

Learning Efficiency of Multi-Agent Information Structures*

Mira Frick Ryota Iijima Yuhta Ishii

First posted version: August 17, 2021
This version: July 22, 2022

Abstract

We study which multi-agent information structures are more effective at eliminating both first-order and higher-order uncertainty, and hence at facilitating efficient play in incomplete-information coordination games. We consider a learning setting à la [Cripps, Ely, Mailath, and Samuelson \(2008\)](#) where players have access to many private signal draws from an information structure. First, we characterize the rate at which players achieve approximate common knowledge of the state, based on a simple learning efficiency index. Notably, this coincides with the rate at which players' first-order uncertainty vanishes, as higher-order uncertainty becomes negligible relative to first-order uncertainty after enough signal draws. Based on this, we show that information structures with higher learning efficiency induce more efficient equilibrium outcomes in coordination games that are played after sufficiently many signal draws. We highlight some robust implications for information design in games played in data-rich environments.

Keywords: higher-order beliefs, common learning, coordination, speed of learning, 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, Michihiro Kandori, Andreas Kleiner, Vijay Krishna, Giacomo Lanzani, Bart Lipman, George Mailath, Stephen Morris, Daisuke Oyama, Antonio Penta, Larry Samuelson, Eddie Schlee, Ron Siegel, Jakub Steiner, Colin Stewart, Philipp Strack, Nick Wu, and Muhamet Yildiz.

1 Introduction

Coordination problems under uncertainty about a payoff-relevant state of the world are ubiquitous in economics, from joint investment decisions and technology adoption to currency attacks, bank runs, and political revolutions. In such settings, there are two obstacles to coordinating on an efficient outcome: players' first-order uncertainty about the state and their higher-order uncertainty about other players' beliefs about the state. Thus, an important question is to understand which information structures are more effective at reducing both forms of uncertainty, and hence at facilitating coordination.

In this paper, we address this question by considering a learning setting, where players have access to many draws of private signals from an information structure (capturing, for instance, that data is “cheap” or abundant). Our starting point is a classic result due to [Cripps, Ely, Mailath, and Samuelson \(2008\)](#), henceforth CEMS, which shows that (under natural conditions) this setting leads to *common learning*: Under any information structure, players achieve approximate common knowledge ([Monderer and Samet, 1989](#)) of the true state as the number of signal draws goes to infinity. Thus, asymptotically, all information structures eliminate both first-order and higher-order uncertainty, but the result is silent about which information structures do so more effectively. To understand the latter, a natural approach is to compare which information structures lead to *faster* common learning, i.e., are more likely to induce approximate common knowledge of the state away from the limit, after any large but finite number of signal draws.

Our first main result conducts such a comparison, by characterizing the *speed of common learning* under each information structure. Our key insight is that all that matters is how fast an information structure eliminates first-order uncertainty: We show that the speed of common learning simply coincides with the speed at which all players individually learn the state, because, under every information structure, higher-order uncertainty vanishes faster than first-order uncertainty. This allows us to characterize the speed of common learning using a simple *multi-agent learning efficiency* index. The index depends only on the statistical informativeness ([Chernoff, 1952](#); [Moscarini and Smith, 2002](#)) of the worst-informed player's private signals; in contrast, the correlation across different players' private signals is irrelevant.

Second, we apply this result to rank information structures in terms of their value in coordination problems. In particular, we show that, for a rich class of games and objectives that are “aligned at certainty,” information structures with higher learning

efficiency lead to better equilibrium outcomes whenever players have access to sufficiently many signal draws. Based on the structure of the learning efficiency index, this result yields some robust implications for information design in coordination games that are played in data-rich settings.

Section 2 introduces the learning setting. An information structure \mathcal{I} specifies a joint distribution over players' private signals in each state, where both states and signals are assumed finite. We consider a setting where players receive t independent draws of private signals from \mathcal{I} , but \mathcal{I} may feature arbitrary correlation across different players' private signals.

Section 3 characterizes the speed of common learning: For each information structure \mathcal{I} , we consider the probability that players have common p -belief (for p arbitrarily close to 1) of the true state after t signal draws from \mathcal{I} , and analyze how fast this converges to one as t grows large. Common p -belief is a much more demanding notion than individual knowledge, as it imposes confidence not only on players' first-order beliefs about the state, but on their infinite hierarchy of higher-order beliefs. However, perhaps surprisingly, Theorem 1 shows that the probability of common p -belief converges to one at the *same* exponential rate at which all players individually learn the state, which is characterized by the aforementioned learning efficiency index. The proof of Theorem 1 relies on a key information-theoretic lemma that uses Kullback-Leibler divergence to formalize that players' higher-order uncertainty vanishes faster than their first-order uncertainty (Lemma 1).

Section 4 augments the learning setting by assuming that, once players have observed many signal draws from an information structure, they face an incomplete-information game. With each game, we associate an objective function over action profiles in each state, capturing, for instance, players' welfare or a designer's preferences. Theorem 2 provides a large-sample ranking over information structures: We identify a class of games and objectives for which information structures with a higher learning efficiency index induce better (Bayes-Nash) equilibrium outcomes whenever players observe sufficiently many signal draws. This class satisfies one substantive assumption, *alignment at certainty*: We require 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. A leading instance of this assumption is when the objective is to maximize utilitarian welfare and the game is a coordination problem, such as the illustrative joint investment example below, coordinated attack games (Example 2), and other important examples in the literature. As we will see, the fact that

the ranking in Theorem 2 applies uniformly to all these environments relies crucially on our finding in Theorem 1 that the speed of common learning coincides with the speed of individual learning.

By focusing on settings where players have access to rich data, our analysis yields some insights into information design in coordination problems that apply robustly, regardless of the specific game being played. First, a designer seeking to facilitate coordination should focus on improving players' information about the state; in contrast, the effect of providing additional signals about other players' signals (that are not directly informative about the state) is negligible. Second, the designer should be "egalitarian," i.e., focus on improving the worst-informed player's information about the state.

Example 1 (Illustrative example: Joint investment). Consider two players, $i = 1, 2$, with symmetric action sets $A_i = \{0, 1\}$. Action 1 represents investment and action 0 no investment. The state $\theta \in \{\underline{\theta}, \bar{\theta}\}$ captures whether the market fundamental is low ($\underline{\theta}$) or high ($\bar{\theta}$) and is drawn according to some non-degenerate prior p_0 . Each player i 's utility takes the form

$$u_i(a, \theta) = \begin{cases} \mathbf{1}_{\{\theta=\bar{\theta}\}} \mathbf{1}_{\{a_{-i}=1\}} - c & \text{if } a_i = 1 \\ 0 & \text{if } a_i = 0. \end{cases}$$

That is, if i invests, she incurs a cost of $c \in (0, 1)$, and the investment is successful (payoff of 1) if and only if the state is $\bar{\theta}$ and her opponent also invests. The payoff to unsuccessful or no investment is 0. Under utilitarian welfare, $\frac{1}{2} (u_1(a, \theta) + u_2(a, \theta))$, the efficient outcome is to play (1, 1) in state $\bar{\theta}$ and (0, 0) in state $\underline{\theta}$. These are strict Nash equilibria under common knowledge of θ , but incomplete information prevents the efficient outcome from being an equilibrium.

Now suppose that, prior to choosing actions, players learn about state θ from repeated signal draws. Our analysis yields a (generically) complete ranking over information structures: Using our learning efficiency index, one can compare how fast players achieve approximate common knowledge of θ under different information structures, 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: Each player i 's private signal realizations x_i are either $\underline{\theta}$ or $\bar{\theta}$, and the joint probabilities of players' signals in state θ are given by

Here, the individual precision parameter $\gamma \in (1/2, 1)$ captures the probability with

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

which each player’s signal matches the state, and the parameter $\rho \in [0, 1]$ captures the extent of correlation across players’ private signals. Higher values of γ help to reduce players’ first-order uncertainty about the state, while ρ influences players’ predictions of their opponent’s signals, i.e., their higher-order uncertainty. Thus, in comparing two information structures parametrized by (γ, ρ) and $(\tilde{\gamma}, \tilde{\rho})$, it might not be obvious how to trade off these two considerations. Indeed, if players observe only a small number of signal draws, whether (γ, ρ) or $(\tilde{\gamma}, \tilde{\rho})$ induces better equilibrium play can vary across different priors p_0 and investment costs c .

However, we will show that our learning efficiency index depends only on γ . Thus, for *any* p_0 and c , higher levels of individual precision γ allow for more efficient equilibrium play when players observe sufficiently many signal draws; in contrast, the effect of correlation ρ becomes negligible as the number of signals grows large. This reflects our key insight that the speed of common learning is the same as the speed of individual learning, because higher-order uncertainty about opponents’ signals vanishes faster than first-order uncertainty about the state. \blacktriangle

1.1 Related Literature

Our paper contributes to the large literature on higher-order beliefs (e.g., Rubinstein, 1989; Carlsson and Van Damme, 1993; Kajii and Morris, 1997; Morris and Shin, 1998; Weinstein and Yildiz, 2007). A central insight in this literature is that higher-order uncertainty about a payoff-relevant state can be an important source of inefficiency in coordination games. This reflects the fact that, even when all players’ first-order uncertainty is small, higher-order uncertainty can be significant. In contrast, we highlight that, in natural learning settings where players have access to rich enough data about the state, higher-order uncertainty vanishes faster than first-order uncertainty and eventually becomes negligible relative to first-order uncertainty.

To make this point, we consider the same learning setting as CEMS.¹ As mentioned, our contribution relative to their paper is to provide a *comparison* of different

¹Other 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.

information structures based on the speed at which they induce approximate common knowledge, and to use this to rank information structures in terms of their value in coordination games. Our proof of Theorem 1 builds on CEMS’ proof approach, but as Section 3.2 illustrates, we refine their analysis by introducing information-theoretic arguments that are crucial for deriving the *rate* of common learning. We obtain a complete ranking over any two information structures whose learning efficiency indices are not equal; more recently, [Awaya and Krishna \(2022\)](#) study the natural complementary case where information structures have common marginal signal distributions (and hence equal learning efficiency indices) but differ in their correlation structure (see the discussion in Section 3.1).

[Moscarini and Smith \(2002\)](#) derive an efficiency index that characterizes the speed of single-agent learning. Our learning efficiency index generalizes theirs to multi-agent settings. Our key finding is that, because higher-order uncertainty vanishes faster than first-order uncertainty, the multi-agent index simply reduces to the slowest agent’s individual learning efficiency index and does not depend on the correlation across different agents’ signals. 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.² 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.

More broadly, we relate to the literature on information design in games (for surveys, see [Bergemann and Morris, 2019](#); [Kamenica, 2019](#)). In contrast to the typical approach in this literature, we assume that players observe many i.i.d. draws from the chosen information structure and we rule out information structures that fully reveal the state. Our ranking over information structures has robust design implications that apply to all games and objectives satisfying alignment at certainty. As is also common in this literature, our analysis assumes a designer-preferred equilibrium selection; see, e.g., [Morris, Oyama, and Takahashi \(2020\)](#) for an alternative approach that considers adversarial equilibrium selection.

²See, 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 \(2019\)](#).

Finally, our exercise relates to the literature on comparisons of information structures. Blackwell (1951) compares information structures in terms of their induced payoffs in all single-agent decision problems. While Blackwell’s order assumes that the agent observes a single signal draw, Moscarini and Smith’s (2002) aforementioned efficiency index extends this order to single-agent settings with many i.i.d. signal draws.³ Extensions of Blackwell’s order to multi-player games have focused on the single signal draw case. Because more information can be harmful in some games (e.g., Hirshleifer, 1971), one needs to restrict the class of games and objectives to avoid obtaining a highly conservative ranking.⁴ In particular, Lehrer, Rosenberg, and Shmaya (2010) focus on common-interest games with utilitarian welfare, while Peški (2008) compares min-max values in zero-sum games.⁵ As we discuss in Remark 1, 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 objectives beyond the common-interest case.

2 Setting

Learning environment. Throughout the paper, we fix a finite set of agents N , 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 $i \in N$, with corresponding set of signal profiles $X := \prod_{i \in N} X_i$, 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.

We consider a setting where agents observe repeated i.i.d. signal draws from an information structure. 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

³Azrieli (2014) and Mu, Pomatto, Strack, and Tamuz (2021) consider a more demanding order that requires the number of signal draws to be uniform across decision problems.

⁴Indeed, Gossner (2000) compares Bayes-Nash equilibrium outcomes for general games and objectives and shows that no two information structures that induce different (higher-order) beliefs can be compared.

⁵Bergemann and Morris (2016) study general games using Bayes-correlated equilibria, which are equivalent to Bayes-Nash equilibria in a setting with a mediator who observes the state and signals. Brooks, Frankel, and Kamenica (2021) compare the informativeness of different agents’ signals within a multi-agent information structure.

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 . Agent i 's observed sequence of private signals is $x_i^t = (x_{i\tau})_{\tau=1,\dots,t}$.

Common learning. CEMS's classic result is that, in this setting, agents *commonly learn* the state, i.e., both their first-order uncertainty about θ and their higher-order uncertainty about other agents' beliefs about θ vanishes as t grows large.

Formally, 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} . That is,

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

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

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

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, 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 . At $C_t^p(E)$, the event E is p -believed, the event $B_t^p(E)$ is p -believed, and so on. The event $C_t^p(\theta)$ for p close to 1 captures that agents have approximate common knowledge of state θ (Monderer and Samet, 1989). *Common learning* requires that the true state is eventually commonly p -believed for p arbitrarily close to 1; that is, for all $p \in (0, 1)$ and $\theta \in \Theta$,

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

CEMS show that when states and signals are finite, as in the current setting, then every information structure \mathcal{I} gives rise to common learning.⁶

⁶See Section 6 for a discussion of more general settings.

3 Multi-Agent Learning Efficiency

3.1 Speed of Common Learning

While CEMS' result shows that, asymptotically, all information structures lead to approximate common knowledge of the state, it says nothing about which information structures do so more effectively. To capture this, a natural approach is to compare how different information structures \mathcal{I} affect the probability $\mathbb{P}_t^{\mathcal{I}}(C_t^p(\theta) \mid \theta)$ of approximate common knowledge at all large but finite t , i.e., to analyze the *rate* of convergence in (1). Our first main result provides a simple characterization of this rate, allowing us to rank information structures in terms of their *learning efficiency*.

We first recall a standard statistical measure that characterizes a *single* agent's rate of individual learning under each information structure \mathcal{I} . Fix any agent i and true state θ . Then, for any state $\theta' \neq \theta$, one can measure how difficult i finds it to statistically distinguish θ' from θ using the **Chernoff distance** (e.g., Cover and Thomas, 1999) between i 's marginal signal distributions in states θ and θ' :

$$d(\mu_i^\theta, \mu_i^{\theta'}) := \min_{\nu_i \in \Delta(X_i)} \max \left\{ \text{KL}(\nu_i, \mu_i^\theta), \text{KL}(\nu_i, \mu_i^{\theta'}) \right\}. \quad (2)$$

Here, $\text{KL}(\nu_i, \mu_i^\theta)$ denotes the Kullback-Leibler (henceforth, KL) divergence of ν_i relative to μ_i^θ .⁷ Observe that any minimizer ν_i of (2) must satisfy $\text{KL}(\nu_i, \mu_i^\theta) = \text{KL}(\nu_i, \mu_i^{\theta'})$. Thus, $d(\mu_i^\theta, \mu_i^{\theta'})$ is the distance from μ_i^θ and $\mu_i^{\theta'}$ to their KL-midpoint, so smaller values of $d(\mu_i^\theta, \mu_i^{\theta'})$ capture that i 's private signal distributions in states θ and θ' are closer to each other. Note that (unlike KL-divergence) the Chernoff distance is symmetric, and that $d(\mu_i^\theta, \mu_i^{\theta'}) > 0$ by the assumption that $\mu_i^\theta \neq \mu_i^{\theta'}$.

Statistical arguments (e.g., Chernoff, 1952) yield the following characterization of i 's speed of individual learning in state θ : For any $p \in (0, 1)$, as $t \rightarrow \infty$, the probability that i achieves individual p -belief of state θ goes to 1 exponentially,

$$\mathbb{P}_t^{\mathcal{I}}(B_{it}^p(\theta) \mid \theta) = 1 - \exp[-\lambda_i^\theta(\mathcal{I})t + o(t)], \quad (3)$$

where the rate of convergence is given by

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

⁷That is, $\text{KL}(\nu_i, \mu_i^\theta) := \sum_{x_i \in X_i} \nu_i(x_i) \log \frac{\nu_i(x_i)}{\mu_i^\theta(x_i)}$. By convention, $0 \log 0 = \frac{0}{0} = 0$ and $\log \frac{1}{0} = \infty$.

Thus, i 's individual learning efficiency under information structure \mathcal{I} is captured by a simple index $\lambda_i^\theta(\mathcal{I})$ that measures how difficult i finds it to distinguish state θ from the state θ' that generates the most similar private signal distribution. Building on this, Moscarini and Smith (2002) show that $\lambda_i^\theta(\mathcal{I})$ quantifies the value of information in single-agent decision problems under large samples of signals and prove that this index extends Blackwell's order:⁸ If i 's marginal signal distributions under \mathcal{I} Blackwell-dominate those under $\tilde{\mathcal{I}}$, then $\lambda_i^\theta(\mathcal{I}) \geq \lambda_i^\theta(\tilde{\mathcal{I}})$ for all θ .

Our first main result is that the rate at which agents commonly learn state θ is given by the *multi-agent learning efficiency* index

$$\lambda^\theta(\mathcal{I}) := \min_{i \in N} \lambda_i^\theta(\mathcal{I}) = \min_{i \in N, \theta' \in \Theta \setminus \{\theta\}} d(\mu_i^\theta, \mu_i^{\theta'}), \quad (4)$$

which simply considers the slowest agent's rate of individual learning.

Theorem 1. *Fix any information structure \mathcal{I} , state $\theta \in \Theta$, and $p \in (0, 1)$. Then individual learning and common learning both occur at rate $\lambda^\theta(\mathcal{I})$, i.e.,*

$$\mathbb{P}_t^{\mathcal{I}}(B_t^p(\theta) \mid \theta) = 1 - \exp[-\lambda^\theta(\mathcal{I})t + o(t)]; \quad (5)$$

$$\mathbb{P}_t^{\mathcal{I}}(C_t^p(\theta) \mid \theta) = 1 - \exp[-\lambda^\theta(\mathcal{I})t + o(t)]. \quad (6)$$

The fact that $\lambda^\theta(\mathcal{I})$ characterizes the rate of individual learning is immediate from (3): Since single-agent learning is exponential, the rate at which all agents achieve p -belief of the true state is determined by the slowest agent's rate of learning.

The substantive part of Theorem 1 is the characterization of the speed of common learning. As highlighted by a rich literature (see Section 1.1), 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.⁹ Based on this, it might be natural to expect common learning to occur more slowly than individual learning. However, Theorem 1 shows that, as $t \rightarrow \infty$, 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 $\lambda^\theta(\mathcal{I})$.¹⁰ As we illustrate

⁸More precisely, they use the index $\min_{\theta' \in \Theta \setminus \{\theta\}} \max_{\kappa \in [0, 1]} -\log \sum_{x_i \in X_i} \mu_i^\theta(x_i)^\kappa \mu_i^{\theta'}(x_i)^{1-\kappa}$, which is equal to $\lambda_i^\theta(\mathcal{I})$ by the variational formula (e.g., Dupuis and Ellis, 2011, Lemma 6.2.3.f).

⁹Relatedly, Kajii and Morris's (1997) critical path theorem yields a lower bound on the probability of $C_t^p(\theta)$ relative to the probability of $B_t^p(\theta)$, but this result only applies when p is small ($p < \frac{1}{|N|}$), reflecting a significant gap between common p -belief and individual p -belief when p is close to 1.

¹⁰The $o(t)$ terms can differ across (5) and (6) and can depend on p_0 , p , and features of \mathcal{I} other than $\lambda^\theta(\mathcal{I})$, but these terms become negligible as $t \rightarrow \infty$.

in Section 3.2, the key insight behind this result is that, as the number of signal draws grows large, agents’ higher-order uncertainty about others’ beliefs vanishes faster than their first-order uncertainty about the state.

The latter insight is also reflected by the structure of the multi-agent learning efficiency index: $\lambda^\theta(\mathcal{I})$ reduces each information structure \mathcal{I} to a simple one-dimensional measure that only focuses on the worst-informed agent i and the state θ' that i finds most difficult to distinguish from the true state θ based on her private signals; in contrast, the correlation across agents’ signals plays no role. For instance, in the illustrative Example 1, where \mathcal{I} is summarized by an individual precision parameter γ and a correlation parameter ρ , we have $\lambda^\theta(\mathcal{I}) = \text{KL}((\frac{1}{2}, \frac{1}{2}), (\gamma, 1 - \gamma))$; this is strictly increasing in γ but does not depend on ρ . When agents observe a small sample of signals, the probability of common p -belief in general depends on various other features of an information structure, including the correlation across agents’ signals. However, Theorem 1 implies that, under sufficiently large samples of signals, these features become irrelevant and λ^θ is all that is needed to compare the probabilities of common p -belief across different information structures:

Corollary 1. *Take any information structures $\mathcal{I}, \tilde{\mathcal{I}}$ and state $\theta \in \Theta$ such that $\lambda^\theta(\mathcal{I}) > \lambda^\theta(\tilde{\mathcal{I}})$. Then, for each $p \in (0, 1)$, there is T such that, for all $t \geq T$,*

$$\mathbb{P}_t^{\mathcal{I}}(C_t^p(\theta) \mid \theta) > \mathbb{P}_t^{\tilde{\mathcal{I}}}(C_t^p(\theta) \mid \theta).$$

Corollary 1 ranks any two information structures whose learning efficiency indices are not exactly tied, which holds for generic pairs of information structures.¹¹ One natural setting this excludes is when \mathcal{I} and $\tilde{\mathcal{I}}$ feature the same marginal signal distributions and differ only in their correlation. Complementary to Corollary 1, [Awaya and Krishna \(2022\)](#) study such settings and show that here higher correlation across agents’ signals can reduce the probability of common p -belief at all large enough t .

3.2 Illustration of Theorem 1

We prove Theorem 1 in Appendices B and E. To illustrate the key insight behind the result, consider the binary information structure from Example 1 with $\gamma = 3/5$ and $\rho = 5/12$.¹² Thus, the signal probabilities conditional on each state θ are:

¹¹Given \mathcal{I} , the set of information structures $\tilde{\mathcal{I}}$ such that $\lambda^\theta(\mathcal{I}) \neq \lambda^\theta(\tilde{\mathcal{I}})$ holds for all θ is open and dense in $\Delta(X)^\Theta$ endowed with the Euclidean topology.

¹²With these parameter values, agents’ signals are negatively correlated conditional on each state, but this feature is not important for our general arguments.

	$x_1 = \theta$	$x_1 \neq \theta$
$x_2 = \theta$	0.25	0.35
$x_2 \neq \theta$	0.35	0.05

Fix any $p \in (0, 1)$. Let $\nu_{it} \in \Delta(X_i)$ denote the empirical distribution of agent i 's signals up to t , and observe that this is a sufficient statistic for i 's (first-order and higher-order) beliefs. Hence, the events $B_t^p(\theta)$ and $C_t^p(\theta)$ can be described as subsets of $\Delta(X_1) \times \Delta(X_2)$. In particular, as depicted in Figure 1 (left), for all large enough t , one can show that $B_t^p(\theta)$ and $C_t^p(\theta)$ are approximated by

$$B_t^p(\theta) \approx \left\{ \nu_{it}(\theta) \in \left(\frac{1}{2}, 1 \right], \forall i = 1, 2 \right\}, \quad C_t^p(\theta) \approx \left\{ \nu_{it}(\theta) \in \left(\frac{1}{2}, \frac{9}{11} \right), \forall i = 1, 2 \right\}. \quad (7)$$

The expression for $B_t^p(\theta)$ is intuitive: at large t , i becomes confident in state θ as long as the majority of i 's signals matches θ . To see the idea behind $C_t^p(\theta)$, note that for any realized signal frequency $\nu_{it}(\theta) = \alpha \in (\frac{1}{2}, 1]$ of agent i and all large enough t , i assigns high probability to j 's realized signal frequency $\nu_{jt}(\theta)$ being approximately¹³

$$\mathbb{E}[\nu_{jt}(\theta) \mid \theta, \nu_{it}(\theta) = \alpha] = \alpha \frac{0.25}{0.6} + (1 - \alpha) \frac{0.35}{0.4}. \quad (8)$$

Observe that (8) exceeds $\frac{1}{2}$ only if $\alpha < \frac{9}{11}$. Thus, for i to be confident both in state θ and in the fact that j is confident in state θ , we need $\nu_{it}(\theta) \in (\frac{1}{2}, \frac{9}{11})$. Conversely, if $\nu_{it}(\theta) \in (\frac{1}{2}, \frac{9}{11})$, then (8) is itself in $(\frac{1}{2}, \frac{9}{11})$. This yields the approximation for $C_t^p(\theta)$.¹⁴

To consider the rate of common learning, assume that the true state is $\bar{\theta}$. By (7), at large t , the event $(C_t^p(\bar{\theta}))^c$ that common p -belief of $\bar{\theta}$ fails can be decomposed into two types of failures:

1. *First-order belief failures:* $(B_t^p(\bar{\theta}))^c \approx \{ \nu_{it}(\bar{\theta}) \leq \frac{1}{2} \text{ for some } i \}$.
2. *Higher-order belief failures:* $B_t^p(\bar{\theta}) \setminus C_t^p(\bar{\theta}) \approx \{ \nu_{it}(\bar{\theta}) \geq \frac{1}{2} \forall i, \nu_{it}(\bar{\theta}) \geq \frac{9}{11} \text{ for some } i \}$.

Reflecting that common p -belief is more demanding than individual p -belief, the second event, $B_t^p(\bar{\theta}) \setminus C_t^p(\bar{\theta})$, remains bounded away from the empty set even as $t \rightarrow \infty$. However, the key insight behind Theorem 1 is that, as t grows large, the probability

¹³This holds because i becomes confident in θ , so i 's beliefs about $\nu_{jt}(\theta)$ concentrate on the expectation $\mathbb{E}[\nu_{jt}(\theta) \mid \theta, \nu_{it}(\theta) = \alpha]$ by a law of large numbers argument.

¹⁴More precisely, based on these observations, one can show that, for all large enough t , there is an event $F_t \approx \{ \nu_{it}(\theta) \in (\frac{1}{2}, \frac{9}{11}], \forall i = 1, 2 \}$ such that $F_t \subseteq B_t^p(\theta)$ and $F_t \subseteq C_t^p(F_t)$ (i.e., F_t is p -evident), which by [Monderer and Samet \(1989\)](#) implies that $F_t \subseteq C_t^p(\theta)$.




Figure 1: *Left: Approximation of $B_t^p(\bar{\theta})$ and $C_t^p(\bar{\theta})$ at large t . Right: Rate of decay of higher-order belief failures (KL-distance of dashed arrows) and first-order belief failures (KL-distance of solid arrows) in state $\bar{\theta}$.*

of higher-order belief failures vanishes much faster than the probability of first-order belief failures, and hence becomes negligible for the rate of common learning.

Formally, we invoke Sanov's theorem from large deviation theory. Letting $\nu_t \in \Delta(X)$ denote the joint empirical distribution of agents' signals, this states that, for any set $D \subseteq \Delta(X)$ that is the closure of its interior,

$$\mathbb{P}_t^{\mathcal{I}}(\nu_t \in D \mid \bar{\theta}) = \exp[-\inf_{\nu \in D} \text{KL}(\nu, \mu^{\bar{\theta}})t + o(t)].$$

That is, as t grows large, the probability of event D vanishes exponentially at rate given by the KL-distance between D and the theoretical signal distribution $\mu^{\bar{\theta}}$. In the current setting, this implies that the probability $\mathbb{P}_t^{\mathcal{I}}(B_t^p(\bar{\theta}) \setminus C_t^p(\bar{\theta}) \mid \bar{\theta})$ of higher-order belief failures vanishes at rate

$$\text{KL}\left(\left(\frac{9}{11}, \frac{2}{11}\right), \mu_i^{\bar{\theta}}\right) = \text{KL}\left(\left(\frac{9}{11}, \frac{2}{11}\right), \left(\frac{3}{5}, \frac{2}{5}\right)\right),$$

as illustrated by either of the dashed distances in Figure 1 (right).¹⁵ In contrast, as the solid distances illustrate, the probability $\mathbb{P}_t^{\mathcal{I}}((B_t^p(\bar{\theta}))^c \mid \bar{\theta})$ of first-order belief failures

¹⁵The $\nu \in \Delta(X)$ that attains the infimum $\inf_{\nu \in B_t^p(\bar{\theta}) \setminus C_t^p(\bar{\theta})} \text{KL}(\nu, \mu^{\bar{\theta}})$ satisfies $\text{marg}_{X_i} \nu = (\frac{9}{11}, \frac{2}{11})$ and $\nu(\cdot | x_i) = \mu^{\bar{\theta}}(\cdot | x_i)$ for each x_i , so $\text{KL}(\nu, \mu^{\bar{\theta}})$ depends only on i 's marginal distributions (by the chain rule for KL-divergence). The arrows in Figure 1 should be interpreted as depicting KL-distances in the space of marginal distributions.

vanishes at rate

$$\text{KL} \left(\left(\frac{1}{2}, \frac{1}{2} \right), \mu_j^{\bar{\theta}} \right) = \text{KL} \left(\left(\frac{1}{2}, \frac{1}{2} \right), \left(\frac{3}{5}, \frac{2}{5} \right) \right) = \lambda^{\bar{\theta}}(\mathcal{I}).$$

Crucially, the latter rate is strictly smaller than the former. Thus, as t grows large, the ratio of $\mathbb{P}_t^{\mathcal{I}}((B_t^p(\bar{\theta}))^c \mid \bar{\theta})$ to $\mathbb{P}_t^{\mathcal{I}}(B_t^p(\bar{\theta}) \setminus C_t^p(\bar{\theta}) \mid \bar{\theta})$ explodes. Hence, higher-order belief failures become negligible relative to first-order belief failures, and the rate of common learning coincides with the rate of individual learning $\lambda^{\bar{\theta}}(\mathcal{I})$.

Finally, to see the more general idea, note that, by (8), $(\frac{1}{2}, \frac{1}{2}) = \mathbb{E}[\nu_{jt} \mid \bar{\theta}, \nu_{it} = (\frac{9}{11}, \frac{2}{11})]$. Thus, the inequality $\text{KL} \left((\frac{1}{2}, \frac{1}{2}), \mu_j^{\bar{\theta}} \right) < \text{KL} \left((\frac{9}{11}, \frac{2}{11}), \mu_i^{\bar{\theta}} \right)$ is an instance of the following general result that plays a crucial role in the proof of Theorem 1:

Lemma 1. *Fix any $\theta \in \Theta$ and distinct $i, j \in N$. For each t and realized empirical signal distribution $\nu_{it} \in \Delta(X_i)$, we have*

$$\text{KL}(\mathbb{E}[\nu_{jt} \mid \theta, \nu_{it}], \mu_j^{\theta}) \leq \text{KL}(\nu_{it}, \mu_i^{\theta}). \quad (9)$$

Moreover, the inequality is strict whenever μ^{θ} has full support and $\nu_i \neq \mu_i^{\theta}$.

In the appendix, we derive Lemma 1 from the chain rule for KL-divergence, a central result in information theory. To interpret (9), note that the RHS captures how much i 's signal observations ν_{it} deviate from i 's theoretical signal distribution μ_i^{θ} in state θ , while the LHS quantifies how much i 's expectation of j 's observations deviates from j 's theoretical signal distribution μ_j^{θ} . Thus, (9) 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 is less “atypical” than i 's own signal observations.¹⁶ Generalizing the above illustration, this can be used to show that, as t grows large, the event that agents learn θ but believe other agents to have incorrect first-order beliefs vanishes faster than the event that agents have incorrect first-order beliefs.

The inequality in Lemma 1 is reminiscent of the “contraction principle” in CEMS, whereby the map $\nu_i \mapsto \mathbb{E}[\mathbb{E}[\nu_{it} \mid \theta, \nu_{jt}] \mid \theta, \nu_{it} = \nu_i]$ is an L^1 -norm contraction on $\Delta(X_i)$ if μ^{θ} has full support (see their Lemma 4). This contraction principle can be used to show that the probability of higher-order belief failures vanishes as $t \rightarrow \infty$, and hence that common learning obtains, but it does not deliver the *rate* at which higher-order

¹⁶For 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., the LHS of (9) is 0). If i and j 's signals are perfectly correlated, then i expects j to observe the same signals as herself, so (9) holds with equality.

belief failures vanish. A key difference of our information-theoretic Lemma 1 is its use of KL-divergence. This is essential for being able to apply large deviation theory (Sanov’s theorem) to obtain this rate, and yields the new insight that common learning occurs just as fast as individual learning.

4 Ranking Information Structures in Coordination Problems

We now return to the question which information structures are more valuable for coordination. For this, we consider incomplete-information games that are played after a large number of signal draws, and we apply Theorem 1 to rank information structures in terms of the induced equilibrium outcomes.

4.1 Games and Objective Functions

A *basic game* \mathcal{G} consists of a finite set of actions A_i for each agent i , with corresponding set of action profiles $A := \prod_{i \in N} A_i$, as well as a utility function $u_i : A \times \Theta \rightarrow \mathbb{R}$ over action profiles and states for each agent i . For each basic game \mathcal{G} and information structure \mathcal{I} , we consider the (static) *incomplete-information game* $\mathcal{G}_t(\mathcal{I})$, where agents’ information is parametrized by t draws of signals from \mathcal{I} . That is, states θ and signal sequences x^t are drawn according to $\mathbb{P}_t^{\mathcal{I}}$, and a strategy $\sigma_{it} : (X_i)^t \rightarrow \Delta(A_i)$ for agent i maps i ’s observed sequence of private signals x_i^t to a mixed action in A_i . Let $\text{BNE}_t(\mathcal{G}, \mathcal{I})$ denote the set of Bayes-Nash equilibria (BNE) of $\mathcal{G}_t(\mathcal{I})$.¹⁷

To compare equilibrium outcomes across different information structures, we associate with any basic game \mathcal{G} an *objective function* $W : A \times \Theta \rightarrow \mathbb{R}$. This 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, $W = \frac{1}{N} \sum_{i \in N} u_i$. However, we also allow for objective functions that do not relate to agents’ utilities in any particular way. We assume that in each state θ , W is maximized by a unique action profile, $\{a^{\theta, W}\} = \operatorname{argmax}_{a \in A} W(a, \theta)$.

For any information structure \mathcal{I} and strategy profile $\sigma_t = (\sigma_{it})_{i \in N}$ of game $\mathcal{G}_t(\mathcal{I})$,

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

¹⁷Strategy profile $\sigma_t = (\sigma_{it})_{i \in N}$ is in $\text{BNE}_t(\mathcal{G}, \mathcal{I})$ if for each $i \in N$, $x_i^t \in X_i^t$, and a_i with $\sigma_{it}(a_i | x_i^t) > 0$, we have $a_i \in \operatorname{argmax}_{a'_i \in A_i} \sum_{\theta \in \Theta, x_{-i}^t \in X_{-i}^t, a_{-i} \in A_{-i}} \mathbb{P}_t^{\mathcal{I}}(\theta, x_{-i}^t | x_i^t) \sigma_{-i}(a_{-i} | x_{-i}^t) u_i(a'_i, a_{-i}, \theta)$.

denotes the induced ex-ante expected value of the objective. The *objective value*

$$W_t(\mathcal{G}, \mathcal{I}) := \sup_{\sigma_t \in \text{BNE}_t(\mathcal{G}, \mathcal{I})} W_t(\sigma_t, \mathcal{I}) \quad (10)$$

is the ex-ante expected value of the objective under the best BNE of $\mathcal{G}_t(\mathcal{I})$ (Remark 1 discusses the focus on best BNE).

We seek to compare the objective values $W_t(\mathcal{G}, \mathcal{I})$ and $W_t(\mathcal{G}, \tilde{\mathcal{I}})$ under any two information structures \mathcal{I} and $\tilde{\mathcal{I}}$ 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 $\text{SNE}(\mathcal{G}, \theta) \subseteq A$ denote the set of strict Nash equilibria of \mathcal{G} under common knowledge of θ .

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

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

When W represents utilitarian welfare, Assumption 1 is satisfied by our motivating application of incomplete-information coordination games, such as the joint investment game in Example 1 and other leading examples in the literature: Here, coordination on the efficient outcome is a strict Nash equilibrium under common knowledge of the state, but first-order and higher-order uncertainty may impede efficient coordination. An extreme special case are common-interest games \mathcal{G} , where $u_i = u_j = W$ for all i, j . However, under common interest, agents' incentives in \mathcal{G} are *fully* aligned with W even away from common knowledge, in the sense that maximization of the expected objective is a BNE of \mathcal{G} under any information structure. This is much stronger than Assumption 1, which only requires alignment at *certainty* and imposes no restriction on agents' incentives in \mathcal{G} or the relationship with W away from common knowledge.¹⁸

Finally, under more general objective functions W , Assumption 1 includes many other games \mathcal{G} . In particular, as long as \mathcal{G} admits a strict Nash equilibrium $a^\theta \in \text{SNE}(\mathcal{G}, \theta)$ in each state, Assumption 1 is trivially satisfied under the objective function $W(a, \theta) = \mathbf{1}_{\{a=a^\theta\}}$. In this case, the objective value $W_t(\mathcal{G}, \mathcal{I})$ simply measures the ex-ante probability that, after t draws of signals from \mathcal{I} , agents are able to play the

¹⁸For example, Assumption 1 allows for environments where, away from the common knowledge limit, improving agents' information can lead to worse equilibrium outcomes; see the discussion of Lehrer, Rosenberg, and Shmaya (2010) in Remark 1.

common knowledge equilibrium a^θ in each state θ .

4.2 Ranking of Information Structures

Under Assumption 1, we now 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 . For expositional simplicity, we additionally assume that maximizing W requires all agents to distinguish all states:

Assumption 2 (Full separation). *For all $i \in N$ 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 Example 1. However, this assumption is not essential for our analysis, and in Appendix C, we extend Theorem 2 below when Assumption 2 is dropped.

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

$$\lambda(\mathcal{I}) := \min_{\theta \in \Theta} \lambda^\theta(\mathcal{I}) = \min_{i \in N, \theta, \theta' \in \Theta, \theta' \neq \theta} d(\mu_i^\theta, \mu_i^{\theta'}). \quad (11)$$

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

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.

In the proof (Appendix D–E), we show that, for every BNE sequence $\sigma_t \in \text{BNE}_t(\mathcal{G}, \mathcal{I})$,

$$1 - \underbrace{\sum_{\theta \in \Theta, x^t \in X^t} \mathbb{P}_t^\mathcal{I}(\theta, x^t) \sigma_t(a^{\theta, W} | x^t)}_{\text{probability of inefficiency}} \geq \exp[-t\lambda(\mathcal{I}) + o(t)], \quad (12)$$

and that (12) holds with equality for some BNE sequence (σ_t^*) . That is, under information structure \mathcal{I} , the index $\lambda(\mathcal{I})$ is the fastest rate at which inefficiency (i.e., not choosing $a^{\theta, W}$ at θ) can vanish in 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 $W_t(\mathcal{G}, \tilde{\mathcal{I}})$.

As the following example illustrates, this argument relies crucially both on our finding that the efficiency index $\lambda(\mathcal{I})$ characterizes the (ex-ante expected) rate of common learning and that this coincides with the rate of individual learning:

Example 2 (Coordinated attack). Consider a coordinated attack game à la [Morris and Shin \(1998\)](#), with binary states $\Theta = \{\bar{\theta}, \underline{\theta}\}$ and binary actions $A_i = \{0, 1\}$. Each agent i 's utility function takes the form

$$u_i(a, \theta) = \begin{cases} \mathbf{1}_{\{\theta=\bar{\theta}\}} \mathbf{1}_{\{\sum_{j \neq i} a_j \geq \underline{k}\}} - c & \text{if } a_i = 1 \\ 0 & \text{if } a_i = 0. \end{cases}$$

Here $c \in (0, 1)$ denotes the cost of attacking ($a_i = 1$) and an attack is successful if and only if the state is $\bar{\theta}$ and at least $\underline{k} \in \{0, 1, \dots, |N| - 1\}$ other agents also attack. Under utilitarian welfare, $W = \frac{1}{N} \sum_{i \in N} u_i$, the efficient action profiles are $a^{\bar{\theta}} = (1, \dots, 1)$ and $a^{\underline{\theta}} = (0, \dots, 0)$. Note that Assumptions 1–2 are satisfied and Example 1 corresponds to the special case with $|N| = 2$ and $\underline{k} = 1$.

To see why ex-post inefficiency vanishes at least as fast as $\lambda(\mathcal{I})$, note that common p -belief of state θ is *sufficient* for coordination on the efficient outcome a^{θ} , as a^{θ} is a strict Nash equilibrium under common knowledge of θ . More precisely, if $p \in (0, 1)$ is sufficiently large, then for every t , there is a BNE σ_t^* under which a^{θ} is played conditional on event $C_t^p(\theta)$. Under sequence (σ_t^*) , inefficiency vanishes at least as fast as the (ex-ante expected) rate of common learning, which is $\lambda(\mathcal{I})$ by Theorem 1.

Why can inefficiency not vanish faster than the rate of common learning? This is less immediate, as common p -belief is in general *not necessary* for coordination on the efficient outcome.¹⁹ Indeed, if $\underline{k} < |N| - 1$, a successful attack does not require all agents to attack, so there can be BNE in which $a^{\bar{\theta}}$ is played without there being common p -belief of state $\bar{\theta}$. However, note that in any BNE, $a^{\bar{\theta}}$ can only be played whenever all agents at least have *individual* p -belief of $\bar{\theta}$ for some $p > 0$. Hence, inequality (12) follows from the fact that the rate $\lambda(\mathcal{I})$ of common learning coincides with the rate of individual learning.²⁰ ▲

¹⁹See, e.g., [Oyama and Takahashi \(2020\)](#) for systematic analysis of this issue.

²⁰When Assumption 2 is dropped, coordination on the efficient outcome need not even require all agents to have individual p -belief of the true state. Reflecting this, the generalization of Theorem 2 in Appendix C employs a modified learning efficiency index that, for any \mathcal{G} and W , captures the rate at which each agent i learns to distinguish those states that entail different efficient actions $a_i^{\theta, W}$ for i .

By focusing on data-rich settings, Theorem 2 yields some robust implications for information design in coordination games (and other environments satisfying Assumptions 1–2) that apply regardless of the specific game being played. Specifically, as long as agents have access to many signal draws, the structure of the index $\lambda(\mathcal{I})$ suggests two general principles for facilitating coordination:

Focus on first-order uncertainty: A designer should focus on improving agents’ information about the state, whereas providing signals about other agents’ signals (that do not convey any additional information about the state) has a negligible effect. Thus, in contrast with the insight in the literature that uncertainty about opponents’ signals can be a significant obstacle to coordination, our results suggest that, in data-rich settings, reducing such higher-order uncertainty should be a second-order concern.

Egalitarianism: A designer should focus on improving the worst-informed agent’s information about the state.

Remark 1. Focus on best equilibrium. The definition of the objective value $W_t(\mathcal{G}, \mathcal{I})$ in (10) considers the best BNE, similar to the assumption of designer-preferred equilibrium selection that is common in the literature on information design (Section 1.1).²¹ Thus, in the context of incomplete-information coordination games, our comparison of information structures isolates the extent to which they reduce inefficiency due to first-order and higher-order uncertainty about the fundamental, rather than due to equilibrium selection. At the same time, Online Appendix G shows that the learning efficiency index also characterizes the rate at which the entire equilibrium set $\text{BNE}_t(\mathcal{G}, \mathcal{I})$ approaches the set of common knowledge equilibria in each state.

Comparison with $t = 1$. Lehrer, Rosenberg, and Shmaya (2010) assume agents observe *one* 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 W . In contrast, 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 misaligned incentives.²² Both (i) and (ii) rely on agents observing sufficiently many signal draws: When $t = 1$, many information structures are incomparable even when focusing on

²¹Our analysis extends to the case where (10) instead considers the worst BNE and Assumption 1 is replaced with the assumption that $W(\cdot, \theta)$ is strictly *minimized* by some action profile in $\text{SNE}(\mathcal{G}, \theta)$ (capturing settings with a strong misalignment between the designer’s objective and agents’ incentives). Here, Theorem 2 (applied to the objective $-W$) implies that information structures with a *lower* learning efficiency index are better for the designer at all large t .

²²For (i), note that Lehrer, Rosenberg, and Shmaya (2010)’s order implies that each agent’s marginal signal distributions under \mathcal{I} Blackwell-dominate those under $\tilde{\mathcal{I}}$, which ensures $\lambda(\mathcal{I}) \geq \lambda(\tilde{\mathcal{I}})$.

common-interest games; moreover, 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 environments that satisfy Assumptions 1–2 but are not common-interest.

Bounds on T . A natural question is how many signal draws are needed for our ranking to apply. In some specific environments, one can bound the number of draws T beyond which the ranking in Theorem 2 applies, but the bound may in general depend on \mathcal{G} , W , the prior p_0 , and \mathcal{I} and $\tilde{\mathcal{I}}$. It is worth noting that the proof of Theorem 2 does *not* require that $W_T(\mathcal{G}, \mathcal{I})$ and $W_T(\mathcal{G}, \tilde{\mathcal{I}})$ are close to the first-best payoff, so the payoff gap under \mathcal{I} vs. $\tilde{\mathcal{I}}$ can in general still be non-negligible at T .²³ \blacktriangle

5 Discussion

5.1 Information Design in Games with Cheap Data

The learning efficiency index can be used to solve constrained information design problems where information comes at a small cost. Beyond the ordinal implications highlighted following Theorem 2, here the cardinal value that $\lambda(\mathcal{I})$ assigns to each information structure is relevant.

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:

$$\max_{\mathcal{I} \in \mathbb{I}, t \in \mathbb{N}} W_t(\mathcal{I}, \mathcal{G}) \text{ s.t. } tc(\mathcal{I}) \leq \kappa. \quad (13)$$

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 $\kappa > 0$.

The preceding analysis implies the following:

Corollary 2. *Fix any \mathcal{G} and W satisfying Assumptions 1–2 and any finite set \mathbb{I} of information structures. Whenever the budget κ is sufficiently large (i.e., information is sufficiently cheap), the designer’s problem (13) simplifies to*

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

²³Relatedly, while Corollary 1 yields T such that the probability of common p -belief of the true state under \mathcal{I} exceeds that under $\tilde{\mathcal{I}}$, these probabilities need not be close to 1 at T .

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, one can explore properties of the optimal information structure, depending on the nature of the cost function c .

5.2 Convergence of Belief Hierarchies

There is a discussion in the literature about which topologies over belief hierarchies are appropriate for measuring proximity to common knowledge. However, Theorem 1 implies that, as far as the speed of convergence to common knowledge in our setting is concerned, the choice of topology may be less important: The learning efficiency index $\lambda^\theta(\mathcal{I})$ characterizes this speed under several commonly used topologies.

Recall that a belief hierarchy for agent i is a sequence $\tau_i := (\tau_i^1, \tau_i^2, \dots) \in Z_i = (Z_i^1, Z_i^2, \dots)$, 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 agent i 's k th order beliefs, subject to standard coherency requirements across the k th order beliefs τ_i^k for different k (e.g., [Brandenburger and Dekel, 1993](#)).²⁴ Given any information structure \mathcal{I} , each realized signal sequence x_i^t induces a belief hierarchy $\tau_i(x_i^t) \in Z_i$ for agent i . Let $\tau_i(\theta) \in Z_i$ denote i 's belief hierarchy when there is common certainty of state θ .

Let ρ_i^{product} denote a metric on Z_i that induces the **product topology** over agent i 's belief hierarchies. For example, define $\rho_i^{\text{product}}(\tau_i, \tilde{\tau}_i) := (1 - \beta) \sum_k \beta^k \rho^k(\tau_i^k, \tilde{\tau}_i^k)$, where $\beta \in (0, 1)$ and ρ^k denotes the Prokhorov metric over k th order beliefs. The literature has pointed out that the product topology may in general be too coarse (e.g., [Lipman, 2003](#); [Weinstein and Yildiz, 2007](#)), and has proposed several alternative metrics that refine this topology. For instance, the metric for the **uniform-weak topology** ([Chen, Di Tillio, Faingold, and Xiong, 2010](#)) is given by $\rho_i^{\text{uniform}}(\tau_i, \tilde{\tau}_i) := \sup_k \rho^k(\tau_i^k, \tilde{\tau}_i^k)$.

Theorem 1 implies the following:

Corollary 3. *Fix any information structure \mathcal{I} and state $\theta \in \Theta$. Under both the product and uniform-weak topologies, the rate of convergence to common certainty of θ is given by $\lambda^\theta(\mathcal{I})$: For all sufficiently small $\varepsilon > 0$, we have*

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

²⁴For 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.

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

Thus, although differences between these topologies can play a significant role in general, these differences do not matter for the speed of convergence to common certainty in the current learning setting.²⁵ The proof of Corollary 3 exploits the fact that common learning has the same rate as individual learning.

5.3 Higher-Order Expectations

Beyond its use in the current paper, Lemma 1 can shed light on 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 $T := \prod_{i \in N} T_i$. 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 agent i induces a conditional distribution $\pi(\cdot \mid t_i) \in \Delta(T)$ over type profiles. By identifying each $t_j \in T_j$ with the point-mass distribution $\delta_{t_j} \in \Delta(T_j)$, we can associate with $\pi(\cdot \mid t_i)$ a sequence of higher-order expectations about other agents’ types. In particular, $\mathbb{E}_{t_i}[t_j] := \sum_{t_j \in T_j} \pi(t_j \mid t_i) \delta_{t_j} \in \Delta(T_j)$ is t_i ’s expectation of j ’s type, $\mathbb{E}_{t_i} \mathbb{E}_{t_j}[t_k] := \sum_{t_j \in T_j, t_k \in T_k} \pi(t_j \mid t_i) \pi(t_k \mid t_j) \delta_{t_k} \in \Delta(T_k)$ is t_i ’s expectation of j ’s expectation of k ’s type, and so on.

A seminal result due to Samet (1998) is that any such sequence of higher-order expectations converges to the prior distribution as the number of iterations grows large. Formally, consider any sequence of agents $i_0, i_1, \dots \in N$ in which all $i \in N$ appear infinitely often and any initial type $t_{i_0} \in T_{i_0}$. Then his result adapted to the current setting implies that²⁶

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

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

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

²⁵The result extends to other topologies considered in the literature, for example, the strategic topology (Dekel, Fudenberg, and Morris, 2006), which is in between the product and uniform-weak topologies.

²⁶See the proof of his Proposition 6.

and iteratively, for each k ,

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

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. While Samet’s insight can be applied to analyze equilibrium behavior in beauty contests in the limit as coordination motives become strong (Golub and Morris, 2017), our non-asymptotic finding may be useful for conducting comparative statics with respect to coordination motives away from the limit.

6 Conclusion

This paper conducted a comparison of multi-agent information structures in a learning setting where players have access to rich data. We showed that the speed of common learning under each information structure coincides with the speed of individual learning and used this to rank information structures in terms of their value in coordination games.

As a natural starting point, we assumed that signal spaces are finite and signals are i.i.d. across draws. This allowed us to build on CEMS’ result that this setting always gives rise to common learning.

With infinite signals, CEMS exhibit an example in which common learning fails even though individual learning is successful; at the same time, there are other natural infinite-signal settings, in particular Gaussian signal structures, that do give rise to common learning.²⁷ Online Appendix H analyzes such Gaussian environments and shows that common and individual learning again occur at the same exponential rate.

A simple setting where signals are not identically distributed across draws is when draws independently alternate across two different information structures \mathcal{I} and $\tilde{\mathcal{I}}$. However, this is equivalent to considering repeated independent draws from the product information structure $\mathcal{I} \times \tilde{\mathcal{I}}$, and thus is a special case of our setting in this paper. Online Appendix I analyzes when the learning efficiency index $\lambda(\mathcal{I} \times \tilde{\mathcal{I}})$ of alternating draws from \mathcal{I} and $\tilde{\mathcal{I}}$ is greater or less than the sum $\lambda(\mathcal{I}) + \lambda(\tilde{\mathcal{I}})$ of their separate indices,

²⁷Recently, Faingold and Tamuz (2022) derive general sufficient conditions for common learning under infinite signals. In contrast, Dogan (2018) shows that with (uncountably) infinite states, common learning fails under mild conditions (even if signals are finite).

shedding light on whether \mathcal{I} and $\tilde{\mathcal{I}}$ are complements or substitutes.

When signals are correlated across draws, there are some known settings in which common learning fails even though individual learning is successful and others in which common learning succeeds (e.g., [Steiner and Stewart, 2011](#); [Cripps, Ely, Mailath, and Samuelson, 2013](#)). We leave the analysis of such settings for future work, in particular, the question whether common learning can be successful but occur at a slower rate than individual learning.

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

A.1 Preliminary Definitions

The following will be used throughout the appendix. As in CEMS, given any information structure \mathcal{I} and agents i and j , we consider the matrix $M_{ij}^\theta \in \mathbb{R}^{X_i \times X_j}$ with (x_i, x_j) -th entry

$$M_{ij}^\theta(x_i, x_j) = \mu^\theta(x_j \mid x_i).$$

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 $\mathbb{E}[\nu_{jt} \mid \theta, \nu_{it}] = \nu_{it} M_{ij}^\theta$ (treating $\nu_{it} \in \Delta(X_i) \subseteq \mathbb{R}^{1 \times X_i}$ as a vector). Moreover, $\mu_i^\theta M_{ij}^\theta = \mu_j^\theta$.

For each $d < \lambda^\theta(\mathcal{I})$ and t , define the event

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

Finally, we call an information structure \mathcal{I} **fully private** if the joint distribution μ^θ has full support on X in all states θ . We call \mathcal{I} **public** if signals are perfectly correlated

across agents.²⁸

A.2 Proof of Lemma 1

Fix $\theta \in \Theta$, distinct $i, j \in N$, 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 $m(\cdot | x_i), m(\cdot | x_j), m'(\cdot | x_i), m'(\cdot | x_j)$ 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

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

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

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

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

To show the strict inequality, suppose that $\nu_i \neq \mu_i^\theta$ and μ_j^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 $x_j \in \text{supp}(\nu_i M_{ij}^\theta)$,

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

where the inequality holds since $M_{ij}^\theta(x_i, x_j), M_{ij}^\theta(x'_i, x_j) > 0$ by the full-support as-

²⁸That is, $X_i = X_j$ for all i, j , and for each $x \in X$ and θ , $\mu^\theta(x) = \begin{cases} \mu_i^\theta(x_i) & \text{if } x_i = x_j \text{ for all } i, j \\ 0 & \text{otherwise} \end{cases}$.

sumption on μ^θ . By Gibbs' inequality, this guarantees

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

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

A.3 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 $t \geq T$, $\theta \in \Theta$, $i \in N$, and x_i^t ,*

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

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–C.1:

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

$$\text{KL}(\nu_{it}, \mu_i^\theta) \leq d \implies \mathbb{P}_t^{\mathcal{I}} \left(\bigcup_{\theta' \in \Pi_i(\theta)} (\{\theta'\} \cap F_{-it}(\theta', d)) | x_i^t \right) \geq p. \quad (14)$$

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

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

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

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

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 $\text{KL}(\cdot, \mu_i)$ 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 $\theta' \in \Theta$,

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

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

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

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

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

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

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

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

Proof of Claim 2. Take any $t \geq 1$ and x_i^t such that $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$. Then for each $\theta' \notin \Pi_i(\theta)$, we have $\text{KL}(\nu_{it}, \mu_i^{\theta'}) > d + \kappa$. Indeed, otherwise $\text{KL}(\nu_{it}, \mu_i^\theta), \text{KL}(\nu_{it}, \mu_i^{\theta'}) \leq d + \kappa < d(\mu_i^\theta, \mu_i^{\theta'})$, contradicting the definition of $d(\mu_i^\theta, \mu_i^{\theta'})$.

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

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

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 $\text{KL}(\nu_{it}, \mu_i^{\theta'}) > d + \kappa$,

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

proving Claim 2. \square

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 $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$, we have

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

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

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 (as defined in Appendix A.1). Appendix E extends the proof to general information structures.

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

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log (1 - \mathbb{P}_t^{\mathcal{I}}(C_t^p(\theta) \mid \theta)) \leq -\lambda^\theta(\mathcal{I}). \quad (15)$$

Take any $d \in (0, \lambda^\theta(\mathcal{I}))$. Applying Lemma A.2 to the case with $\Pi_i(\theta) = \{\theta\}$ for each $i \in N$, there exists $T > 0$ such that, for all $t \geq T$, (i) $F_t(\theta, d) \subseteq B_t^p(\theta)$, and (ii) $F_t(\theta, d) \subseteq B_t^p(F(\theta, d))$, i.e., $F_t(\theta, d)$ is p -evident. Thus, by Monderer and Samet (1989), we have $F_t(\theta, d) \subseteq C_t^p(\theta)$ for all $t \geq T$. Therefore,

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

where the last equality follows from Sanov's theorem. Since this holds for all $d < \lambda^\theta(\mathcal{I})$, this establishes (15).

We next establish that

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log (1 - \mathbb{P}_t^{\mathcal{I}}(B_t^q(\theta) \mid \theta)) \geq -\lambda^\theta(\mathcal{I}). \quad (16)$$

Take $i \in N$ 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 $\text{KL}(\nu_i, \mu_i^\theta) = \text{KL}(\nu_i, \mu_i^{\theta'}) < d$. Hence, for some ν_i' close to ν_i ,

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

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

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

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

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

where the final inequality holds by Sanov's theorem. Since this is true for all $d > \lambda^\theta(\mathcal{I})$, this establishes (16). \square

C Ranking of Information Structures without Assumption 2

This appendix generalizes the ranking in Theorem 2 when Assumption 2 is dropped, i.e., playing the W -optimal action profile $a^{\theta, W}$ need not require all agents to distinguish all states. The idea is to construct generalized learning efficiency indices that account for the presence of “equivalent” states for some players.

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

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

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 $\Pi^W := (\Pi_i^W)_{i \in N}$ denote the collection

of all agents' partitions.

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

$$\lambda(\mathcal{I}, \Pi) := \min_{i \in N, \theta, \theta' \in \Theta, \theta' \notin \Pi_i(\theta)} d(\mu_i^\theta, \mu_i^{\theta'}).^{29}$$

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 (11). 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 or public (as defined in Appendix A.1).

Theorem C.1. *Fix any collection $\Pi = (\Pi_i)_{i \in N}$ of partitions over Θ . Take any information structures \mathcal{I} and $\tilde{\mathcal{I}}$, each of which is either fully private or public, with $\lambda(\mathcal{I}, \Pi) > \lambda(\tilde{\mathcal{I}}, \Pi)$. 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 $t \geq T$.*

Theorem C.1 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. The proof of Theorem C.1 is in Appendix D.

D Proof of Theorem 2 (Fully Private Case) and Theorem C.1

Below we prove Theorem C.1. 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 E 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.

D.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

²⁹Slightly abusing notation, we set the index to be ∞ when Π is degenerate (i.e., $\Pi_i(\theta) = \Theta$ for all i).

on this probability for arbitrary sequences of strategy profiles (σ_t) :

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

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

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

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

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

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

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

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 ,

$$\begin{aligned} \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t | \theta) \tilde{\sigma}_{it}(a_i^\theta | x_i^t) &\leq \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t | \theta) \sigma_{it}^*(a_i^\theta | x_i^t) \\ \text{or } \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t | \theta') \tilde{\sigma}_{it}(a_i^{\theta'} | x_i^t) &\leq \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t | \theta') \sigma_{it}^*(a_i^{\theta'} | x_i^t). \end{aligned} \tag{17}$$

Hence,

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

where the first inequality follows from the construction of $(\tilde{\sigma}_{it})$ and the second inequality uses (17). 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),³⁰ i.e.,

$$\begin{aligned}
\lim_{t \rightarrow \infty} \frac{1}{t} \log \left(1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t | \theta) \sigma_{it}^*(a_i^\theta | x_i^t) \right) &= \lim_{t \rightarrow \infty} \frac{1}{t} \log \left(1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t | \theta') \sigma_{it}^*(a_i^{\theta'} | x_i^t) \right) \\
&= -d(\mu_i^\theta, \mu_i^{\theta'}).
\end{aligned}$$

This implies that

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

as claimed. \square

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

Lemma D.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 $(\sigma_t) \in \text{BNE}_t(\mathcal{G}, \mathcal{I})$ such*

³⁰This in turn follows from a simple application of Sanov's theorem.

that, for all $\theta \in \Theta$,

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

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 $d < \lambda(\mathcal{I}, \Pi^W) := \min_{i \in N, \theta \in \Theta, \theta' \notin \Pi_i(\theta)} d(\mu_i^\theta, \mu_i^{\theta'})$. Let $\Sigma_{it}(d)$ denote the set of i 's strategies at t such that $\sigma_{it}(a_i^\theta | x_i^t) = 1$ whenever $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$. This set is well-defined by the choice of d , i.e., there is no $\nu_i \in \Delta(X_i)$ such that $\text{KL}(\nu_i, \mu_i^\theta), \text{KL}(\nu_i, \mu_i^{\theta'}) \leq d$ 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 (14) holds for all i, θ , and $t \geq T$. Thus, for all $t \geq T$, each agent i 's best response against any strategy profile in $\prod_{j \neq i} \Sigma_{jt}(d)$ must be in $\Sigma_{it}(d)$, because whenever $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$, 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 $\prod_i \Sigma_{it}(d)$ (which is convex), 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., $\mathbb{P}_t^{\mathcal{I}}(\cdot | x_i^t) = \mathbb{P}_t^{\mathcal{I}}(\cdot | x_j^t)$ for all i, j , and t . Moreover, $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d \iff \text{KL}(\nu_{jt}, \mu_j^\theta) \leq d$ 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

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

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 $t \rightarrow \infty$,

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

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

D.2 Remaining Proof

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 $\theta \in \Theta$,

$$W(a^\theta, \theta) - \sum_{x^t, a} \mathbb{P}_t^{\mathcal{I}}(x^t | \theta) \sigma_t(a | x^t) W(a, \theta) \leq c \left(1 - \sum_{x^t} \mathbb{P}_t^{\mathcal{I}}(x^t | \theta) \sigma_t(a^\theta | x^t) \right), \quad (18)$$

$$W(a^\theta, \theta) - \sum_{\tilde{x}^t, a} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}^t | \theta) \tilde{\sigma}_t(a | \tilde{x}^t) W(a, \theta) \geq \tilde{c} \left(1 - \sum_{\tilde{x}^t} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}^t | \theta) \tilde{\sigma}_t(a^\theta | \tilde{x}^t) \right). \quad (19)$$

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

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

which by (18) implies

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \sum_{\theta} p_0(\theta) \left(W(a^\theta, \theta) - \sum_{x^t} \mathbb{P}_t^{\mathcal{I}}(x^t | \theta) \sigma_t(a^\theta | x^t) W(a, \theta) \right) \leq -\lambda(\mathcal{I}, \Pi). \quad (20)$$

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

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

which by (19) implies

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \sum_{\theta} p_0(\theta) \left(W(a^\theta, \theta) - \sum_{\tilde{x}^t} \mathbb{P}_{\tilde{\mathcal{I}}}^{\tilde{x}^t}(\tilde{x}^t | \theta) \tilde{\sigma}_t(a^\theta | \tilde{x}^t) W(a, \theta) \right) \geq -\lambda(\tilde{\mathcal{I}}, \Pi). \quad (21)$$

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 (20) and (21) and the assumption that $\lambda(\mathcal{I}, \Pi) > \lambda(\tilde{\mathcal{I}}, \Pi)$. \square

E 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 $S \subseteq N$, 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 $\varepsilon_1, \dots, \varepsilon_{|N|} \in [0, d)$, define the following event:

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

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

The main step in extending the proofs of Theorems 1–2 is the following result,

which we prove in Appendix E.1.

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

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

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

Using Proposition E.1, the proof of Theorem 1 extends as follows. It suffices to prove (15) for general \mathcal{I} , as the argument for (16) in Appendix B did not rely on the full-support assumption. To prove (15), 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 $F_t(\theta, d, \varepsilon_1, \dots, \varepsilon_{|N|})$ constructed in Proposition E.1 satisfy

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

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

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

Since this holds for all $d < \lambda^\theta(\mathcal{I})$, (15) follows.

To extend the proof of Theorem 2, it is sufficient to establish Lemma D.2 for general \mathcal{I} under Assumption 2, as the remaining steps of the proof in Appendix D 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 D.2, and take any $\varepsilon \in (0, d)$. Applying Proposition E.1 and following the same steps as in the original proof of Lemma D.2, we construct a BNE sequence (σ_t) such that for all large enough t and each θ , we have $\sigma_t(a^\theta \mid x^t) = 1$ at all signal sequences $x^t \in F_t(\theta, d, \varepsilon_1, \dots, \varepsilon_{|N|})$. Thus,

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

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

E.1 Proof of Proposition E.1

E.1.1 Generalization of Lemma 1

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

$$(\nu M_i^\theta)(x_i, x_{-i}) := \nu_i(x_i)\mu^\theta(x_{-i}|x_i), \quad \text{for all } (x_i, x_{-i}) \in X^\theta. \quad (22)$$

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 E.1. *Take any $\nu \in \Delta(X^\theta)$, $i \in N$, and $S \subseteq N$. Then $\text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) \leq \text{KL}\left(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta\right)$. Moreover, the inequality is an equality only if $\nu_{H_i^\theta}(\cdot|h) = \mu_{H_i^\theta}^\theta(\cdot|h)$ 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

$$\begin{aligned} \text{KL}(\nu M_i^\theta, \mu^\theta) &= \text{KL}\left((\nu M_i^\theta)_{H_i^\theta}, \mu_{H_i^\theta}^\theta\right) + \sum_{h \in H_i^\theta} (\nu M_i^\theta)_{H_i^\theta}(h) \text{KL}((\nu M_i^\theta)(\cdot|h), \mu^\theta(\cdot|h)) \\ &= \text{KL}\left(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta\right), \end{aligned} \quad (23)$$

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 (22). The chain rule also implies that

$$\begin{aligned} \text{KL}(\nu M_i^\theta, \mu^\theta) &= \text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) + \sum_{h \in H_S^\theta} (\nu M_i^\theta)_{H_S^\theta}(h) \text{KL}((\nu M_i^\theta)(\cdot|h) | \mu^\theta(\cdot|h)) \\ &\geq \text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right). \end{aligned} \quad (24)$$

Combining (23)–(24) yields $\text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) \leq \text{KL}\left(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta\right)$.

For the “moreover” part, suppose that $\text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) = \text{KL}\left(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta\right)$. Then, by (23)–(24), 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 $(\nu M_i^\theta)_{H_i^\theta}(h) > 0$, (22) 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 $(\nu M_i^\theta)_{H_i^\theta}(\cdot|h) = \mu_{H_i^\theta}^\theta(\cdot|h)$. But by (22), $(\nu M_i^\theta)_{H_i^\theta} = \nu_{H_i^\theta}$ and $(\nu M_i^\theta)_{H_i^\theta \wedge H_S^\theta} = \nu_{H_i^\theta \wedge H_S^\theta}$. Thus, $\nu_{H_i^\theta}(\cdot|h) = \mu_{H_i^\theta}^\theta(\cdot|h)$ for all $h \in H_i^\theta \wedge H_S^\theta$ with $\nu_{H_i^\theta \wedge H_S^\theta}(h) > 0$. \square

Lemma E.1 yields the following corollary:

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

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

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

Proof. Consider any $S \subseteq N$, $i \notin S$, and $\nu \in \Delta(X^\theta)$ with $\text{KL}(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta) \leq d$ and $\max_{|S'|=|S|+1} \text{KL}(\nu_{H_{S'}^\theta}, \mu_{H_{S'}^\theta}^\theta) \leq d - \varepsilon$. It suffices to prove that $\text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) < d$, 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 E.1 implies $\text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) \leq \text{KL}(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta) \leq d$. Thus, we can focus on the case in which $\text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) = \text{KL}(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta)$. In this case,

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

where the first equality uses the chain rule and the second one holds by the ‘‘moreover’’ part of Lemma E.1. \square

E.1.2 Completing the Proof

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

$$\text{KL}(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta) \leq d \quad \text{and} \quad \text{KL}(\nu_{H_N^\theta}, \mu_{H_N^\theta}^\theta) \leq d - \varepsilon,$$

we have $\text{KL}\left((\nu M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) < d - \rho_{|N|-1}$.

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

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

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

This yields a sequence

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

with the property that whenever

$$\text{KL}(\nu_{H_S^\theta}, \mu_{H_S^\theta}^\theta) \leq d - \varepsilon_{|S|} \text{ for all } S \subseteq N,$$

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

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

$$\mathbb{P}_t^{\mathcal{I}}\left(\left\{\text{KL}\left((\nu_t)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) \leq d - \varepsilon_{|S|}\right\} \mid x_i^t, \theta\right) \geq p \quad (25)$$

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

To show (25), fix any $p \in (0, 1)$. First, consider $i \in N$ and $S \subseteq N$ 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 $x^t \in F_t(\theta, d, \varepsilon_1, \dots, \varepsilon_{|N|})$, then

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

as required.

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

$$\text{KL}\left((\tilde{\nu}_t M_i^\theta)_{H_S^\theta}, \mu_{H_S^\theta}^\theta\right) \leq d - \rho_{|S|}. \quad (26)$$

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

$\Delta(X^\theta)$,

$$\text{KL}(\nu'_{H_S^\theta}, \mu_{H_S^\theta}^\theta) \leq d - \rho_{|S|} \text{ and } \|\nu' - \nu\| < \kappa \implies \text{KL}(\nu_{H_S^\theta}, \mu_{H_S^\theta}^\theta) \leq d - \varepsilon_{|S|}. \quad (27)$$

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

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

Combined with (26)–(27), this implies that (25) holds for every $t \geq T$ and signal sequence $x^t \in F_t(\theta, d, \varepsilon_1, \dots, \varepsilon_{|N|})$. \square

F Proofs for Section 5

F.1 Proof of Corollary 2

Fix any $\mathcal{I} \in \mathbb{I}$. Let t_κ denote an optimal number of signal draws from \mathcal{I} under budget κ . The analysis in Section 4.2 implies that, for every BNE sequence $\sigma_{t_\kappa} \in \text{BNE}_{t_\kappa}(\mathcal{G}, \mathcal{I})$,

$$1 - \sum_{\theta \in \Theta, x^{t_\kappa} \in X^{t_\kappa}} \mathbb{P}_{t_\kappa}^{\mathcal{I}}(\theta, x^{t_\kappa}) \sigma_{t_\kappa}(a^{\theta, W} \mid x^{t_\kappa}) \geq \exp[-t_\kappa \lambda(\mathcal{I}) + o(t_\kappa)], \quad (28)$$

and that (28) holds with equality for some BNE sequence $(\sigma_{t_\kappa}^*)$. Note that $\lim_{\kappa \rightarrow \infty} t_\kappa = \infty$ holds by optimality, as otherwise the designer's value is bounded away from the first-best payoff as $\kappa \rightarrow \infty$. Thus, maximizing the rate of convergence in the RHS of (28) under the budget constraint implies $\lim_{\kappa \rightarrow \infty} t_\kappa / \kappa = 1/c(\mathcal{I})$. Hence, the difference between the first-best payoff $\sum_\theta p_0(\theta) W(a^{\theta, W}, \theta)$ and the designer's value under each information structure \mathcal{I} takes the form $\exp[-\kappa \frac{\lambda(\mathcal{I})}{c(\mathcal{I})} + o(\kappa)]$.

Since \mathbb{I} is finite, there then exists κ^* such that for all $\kappa \geq \kappa^*$ and $\mathcal{I}, \mathcal{I}' \in \mathbb{I}$ with $\frac{\lambda(\mathcal{I})}{c(\mathcal{I})} > \frac{\lambda(\mathcal{I}')}{c(\mathcal{I}')}$, it is suboptimal for the designer to choose \mathcal{I}' . \square

F.2 Proof of Corollary 3

The convergence under the product topology cannot be faster than $\lambda^\theta(\mathcal{I})$. To see this, fix any $\varepsilon \leq \beta(1 - \beta)$. Then there exists $p \in (0, 1)$ such that

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

The convergence under the uniform-weak topology cannot be slower than $\lambda^\theta(\mathcal{I})$. To see this, fix any $\varepsilon > 0$. 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 . Thus,

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

Finally, by definition, convergence under the uniform-weak topology cannot be faster than under the product topology. \square

References

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2016): “Fragility of asymptotic agreement under Bayesian learning,” *Theoretical Economics*, 11(1), 187–225.
- AWAYA, Y., AND V. KRISHNA (2022): “Commonality of Information and Commonality of Beliefs,” working paper.
- 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.
- (2019): “Information design: A unified perspective,” *Journal of Economic Literature*, 57(1), 44–95.
- 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.
- BROOKS, B., A. FRANKEL, AND E. KAMENICA (2021): “Information Hierarchies,” *Econometrica*.
- CARLSSON, H., AND E. VAN DAMME (1993): “Global games and equilibrium selection,” *Econometrica*, pp. 989–1018.
- CHEN, Y.-C., A. DI TILLIO, E. FAINGOLD, AND S. XIONG (2010): “Uniform topologies on types,” *Theoretical Economics*, 5(3), 445–478.
- CHERNOFF, H. (1952): “A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations,” *The Annals of Mathematical Statistics*, pp. 493–507.
- 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.
- FAINGOLD, E., AND O. TAMUZ (2022): “Statistical Foundations of Common Knowledge,” working paper.
- 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.
- HIRSHLEIFER, J. (1971): “The Private and Social Value of Information and the Reward to Inventive Activity,” *The American Economic Review*, 61(4), 561–574.
- KAJII, A., AND S. MORRIS (1997): “The robustness of equilibria to incomplete information,” *Econometrica*, pp. 1283–1309.
- KAMENICA, E. (2019): “Bayesian persuasion and information design,” *Annual Review of Economics*, 11, 249–272.
- 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., D. OYAMA, AND S. TAKAHASHI (2020): “Implementation via information design in binary-action supermodular games,” working paper.

- MORRIS, S., AND H. S. SHIN (1998): "Unique equilibrium in a model of self-fulfilling currency attacks," *American Economic Review*, pp. 587–597.
- (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.
- OYAMA, D., AND S. TAKAHASHI (2020): "Generalized Belief Operator and Robustness in Binary-Action Supermodular Games," *Econometrica*, 88(2), 693–726.
- PEŒSKI, M. (2008): "Comparison of information structures in zero-sum games," *Games and Economic Behavior*, 62(2), 732–735.
- ROSENBERG, D., AND N. VIEILLE (2019): "On the efficiency of social learning," *Econometrica*, 87(6), 2141–2168.
- RUBINSTEIN, A. (1989): "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"," *The American Economic Review*, pp. 385–391.
- SAMET, D. (1998): "Iterated expectations and common priors," *Games and Economic Behavior*, 24(1-2), 131–141.
- STEINER, J., AND C. STEWART (2011): "Communication, timing, and common learning," *Journal of Economic Theory*, 146(1), 230–247.
- 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

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 ,

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

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

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

We also denote by $\text{NE}(\mathcal{G})$ 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 $\text{BNE}_t^\varepsilon(\mathcal{G}, \mathcal{I}) \subseteq \Delta(\Theta \times A)$ 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} ,*

$$\sup_{m_t \in \text{BNE}_t(\mathcal{G}, \mathcal{I})} \inf_{m \in \text{CE}^\varepsilon(\mathcal{G})} \|m_t - m\| \leq \exp[-t\lambda(\mathcal{I}) + o(t)], \quad (29)$$

$$\sup_{m \in \text{NE}(\mathcal{G})} \inf_{m_t \in \text{BNE}_t^\varepsilon(\mathcal{G}, \mathcal{I})} \|m_t - m\| \leq \exp[-t\lambda(\mathcal{I}) + o(t)]. \quad (30)$$

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

By (29), 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 (30), $\lambda(\mathcal{I})$ 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 E. Fix any $\varepsilon > 0$ and basic game \mathcal{G} .

Inequality (29): 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 $t \geq T$, $i \in N$, and $\theta \in \Theta$, whenever $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$, then

$$\mathbb{P}_t^{\mathcal{I}}(\{\theta\} \cap F_t(\theta, d) \mid x_i^t) \geq p. \quad (31)$$

Take any $t \geq T$, BNE σ_t of $\mathcal{G}_t(\mathcal{I})$, $i \in N$, $\theta \in \Theta$, and $x_i^t \in X_i^t$ such that $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$. Then for any a_i with $\sigma_{it}(a_i \mid x_i^t) > 0$, the fact that σ_t is a BNE implies that, for all $a'_i \in A_i$,

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

which, by (31) and the choice of p , implies that

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

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 $t \geq T$,

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

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

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

Since this holds for any $d < \lambda(\mathcal{I})$, this proves inequality (29).

Inequality (30): 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 $t \geq T$, $i \in N$, and $\theta \in \Theta$, whenever $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$, then (31) holds.

Take any $m \in \text{NE}(\mathcal{G})$, 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 $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$. 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 $\text{KL}(\nu_{it}, \mu_i^\theta) > d$ is interim optimal against σ_{-it} . Moreover, for $t \geq T$, (31) and the choice of p ensure that each player i 's action conditional on a signal sequence x_i^t with $\text{KL}(\nu_{it}, \mu_i^\theta) \leq d$ (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 $t \geq T$,

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

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

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

Since this holds for any $d < \lambda(\mathcal{I})$, this proves inequality (30).

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 $m \in \text{CE}^\varepsilon(\mathcal{G}) \cup \text{NE}(\mathcal{G})$, we have $m(a_i|\theta) = m(a'_i|\theta') = 1$. Thus,

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

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

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

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

Thus, as $t \rightarrow \infty$,

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

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

as claimed. \square

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 $\theta = \underline{\theta}, \bar{\theta}$; extending to more players/states is straightforward. Assume that conditional on state θ , signal profiles are drawn i.i.d. according to

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

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

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

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

where the final equality holds by Cramér's theorem.³¹ Thus, as $t \rightarrow \infty$,

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

To calculate the speed of common learning in state $\bar{\theta}$, 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

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

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

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

³¹Indeed, since m_{it} is the sample mean of i.i.d. draws from $\mathcal{N}(m_i^\theta, (\sigma_i)^2)$, Cramér's theorem implies that $\lim_{t \rightarrow \infty} -\frac{1}{t} \log \mathbb{P}_t [m_{it} < \frac{1}{2} \mid \bar{\theta}] = I(\frac{1}{2})$, 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 $\mathcal{N}(m_i^\theta, (\sigma_i)^2)$.

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

$$|\mathbb{E}[m_{1t}|m_{2t}, \bar{\theta}] - 1| \leq |\rho|d\frac{\sigma_1}{\sigma_2}, \quad |\mathbb{E}[m_{2t}|m_{1t}, \bar{\theta}] - 1| \leq |\rho|d.$$

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

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

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

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

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

I Information Structures as Complements vs. Substitutes

Our analysis suggests a novel formalization of when two information structures \mathcal{I} and $\tilde{\mathcal{I}}$ are complements or substitutes.³² We extend our baseline setting with repeated draws from a single information structure \mathcal{I} by 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 1. 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 rate of common learning under the combined information structure $\mathcal{I} \times \tilde{\mathcal{I}}$ is more than

³²[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.

twice as fast as the rate of common learning under \mathcal{I} or $\tilde{\mathcal{I}}$ alone.³³ 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' ,

$$d(\mu_i^\theta \times \tilde{\mu}_i^\theta, \mu_i^{\theta'} \times \tilde{\mu}_i^{\theta'}) \leq d(\mu_i^\theta, \mu_i^{\theta'}) + d(\tilde{\mu}_i^\theta, \tilde{\mu}_i^{\theta'}). \quad (32)$$

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.³⁴

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 I.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 (32) holds with equality. Thus, only the second channel is relevant and \mathcal{I} and $\tilde{\mathcal{I}}$ are complements. \blacktriangle

³³That 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 .

³⁴Formally, observe that $d(\mu_i^\theta, \mu_i^{\theta'}) = \min_{\nu_i \in \Delta(X_i)} \text{KL}(\nu_i, \mu_i^\theta)$ s.t. $\text{KL}(\nu_i, \mu_i^\theta) = \text{KL}(\nu_i, \mu_i^{\theta'})$. Combined with the fact that KL-divergence is additive across independent distributions, this yields

$$\begin{aligned} d(\mu_i^\theta \times \tilde{\mu}_i^\theta, \mu_i^{\theta'} \times \tilde{\mu}_i^{\theta'}) &= \min_{\nu_i \in \Delta(X_i), \tilde{\nu}_i \in \Delta(\tilde{X}_i)} \text{KL}(\nu_i, \mu_i^\theta) + \text{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta) \\ &\text{s.t. } \text{KL}(\nu_i, \mu_i^\theta) + \text{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta) = \text{KL}(\nu_i, \mu_i^{\theta'}) + \text{KL}(\tilde{\nu}_i, \tilde{\mu}_i^{\theta'}). \end{aligned}$$

This implies (32), because $\text{KL}(\nu_i, \mu_i^\theta) + \text{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta) = \text{KL}(\nu_i, \mu_i^{\theta'}) + \text{KL}(\tilde{\nu}_i, \tilde{\mu}_i^{\theta'})$ is possible even if $\text{KL}(\nu_i, \mu_i^\theta) \neq \text{KL}(\nu_i, \mu_i^{\theta'})$ and $\text{KL}(\tilde{\nu}_i, \tilde{\mu}_i^\theta) \neq \text{KL}(\tilde{\nu}_i, \tilde{\mu}_i^{\theta'})$.