda(\mathcal{I}, \Pi) > \lambda(\tilde{\mathcal{I}}, \Pi)$ holds for all partitions Π by the observation in the previous paragraph. Thus, the conclusion follows from Corollary 1.

To show that (i) implies (ii), we prove the contrapositive. Suppose that we have $\min_{i\in I} d(\mu_i^{\theta}, \mu_i^{\theta+1}) \leq \min_{i\in I} d(\tilde{\mu}_i^{\theta}, \tilde{\mu}_i^{\theta+1})$ for some θ, where the inequality must be strict by the assumption that $\min_{i\in I} d(\mu_i^{\theta}, \mu_i^{\theta+1}) \neq \min_{i\in I} d(\tilde{\mu}_i^{\theta}, \tilde{\mu}_i^{\theta+1})$. Then take any basic game $\mathcal{G}$ and objective function W such that, for all i, $\Pi_i^W(\theta') = \{1, \ldots, \theta\}$ for each $\theta' \leq \theta$ and $\Pi_i^W(\theta') = \{\theta + 1, \ldots, n\}$ for each $\theta' > \theta$. By the observation in the first paragraph, we have $\lambda(\mathcal{I}, \Pi^W) = \min_{i\in I} d(\mu_i^{\theta}, \mu_i^{\theta+1})$ and $\lambda(\tilde{\mathcal{I}}, \Pi^W) = \min_{i\in I} d(\tilde{\mu}_i^{\theta}, \tilde{\mu}_i^{\theta+1})$. Thus, by Theorem 3, there exists T such that $W_t(\tilde{\mathcal{I}}, \mathcal{G}) > W_t(\mathcal{I}, \mathcal{G})$ for all $t > T$. □

G Convergence of Equilibrium Sets

In Section 4, we focused on equilibria of $\mathcal{G}_t(\mathcal{I})$ that maximize the expected objective. In this section, we show that the learning efficiency index also captures how fast the whole equilibrium set of $\mathcal{G}_t(\mathcal{I})$ converges to the set of common knowledge equilibria.

Formally, given any basic game $\mathcal{G}$, $m \in \Delta(A)$ is an ***ε-correlated equilibrium*** at θ if, for each i,

$$m(a_i) > 0 \implies \sum_{a_{-i}} m(a_{-i}|a_i)\left(u_i(a_i, a_{-i}, \theta) - u_i(a_i', a_{-i}, \theta)\right) \geq -\varepsilon, \forall a_i' \in A_i.$$

Let $\mathrm{CE}^{\theta,\varepsilon}(\mathcal{G})$ denote the set of ε-correlated equilibria at θ, and $\mathrm{CE}^{\varepsilon}(\mathcal{G})$ the set of joint distributions over states and actions induced by ε-correlated equilibria at each state, i.e.,

$$\mathrm{CE}^{\varepsilon}(\mathcal{G}) := \{m \in \Delta(\Theta \times A) : m(\theta) = p_0(\theta), m(\cdot|\theta) \in \mathrm{CE}^{\theta,\varepsilon}(\mathcal{G}),\ \forall \theta \in \Theta\}.$$

We also denote by $\mathrm{NE}(\mathcal{G})$ the set of joint distributions over states and actions induced by Nash equilibria at each state, defined in the usual manner.

Define the set of ε-Bayes Nash equilibria of $\mathcal{G}_t(\mathcal{I})$ analogously. Finally, abusing notation relative to the main text, let $\mathrm{BNE}^{\varepsilon}_t(\mathcal{G}, \mathcal{I}) \subseteq \Delta(\Theta \times A)$ denote the set of joint distributions over states and actions induced by ε-Bayes Nash equilibria of $\mathcal{G}_t(\mathcal{I})$.

Corollary G.1. *Take any information structure $\mathcal{I}$ and any $\varepsilon > 0$. For any basic game $\mathcal{G}$, as $t \to \infty$,*

$$\sup_{m_t \in \mathrm{BNE}_t(\mathcal{G},\mathcal{I})} \inf_{m \in \mathrm{CE}^{\varepsilon}(\mathcal{G})} \|m_t - m\| \leq \exp[-t\lambda(\mathcal{I}) + o(t)], \tag{31}$$

$$\sup_{m \in \mathrm{NE}(\mathcal{G})} \inf_{m_t \in \mathrm{BNE}^{\varepsilon}_t(\mathcal{G},\mathcal{I})} \|m_t - m\| \leq \exp[-t\lambda(\mathcal{I}) + o(t)]. \tag{32}$$

Moreover, for some basic game $\mathcal{G}$, both inequalities hold with equality.

By (31), the ex-ante learning efficiency index $\lambda(\mathcal{I})$ lower-bounds the speed at which every BNE outcome at large t can be approximated by some ε-correlated equilibrium in the complete information limit. Note that we employ ε-correlated equilibria in the limit instead of ε-Nash equilibria; this is because, even though players achieve approximate common knowledge at large t, signal distributions in general introduce correlation into their action choices. By (32), $\lambda(\mathcal{I})$ also lower-bounds the speed at which every Nash equilibrium outcome in the complete information limit can be approximated by some ε-BNE outcome at large t. Finally, both bounds are tight.

Proof of Corollary G.1. For simplicity, we focus on the case where each joint distribution $\mu^{\theta} \in \Delta(X)$ has full support; the extension to general information structures $\mathcal{I}$ follows similar arguments as in Appendix D. Fix any $\varepsilon > 0$ and basic game $\mathcal{G}$.

Inequality (31): Pick $p \in (0, 1)$ large enough that

$$p\varepsilon \geq (1 - p) \max_{i,a_i,a_i',a_{-i},\theta} |u_i(a_i, a_{-i}, \theta) - u_i(a_i', a_{-i}, \theta)|.$$

Take any $d < \lambda(\mathcal{I})$. By Lemma A.2, there exists T such that for all $t \geq T$, $i \in I$, and $\theta \in \Theta$, whenever $\mathrm{KL}(\nu_{it}, \mu^{\theta}_i) \leq d$, then

$$\mathbb{P}^{\mathcal{I}}_t \left(\{\theta\} \cap F_t(\theta, d) \mid x^t_i\right) \geq p. \tag{33}$$

Take any $t \geq T$, BNE σ_t of $\mathcal{G}_t(\mathcal{I})$, $i \in I$, $\theta \in \Theta$, and $x^t_i \in X^t_i$ such that $\mathrm{KL}(\nu_{it}, \mu^{\theta}_i) \leq d$. Then for any a_i with $\sigma_{it}(a_i|x^t_i) > 0$, the fact that σ_t is a BNE implies that, for all

$a_i' \in A_i$,

$$\sum_{\theta' \in \Theta, x_{-i}^t \in X_{-i}^t} \mathbb{P}_t^{\mathcal{I}}(\theta', x_{-i}^t | x_i^t) \left(u_i(a_i, \sigma_{-i}(x_{-i}^t), \theta') - u_i(a_i', \sigma_{-i}(x_{-i}^t), \theta') \right) \geq 0,$$

which, by (33) and the choice of p, implies that

$$\sum_{\theta' \in \Theta, x_{-i}^t \in X_{-i}^t} \mathbb{P}_t^{\mathcal{I}}(\theta', x_{-i}^t | x_i^t, \{\theta\} \cap F_t(\theta, d)) \left(u_i(a_i, \sigma_{-i}(x_{-i}^t), \theta') - u_i(a_i', \sigma_{-i}(x_{-i}^t), \theta') \right) \geq -\varepsilon.$$

That is, for all $t \geq T$ and $\theta \in \Theta$, the action distribution induced by any BNE of $\mathcal{G}_t(\mathcal{I})$ conditional on the event $\{\theta\} \cap F_t(\theta, d)$ is an ε-correlated equilibrium at θ.

Thus, for all $t \geq T$,

$$\sup_{m_t \in \mathrm{BNE}_t(\mathcal{G}, \mathcal{I})} \inf_{m \in \mathrm{CE}^{\varepsilon}(\mathcal{G})} \|m_t - m\| \leq \max_{\theta \in \Theta} p_0(\theta) \left(1 - \mathbb{P}_t^{\mathcal{I}}(F_t(\theta, d) | \theta) \right).$$

By Sanov's theorem, this implies that, as $t \to \infty$,

$$\sup_{m_t \in \mathrm{BNE}_t(\mathcal{G}, \mathcal{I})} \inf_{m \in \mathrm{CE}^{\varepsilon}(\mathcal{G})} \|m_t - m\| \leq \exp[-td + o(t)].$$

Since this holds for any $d < \lambda(\mathcal{I})$, this proves inequality (31).

Inequality (32): Pick $p \in (0, 1)$ large enough that

$$\varepsilon \geq (1 - p) \max_{i, a_i, a_i', a_{-i}, \theta} |u_i(a_i, a_{-i}, \theta) - u_i(a_i', a_{-i}, \theta)|.$$

Take any $d < \lambda(\mathcal{I})$. By Lemma A.2, there exists T such that for all $t \geq T$, $i \in I$, and $\theta \in \Theta$, whenever $\mathrm{KL}(\nu_{it}, \mu_i^{\theta}) \leq d$, then (33) holds.

Take any $m \in \mathrm{NE}(\mathcal{G})$, and let $\alpha_i^{\theta} \in \Delta(A_i)$ denote the corresponding Nash equilibrium strategy of player i at θ. Let $\Sigma_{it}(d)$ denote the set of i's strategies σ_{it} in $\mathcal{G}_t(\mathcal{I})$ such that, for each θ, $\sigma_{it}(\cdot | x_i^t) = \alpha_i^{\theta}(\cdot)$ whenever $\mathrm{KL}(\nu_{it}, \mu_i^{\theta}) \leq d$. By Kakutani's fixed-point theorem applied to the best-response correspondences on the restricted strategy space $\prod_i \Sigma_{it}(d)$, there exists a strategy profile $\sigma_t \in \prod_i \Sigma_{it}(d)$ such that each player i's action conditional on a signal sequence x_i^t with $\mathrm{KL}(\nu_{it}, \mu_i^{\theta}) > d$ is interim optimal against σ_{-it}. Moreover, for $t \geq T$, (33) and the choice of p ensure that each player i's action conditional on a signal sequence x_i^t with $\mathrm{KL}(\nu_{it}, \mu_i^{\theta}) \leq d$ (i.e., the support of α_i^{θ}) is ε-interim optimal against σ_{-it}. Thus, σ_t is an ε-BNE of $\mathcal{G}_t(\mathcal{I})$.

Hence, for all $t \geq T$,

$$\sup_{m_t \in \mathrm{NE}(\mathcal{G}, \mathcal{I})} \inf_{m \in \mathrm{BNE}_t^{\varepsilon}(\mathcal{G}, \mathcal{I})} \|m_t - m\| \leq \max_{\theta \in \Theta} p_0(\theta) \left(1 - \mathbb{P}_t^{\mathcal{I}}(F_t(\theta, d) | \theta) \right).$$

By Sanov's theorem, this implies that, as $t \to \infty$,

$$\sup_{m_t \in \mathrm{NE}(\mathcal{G},\mathcal{I})} \inf_{m \in \mathrm{BNE}^{\varepsilon}_t(\mathcal{G},\mathcal{I})} \|m_t - m\| \leq \exp[-dt + o(t)].$$

Since this holds for any $d < \lambda(\mathcal{I})$, this proves inequality (32).

Equality for some $\mathcal{G}$: Take i and θ, θ' such that $d(\mu_i^{\theta}, \mu_i^{\theta'}) = \lambda(\mathcal{I})$. Then consider a basic game $\mathcal{G}$ such that $A_i = \{a_i, a_i'\}$ and, for all a_{-i},

$$u_i(a_i, a_{-i}, \theta) - u_i(a_i', a_{-i}, \theta) = 2\varepsilon = u_i(a_i', a_{-i}, \theta') - u_i(a_i, a_{-i}, \theta').$$

This implies that, for any $m \in \mathrm{CE}^{\varepsilon}(\mathcal{G}) \cup \mathrm{NE}(\mathcal{G})$, we have $m(a_i|\theta) = m(a_i'|\theta') = 1$. Thus,

$$\sup_{m_t \in \mathrm{BNE}_t(\mathcal{G},\mathcal{I})} \inf_{m \in \mathrm{CE}^{\varepsilon}(\mathcal{G})} \|m_t - m\| \geq \sup_{m_t \in \mathrm{BNE}_t(\mathcal{G},\mathcal{I})} \max\{p_0(\theta)(1 - m_t(a_i|\theta)), p_0(\theta')(1 - m_t(a_i'|\theta'))\},$$

$$\sup_{m \in \mathrm{NE}(\mathcal{G})} \inf_{m_t \in \mathrm{BNE}^{\varepsilon}_t(\mathcal{G},\mathcal{I})} \|m_t - m\| \geq \inf_{m_t \in \mathrm{BNE}^{\varepsilon}_t(\mathcal{G},\mathcal{I})} \max\{p_0(\theta)(1 - m_t(a_i|\theta)), p_0(\theta')(1 - m_t(a_i'|\theta'))\}.$$

For any sequence (m_t) of distributions induced by ε-BNE (or any strategy profiles more generally), the proof of Lemma C.1 adapted to the current notation shows that

$$\liminf_{t \to \infty} \frac{1}{t} \log\left(\max\{1 - m_t(a_i|\theta), 1 - m_t(a_i'|\theta')\}\right) \geq -d(\mu_i^{\theta}, \mu_i^{\theta'}).$$

Thus, as $t \to \infty$,

$$\sup_{m_t \in \mathrm{BNE}_t(\mathcal{G},\mathcal{I})} \inf_{m \in \mathrm{CE}^{\varepsilon}(\mathcal{G})} \|m_t - m\| \geq \exp[-td(\mu_i^{\theta}, \mu_i^{\theta'}) + o(t)],$$

$$\sup_{m \in \mathrm{NE}(\mathcal{G})} \inf_{m_t \in \mathrm{BNE}^{\varepsilon}_t(\mathcal{G},\mathcal{I})} \|m_t - m\| \geq \exp[-td(\mu_i^{\theta}, \mu_i^{\theta'}) + o(t)],$$

as claimed. □

H Gaussian Signals

We show that the speed of common learning also coincides with the speed of individual learning in the following infinite-signal Gaussian environment. For simplicity, consider two players $i = 1, 2$ and two states $\theta = \underline{\theta}, \overline{\theta}$; extending to more players/states is straightforward. Assume that conditional on state θ, signal profiles are drawn i.i.d. according to

$$(x_{1t}, x_{2t}) \sim \mathcal{N}\left((m_1^{\theta}, m_2^{\theta}), \Sigma\right), \quad \Sigma = \begin{pmatrix} (\sigma_1)^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & (\sigma_2)^2 \end{pmatrix}$$

for some $\rho \in (-1,1)$. Up to applying an affine transformation of signals, we can assume without loss of generality that $m_i^{\underline{\theta}} = 0$, $m_i^{\overline{\theta}} = 1$ for $i = 1,2$.

Consider $m_{it} := \sum_{s=1}^{t} \frac{x_{is}}{t}$, which is a sufficient statistic for player i's (higher-order) beliefs. Conditional on state θ, (m_{1t}, m_{2t}) is distributed Gaussian with mean (m_1^θ, m_2^θ) and covariance matrix $\frac{1}{t}\Sigma$. Moreover, by the law of large numbers, $m_{it} \to m_i^\theta$ almost surely conditional on state θ. For any sufficiently large t, if $m_{it} < \frac{1}{2}$ (resp. $m_{it} > \frac{1}{2}$), then i's belief concentrates on state $\underline{\theta}$ (resp. $\overline{\theta}$).

Fix any $p \in (0,1)$ and consider state $\overline{\theta}$; the argument in state $\underline{\theta}$ is analogous. To calculate the speed of individual learning in state $\overline{\theta}$, note that

$$\lim_{t\to\infty} -\frac{1}{t} \log\left(1 - \mathbb{P}_t[B_i^p(\overline{\theta}) \mid \overline{\theta}]\right) = \lim_{t\to\infty} -\frac{1}{t} \log \mathbb{P}_t\left[m_{it} < \frac{1}{2} \mid \overline{\theta}\right] = \frac{1}{8(\sigma_i)^2},$$

where the final equality holds by Cramér's theorem.[27] Thus, as $t \to \infty$,

$$\mathbb{P}_t[B_t^p(\overline{\theta}) \mid \overline{\theta}] = 1 - \exp\left[\frac{-1}{8\max_i(\sigma_i)^2} t + o(t)\right].$$

To calculate the speed of common learning in state $\overline{\theta}$, assume without loss that $\sigma_1 \leq \sigma_2$, i.e., player 1's rate of individual learning is faster. For each $d \in (0, 1/2)$, consider the event

$$F_t(d, \overline{\theta}) = \left\{|m_{1t} - 1| \leq d\frac{\sigma_1}{\sigma_2}\right\} \cap \left\{|m_{2t} - 1| \leq d\right\}.$$

Observe that $F_t(d, \overline{\theta}) \subseteq B_t^p(\overline{\theta})$ for all sufficiently large t. Next, we show that this event is p-evident. Indeed, note that for each i, we have

$$\left|\mathbb{E}[m_{-it} | m_{it}, \overline{\theta}] - 1\right| = |\rho| \frac{\sigma_{-i}}{\sigma_i} |m_{it} - 1|.$$

Thus, conditional on event $F_t(d, \overline{\theta})$, we have

$$\left|\mathbb{E}[m_{1t} | m_{2t}, \overline{\theta}] - 1\right| \leq |\rho| d \frac{\sigma_1}{\sigma_2}, \qquad \left|\mathbb{E}[m_{2t} | m_{1t}, \overline{\theta}] - 1\right| \leq |\rho| d.$$

Since i's estimate of m_{-it} given m_{it} and $\overline{\theta}$ becomes arbitrarily precise as t grows large (i.e., the conditional variance $\frac{1}{t}(1-\rho^2)\sigma_{-i}^2 \to 0$), this guarantees that event $F_t(d, \overline{\theta})$ is p-evident for all sufficiently large t. Hence, by Monderer and Samet (1989),

[27] Indeed, since m_{it} is the sample mean of i.i.d. draws from $\mathcal{N}(m_i^{\overline{\theta}}, (\sigma_i)^2)$, Cramér's theorem implies that $\lim_{t\to\infty} -\frac{1}{t} \log \mathbb{P}_t\left[m_{it} < \frac{1}{2} \mid \overline{\theta}\right] = I(\frac{1}{2})$, where $I(a) := \sup_\lambda (\lambda a - \log M(\lambda)) = \frac{(a-\mu_i^\theta)^2}{2(\sigma_i^\theta)^2}$ and $M(\lambda) = \exp[\lambda m_i^\theta + \frac{\lambda^2(\sigma_i)^2}{2}]$ is the moment generating function of $\mathcal{N}(m_i^\theta, (\sigma_i)^2)$.

$F_t(d, \overline{\theta}) \subseteq C_t^p(\overline{\theta})$ for all sufficiently large t. Thus, Cramér's theorem implies that

$$\begin{aligned} & \liminf_{t\to\infty} -\frac{1}{t} \log \left(1 - \mathbb{P}_t[C^p(\overline{\theta}) \mid \overline{\theta}]\right) \\ \geq \quad & \lim_{t\to\infty} -\frac{1}{t} \log \mathbb{P}_t \left[|m_{1t} - 1| > d\frac{\sigma_1}{\sigma_2} \text{ or } |m_{2t} - 1| > d \mid \overline{\theta} \right] = \frac{d^2}{2(\sigma_2)^2}. \end{aligned}$$

Since d can be chosen arbitrarily close to $\frac{1}{2}$, it follows that

$$\mathbb{P}_t[C_t^p(\overline{\theta}) \mid \overline{\theta}] = 1 - \exp\left[\frac{-1}{8\max_i(\sigma_i)^2} t + o(t)\right],$$

i.e., as $t \to \infty$, common learning and individual learning occur at the same exponential rate.