


LEARNING EFFICIENCY OF MULTIAGENT  
INFORMATION STRUCTURES

By

Mira Frick, Ryota Iijima, and Yuhta Ishii

COWLES FOUNDATION PAPER NO. 1860



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS  
YALE UNIVERSITY  
Box 208281  
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

# Learning Efficiency of Multiagent Information Structures

---

Mira Frick

*Yale University*

Ryota Iijima

*Yale University*

Yuhta Ishii

*Pennsylvania State University*

Which information structures are more effective at eliminating first- and higher-order uncertainty and hence at facilitating efficient play in coordination games? We consider a learning setting where players observe many private signals about the state. First, we characterize multiagent learning efficiency, that is, the rate at which players approximate common knowledge. We find that this coincides with the rate at which first-order uncertainty disappears, as higher-order uncertainty vanishes faster than first-order uncertainty. Second, we show that with enough signal draws, information structures with higher learning efficiency induce higher equilibrium welfare. We highlight information design implications for games in data-rich environments.

For helpful comments that substantially improved the paper, we thank the editor, three anonymous referees, numerous seminar audiences, as well as Nageeb Ali, Yu Awaya, Dirk Bergemann, Tianxiang Cong, Roberto Corrao, Martin Cripps, 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. Frick gratefully acknowledges the financial support from a Sloan

Electronically published November 3, 2023

*Journal of Political Economy*, volume 131, number 12, December 2023.

© 2023 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/725171>

## I. 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, e.g., that data is cheap or abundant). Our starting point is a classic result due to Cripps et al. (2008), 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—that is, 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 multiagent 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

---

Research Fellowship. Frick and Iijima thank Stanford University for hospitality while part of this research was conducted. This paper was edited by Emir Kamenica.

outcomes whenever players have access to sufficiently many signal draws. Our characterization of the speed of common learning is essential for deriving this ranking. Moreover, on the basis of the structure of the learning efficiency index, this ranking yields some robust implications for information design in coordination games that are played in data-rich settings.

Section II 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 III 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 1 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 also on their infinite hierarchy of higher-order beliefs. However, perhaps surprisingly, theorem 1 shows that the probability of common  $p$ -belief converges to 1 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 (KL) divergence to formalize that players' higher-order uncertainty vanishes faster than their first-order uncertainty (lemma 1).

Section IV 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, that is, 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 nondegenerate prior  $p_\theta$ . 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,  $(1/2)(u_1(a, \theta) + u_2(a, \theta))$ , 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  $\theta$ , but incomplete information prevents the efficient outcome from being an equilibrium.

Now suppose that prior to choosing actions, players learn about state  $\theta$  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  $\theta$  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  $\theta$  are summarized by the following table:

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

Here, the individual precision parameter  $\gamma \in (1/2, 1)$  captures the probability with 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  $\gamma$  help to reduce players' first-order uncertainty about the state, while  $\rho$  influences players' predictions of their opponent's signals, that is, their higher-order uncertainty. Thus, in comparing two information structures parametrized by  $(\gamma, \rho)$  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  $(\gamma, \rho)$  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  $\gamma$ . Thus, for any  $p_0$  and  $c$ , higher levels of individual precision  $\gamma$  allow for more efficient equilibrium play when players observe sufficiently many signal draws; in contrast, the effect of correlation  $\rho$  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.

*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 Cripps et al. (2008).<sup>1</sup> As mentioned, our contribution relative to their paper is to provide a comparison of different 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 Cripps et al.'s (2008) proof approach, but as section III.B 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 indexes are not equal;

<sup>1</sup> Other papers (e.g., Steiner and Stewart 2011; Cripps et al. 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.

more recently, Awaya and Krishna (2022) study the natural complementary case where information structures have common marginal signal distributions (and hence equal learning efficiency indexes) but differ in their correlation structure (see discussion in sec. III.A).

Moscarini and Smith (2002) derive an efficiency index that characterizes the speed of single-agent learning. Our learning efficiency index generalizes theirs to multiagent settings. Our key finding is that because higher-order uncertainty vanishes faster than first-order uncertainty, the multiagent 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.<sup>2</sup> A notable exception is Harel et al. (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 independently and identically distributed (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. Our analysis assumes a designer-preferred equilibrium selection; remark 1 discusses the importance of this assumption for theorem 2.

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.<sup>3</sup> Extensions of Blackwell's order to multiplayer 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

<sup>2</sup> See, e.g., Vives (1993), Duffie and Manso (2007), Hann-Caruthers, Martynov, and Tamuz (2018), Dasaratha and He (2019), Rosenberg and Vieille (2019), and Liang and Mu (2020).

<sup>3</sup> Azrieli (2014) and Mu et al. (2021) consider a more demanding order that requires the number of signal draws to be uniform across decision problems.

objectives to avoid obtaining a highly conservative ranking.<sup>4</sup> In particular, Lehrer, Rosenberg, and Shmaya (2010) focus on common interest games with utilitarian welfare, while Peşki (2008) compares minmax values in zero-sum games.<sup>5</sup> 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.

## II. Setting

### A. Learning Environment

Throughout the paper, we fix a finite set of agents  $N$ , a finite set of states  $\Theta$ , 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  $\theta$ . We assume that for all agents  $i$  and states  $\theta$ ,  $\mu_i^\theta$  has full support and  $\mu_i^\theta \neq \mu_i^{\theta'}$  for all  $\theta' \neq \theta$ . Note that the joint distribution  $\mu^\theta$  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  $\mathbb{P}_i^t \in \Delta(\Theta \times X^t)$  denote the probability distribution over states and sequences of signal profiles that results when the state  $\theta$  is drawn according to prior  $p_0$  and, conditional on each state  $\theta$ , a sequence  $x^t = (x_\tau)_{\tau=1, \dots, t}$  of signal profiles is generated according to  $t$  independent draws from  $\mu^\theta$ . Agent  $i$ 's observed sequence of private signals is  $x_i^t = (x_{i\tau})_{\tau=1, \dots, t}$ .

### B. Common Learning

Cripps et al.'s (2008) classic result is that in this setting, agents commonly learn the state; that is, both their first-order uncertainty about  $\theta$  and their higher-order uncertainty about other agents' beliefs about  $\theta$  vanishes as  $t$  grows large.

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

<sup>5</sup> Bergemann and Morris (2016) study general games using Bayes-correlated equilibria, which are equivalent to BNE 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 multiagent information structure.



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

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

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

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

where, slightly abusing notation, we also use  $\theta$  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

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

denote the event that  $E$  is *commonly  $p$ -believed* at  $t$ , where  $(B_i^p)^1(E) := B_i^p(E)$  and  $(B_i^p)^k(E) := B_i^p((B_i^p)^{k-1}(E))$  for all  $k \geq 2$ . At  $C_i^p(E)$ , the event  $E$  is  $p$ -believed, the event  $B_i^p(E)$  is  $p$ -believed, and so on. The event  $C_i^p(\theta)$  for  $p$  close to 1 captures that agents have approximate common knowledge of state  $\theta$  (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  $\theta \in \Theta$ ,

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

Common learning is a straightforward consequence of individual learning when agents' private signals in  $\mathcal{I}$  are either independently distributed or perfectly correlated. On the other hand, if  $\mathcal{I}$  displays intermediate levels of correlation, this raises the possibility that when agent  $i$  has seen a history that results in her assigning probability greater than (but close to)  $p$  to state  $\theta$ ,  $i$  may put a significant probability on  $j$  having seen a history that results in a posterior less than  $p$  on  $\theta$ . Nonetheless, Cripps et al. (2008) show that when states and signals are finite, as in the current setting, then every information structure  $\mathcal{I}$  gives rise to common learning.<sup>6</sup>

<sup>6</sup> See sec. VI for a discussion of more general settings.

### III. Multiagent Learning Efficiency

#### A. Speed of Common Learning

While Cripps et al.'s (2008) 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  $\mathbb{P}_i^{\mathcal{I}}(C_i^b(\theta) \mid \theta)$  of approximate common knowledge at all large but finite  $t$ , that is, 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  $\theta$ . Then, for any state  $\theta' \neq \theta$ , one can measure how difficult  $i$  finds it to statistically distinguish  $\theta'$  from  $\theta$  using the Chernoff distance (e.g., Cover and Thomas 1999) between  $i$ 's marginal signal distributions in states  $\theta$  and  $\theta'$ :

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

Here,  $\text{KL}(\nu_i, \mu_i^\theta)$  denotes the KL divergence of  $\nu_i$  relative to  $\mu_i^\theta$ .<sup>7</sup> Observe that any minimizer  $\nu_i$  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  $\mu_i^\theta$  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  $\theta$  and  $\theta'$  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  $\theta$ : for any  $p \in (0, 1)$ , as  $t \rightarrow \infty$ , the probability that  $i$  achieves individual  $p$ -belief of state  $\theta$  goes to 1 exponentially,<sup>8</sup>

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

where the rate of convergence is given by

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

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

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

<sup>8</sup> Here  $o(t)$  denotes a sublinear term, i.e.,  $o(t) = f(t)$  for some function  $f$ , with  $\lim_{t \rightarrow \infty} f(t)/t = 0$ .

to distinguish state  $\theta$  from the state  $\theta'$  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:<sup>9</sup> 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  $\theta$ .

Our first main result is that the rate at which agents commonly learn state  $\theta$  is given by the *multiagent learning efficiency* index

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

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  $\underline{\lambda}^\theta(\mathcal{I})$ , that is,

$$\mathbb{P}_i^{\mathcal{I}}(B_i^p(\theta) \mid \theta) = 1 - \exp[-\underline{\lambda}^\theta(\mathcal{I})t + o(t)], \text{ and} \quad (5)$$

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

The fact that  $\underline{\lambda}^\theta(\mathcal{I})$  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 "Related literature" in sec. I), common  $p$ -belief is a much more demanding requirement than individual  $p$ -belief:  $C_i^p(\theta)$  imposes confidence not only on agents' first-order beliefs about the state but also on their entire infinite hierarchy of higher-order beliefs.<sup>10</sup> On the basis of 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  $\theta$  both tend to 1 at the same exponential rate  $\underline{\lambda}^\theta(\mathcal{I})$ .<sup>11</sup> As we illustrate in section III.B, the key observation 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.

<sup>9</sup> More precisely, they use the index  $\min_{\theta' \in \Theta \setminus \{\theta\}} \max_{x \in [0,1]} -\log \sum_{x_i \in X} \mu_i^\theta(x_i)^x \mu_i^{\theta'}(x_i)^{1-x}$ , which is equal to  $\lambda_i^\theta(\mathcal{I})$  by the variational formula (e.g., Dupuis and Ellis 2011, lemma 6.2.3.f).

<sup>10</sup> Relatedly, Kajii and Morris's (1997) critical path theorem yields a lower bound on the probability of  $C_i^p(\theta)$  relative to the probability of  $B_i^p(\theta)$ , but this result applies only when  $p$  is small ( $p < 1/|N|$ ), reflecting a significant gap between common  $p$ -belief and individual  $p$ -belief when  $p$  is close to 1.

<sup>11</sup> The  $o(t)$  terms can differ across (5) and (6) and can depend on  $p_0$ ,  $p$ , and features of  $\mathcal{I}$  other than  $\underline{\lambda}^\theta(\mathcal{I})$ , but these terms become negligible as  $t \rightarrow \infty$ .

The latter observation is also reflected by the structure of the multi-agent learning efficiency index:  $\underline{\lambda}^\theta(\mathcal{I})$  reduces each information structure  $\mathcal{I}$  to a simple one-dimensional measure that focuses on only the worst-informed agent  $i$  and the state  $\theta'$  that  $i$  finds most difficult to distinguish from the true state  $\theta$  on the basis of 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  $\gamma$  and a correlation parameter  $\rho$ , we have  $\underline{\lambda}^\theta(\mathcal{I}) = \text{KL}((1/2, 1/2), (\gamma, 1 - \gamma))$ ; this is strictly increasing in  $\gamma$  but does not depend on  $\rho$ . 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  $\underline{\lambda}^\theta$  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  $\underline{\lambda}^\theta(\mathcal{I}) > \underline{\lambda}^\theta(\tilde{\mathcal{I}})$ . Then, for each  $p \in (0, 1)$ , there is  $T$  such that for all  $t \geq T$ ,

$$\mathbb{P}_i^t(C_i^p(\theta) \mid \theta) > \mathbb{P}_i^{\tilde{t}}(C_i^p(\theta) \mid \theta).$$

Corollary 1 ranks any two information structures whose learning efficiency indexes are not equal, which holds for generic pairs of information structures.<sup>12</sup> 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, Aways 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  $t$ .

### B. Illustration of Theorem 1

We prove theorem 1 in appendixes B and E. To illustrate the key insight behind the result, consider the binary information structure from example 1 with  $\gamma = 3/5$  and  $\rho = 5/12$ .<sup>13</sup> Thus, the signal probabilities conditional on each state  $\theta$  are as follows:

	$x_1 = \theta$	$x_1 \neq \theta$
$x_2 = \theta$	.25	.35
$x_2 \neq \theta$	.35	.05

<sup>12</sup> Given  $\mathcal{I}$ , the set of information structures  $\tilde{\mathcal{I}}$  such that  $\underline{\lambda}^\theta(\mathcal{I}) \neq \underline{\lambda}^\theta(\tilde{\mathcal{I}})$  holds for all  $\theta$  is open and dense in  $\Delta(X)^\Theta$  endowed with the Euclidean topology.

<sup>13</sup> With these parameter values, agents' signals are negatively correlated conditional on each state, but this feature is not important for our general arguments.

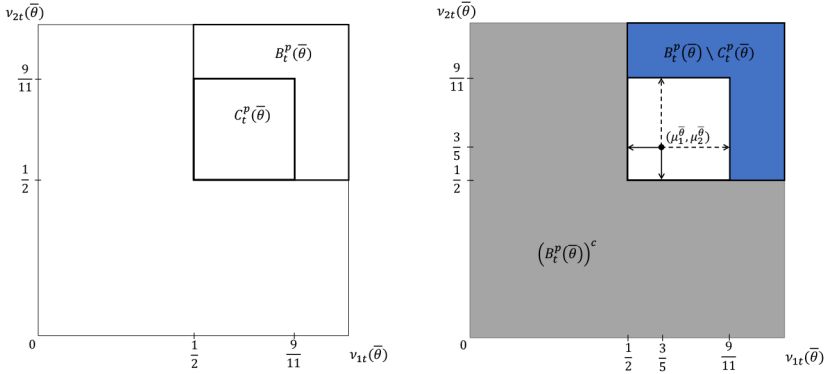


FIG. 1.—*Left*, approximation of  $B_i^p(\bar{\theta})$  and  $C_i^p(\bar{\theta})$  at large  $t$ . *Right*, rate of decay of higher-order belief failures (KL distance of dashed arrows) and first-order belief failures (KL distance of solid arrows) in state  $\bar{\theta}$ .

Fix any  $p \in (0, 1)$ . Let  $\nu_{it} \in \Delta(X_i)$  denote the empirical distribution of agent  $i$ 's signals up to  $t$ ; that is,  $\nu_{it}(\theta) := (1/t)\sum_{\tau=1}^t \mathbf{1}_{\{x_{i\tau}=\theta\}}$  is the frequency with which  $i$ 's realized signals are equal to  $\theta$ . Observe that  $\nu_{it}$  is a sufficient statistic for  $i$ 's (first-order and higher-order) beliefs. Hence, the events  $B_i^p(\theta)$  and  $C_i^p(\theta)$  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  $B_i^p(\theta)$  and  $C_i^p(\theta)$  are approximated by

$$\begin{aligned}
 B_i^p(\theta) &\approx \left\{ \nu_{it}(\theta) \in \left( \frac{1}{2}, 1 \right], \forall i = 1, 2 \right\}, \\
 C_i^p(\theta) &\approx \left\{ \nu_{it}(\theta) \in \left( \frac{1}{2}, \frac{9}{11} \right), \forall i = 1, 2 \right\}.
 \end{aligned}
 \tag{7}$$

The expression for  $B_i^p(\theta)$  is intuitive: at large  $t$ ,  $i$  becomes confident in state  $\theta$  as long as the majority of  $i$ 's signals matches  $\theta$ . To see the idea behind  $C_i^p(\theta)$ , note that for any realized signal frequency  $\nu_{it}(\theta) = \alpha \in (1/2, 1]$  of agent  $i$  and all large enough  $t$ ,  $i$  assigns high probability to  $j$ 's realized signal frequency  $\nu_{jt}(\theta)$  being approximately<sup>14</sup>

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

Observe that (8) exceeds  $1/2$  only if  $\alpha < 9/11$ . Thus, for  $i$  to be confident both in state  $\theta$  and in the fact that  $j$  is confident in state  $\theta$ , we need  $\nu_{it}(\theta) \in$

<sup>14</sup> This holds because  $i$  becomes confident in  $\theta$ , 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.

$(1/2, 9/11)$ . Conversely, if  $\nu_{ii}(\theta) \in (1/2, 9/11)$ , then (8) is itself in  $(1/2, 9/11)$ . This yields the approximation for  $C_i^p(\theta)$ .<sup>15</sup>

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

1. First-order belief failures:  $(B_i^p(\bar{\theta}))^c \approx \{\nu_{ii}(\bar{\theta}) \leq 1/2 \text{ for some } i\}$ .
2. Higher-order belief failures:  $B_i^p(\bar{\theta}) \setminus C_i^p(\bar{\theta}) \approx \{\nu_{ii}(\bar{\theta}) \geq 1/2 \forall i, \nu_{ii}(\bar{\theta}) \geq 9/11 \text{ for some } i\}$ .

Reflecting that common  $p$ -belief is more demanding than individual  $p$ -belief, the second event,  $B_i^p(\bar{\theta}) \setminus C_i^p(\bar{\theta})$ , 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 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. If we let  $\nu_t \in \Delta(X)$  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,

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

That is, as  $t$  grows large, the probability of event  $D$  vanishes exponentially at the 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  $\mathbb{P}_t^X(B_i^p(\bar{\theta}) \setminus C_i^p(\bar{\theta}) \mid \bar{\theta})$  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*).<sup>16</sup> In contrast, as the solid distances illustrate, the probability  $\mathbb{P}_t^X((B_i^p(\bar{\theta}))^c \mid \bar{\theta})$  of first-order belief failures vanishes at rate

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

Crucially, the latter rate is strictly smaller than the former. Thus, as  $t$  grows large, the ratio of  $\mathbb{P}_t^X((B_i^p(\bar{\theta}))^c \mid \bar{\theta})$  to  $\mathbb{P}_t^X(B_i^p(\bar{\theta}) \setminus C_i^p(\bar{\theta}) \mid \bar{\theta})$  explodes.

<sup>15</sup> More precisely, on the basis of these observations, one can show that for all large enough  $t$ , there is an event  $F_t \approx \{\nu_{ii}(\theta) \in (1/2, 9/11), \forall i = 1, 2\}$  such that  $F_t \subseteq B_i^p(\theta)$  and  $F_t \subseteq B_i^p(F_t)$  (i.e.,  $F_t$  is  $p$ -evident), which by Monderer and Samet (1989) implies that  $F_t \subseteq C_i^p(\theta)$ .

<sup>16</sup> The  $\nu \in \Delta(X)$  that attains the infimum  $\inf_{\nu \in B_i^p(\bar{\theta}) \setminus C_i^p(\bar{\theta})} \text{KL}(\nu, \mu^{\bar{\theta}})$  satisfies  $\text{marg}_X \nu = (9/11, 2/11)$  and  $\nu(\cdot \mid x_i) = \mu^{\bar{\theta}}(\cdot \mid x_i)$  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 fig. 1 should be interpreted as depicting KL distances in the space of marginal distributions.

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  $\underline{\lambda}^\theta(\mathcal{I})$ .

Finally, to see the more general idea, note that by (8),  $(1/2, 1/2) = \mathbb{E}[v_{jt} \mid \bar{\theta}, v_{it} = (9/11, 2/11)]$ . Thus, the inequality  $\text{KL}((1/2, 1/2), \mu_j^\theta) < \text{KL}((9/11, 2/11), \mu_i^\theta)$  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  $v_{it} \in \Delta(X_i)$ , we have

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

Moreover, the inequality is strict whenever  $\mu^\theta$  has full support and  $v_{it} \neq \mu_i^\theta$ .

In appendix A, we derive lemma 1 from the chain rule for KL divergence, a central result in information theory. To interpret (9), note that the right-hand side captures how much  $i$ 's signal observations  $v_{it}$  deviate from  $i$ 's theoretical signal distribution  $\mu_i^\theta$  in state  $\theta$ , while the left-hand side quantifies how much  $i$ 's expectation of  $j$ 's observations deviates from  $j$ 's theoretical signal distribution  $\mu_j^\theta$ . 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  $\theta$ ) this estimate is less atypical than  $i$ 's own signal observations.<sup>17</sup> Generalizing the above illustration, we can use this to show that as  $t$  grows large, the event that agents learn  $\theta$  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 Cripps et al. (2008), whereby the map  $v_i \mapsto \mathbb{E}[\mathbb{E}[v_{it} \mid \theta, v_{jt}] \mid \theta, v_{it} = v_i]$  is an  $L^1$ -norm contraction on  $\Delta(X_i)$  if  $\mu^\theta$  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 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.

#### IV. Ranking Information Structures in Coordination Problems

We now return to the question of which information structures are more valuable for coordination. For this, we consider incomplete information

<sup>17</sup> For example, if  $i$  and  $j$ 's signals are independent, then regardless of her own observations,  $i$ 's estimate of  $j$ 's observations is always the theoretical distribution (i.e., the left-hand side 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.

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.

A. *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*  $\mathcal{G}_i(\mathcal{I})$ , where agents’ information is parametrized by the full-support common prior  $p_0 \in \Delta(\Theta)$  and  $t$  draws of signals from  $\mathcal{I}$ . That is, states  $\theta$  and signal sequences  $x^t$  are drawn according to  $\mathbb{P}_i^{\mathcal{I}}$ , 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  $\text{BNE}_i(\mathcal{G}, \mathcal{I})$  denote the set of Bayes-Nash equilibria (BNE) of  $\mathcal{G}_i(\mathcal{I})$ .

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,  $W = (1/|N|)\sum_{i \in N} u_i$ . However, we also allow for objective functions that do not relate to agents’ utilities in any particular way. We assume that in each state  $\theta$ ,  $W$  is maximized by a unique action profile,  $\{a^{\theta, W}\} = \text{argmax}_{a \in A} W(a, \theta)$ .

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

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

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

$$W_i(\mathcal{G}, \mathcal{I}) := \sup_{\sigma_i \in \text{BNE}_i(\mathcal{G}, \mathcal{I})} W_t(\sigma_i, \mathcal{I}) \tag{10}$$

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

We seek to compare the objective values  $W_i(\mathcal{G}, \mathcal{I})$  and  $W_i(\mathcal{G}, \tilde{\mathcal{I}})$  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  $\theta$ .

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  $\theta$ , 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  $\theta$ .

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.<sup>18</sup>

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 common knowledge equilibrium  $a^\theta$  in each state  $\theta$ .

### B. Ranking of Information Structures

Under assumption 1, we now rank information structures  $\mathcal{I}$  and  $\tilde{\mathcal{I}}$  in terms of their objective values  $W_t(\mathcal{G}, \mathcal{I})$  and  $W_t(\mathcal{G}, \tilde{\mathcal{I}})$  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 when assumption 2 is dropped.

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

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

<sup>18</sup> For 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.

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

**THEOREM 2.** Take any information structures  $\mathcal{I}, \tilde{\mathcal{I}}$  with  $\underline{\lambda}(\mathcal{I}) > \underline{\lambda}(\tilde{\mathcal{I}})$ . For every basic game  $\mathcal{G}$  and objective function  $W$  satisfying assumptions 1 and 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 and 2, the learning efficiency index eventually permits a generically complete ranking over information structures: except when the efficiency indexes  $\underline{\lambda}(\mathcal{I})$  and  $\underline{\lambda}(\tilde{\mathcal{I}})$  are equal,  $\mathcal{I}$  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 (apps. D, E), we show that for every BNE sequence  $\sigma_t \in \text{BNE}_t(\mathcal{G}, \mathcal{I})$ ,

$$1 - \underbrace{\sum_{\theta \in \Theta, x' \in X'} \mathbb{P}_t^\mathcal{I}(\theta, x') \sigma_t(a^{\theta, W} \mid x')}_{\text{probability of inefficiency}} \geq \exp[-t\underline{\lambda}(\mathcal{I}) + o(t)] \tag{12}$$

and that (12) holds with equality for some BNE sequence ( $\sigma_t^*$ ). That is, under information structure  $\mathcal{I}$ , the index  $\underline{\lambda}(\mathcal{I})$  is the fastest rate at which inefficiency (i.e., not choosing  $a^{\theta, W}$  at some  $\theta$ ) can vanish in equilibrium. If  $\underline{\lambda}(\mathcal{I}) > \underline{\lambda}(\tilde{\mathcal{I}})$ , 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  $\sum_\theta p_\theta(\theta) W(a^{\theta, W}, \theta)$  faster than does  $W_t(\mathcal{G}, \tilde{\mathcal{I}})$ . Notably, the index  $\underline{\lambda}(\mathcal{I})$  does not depend on the prior  $p_\theta$ . This is because, conditional on each state  $\theta$ , inefficiency vanishes at the exponential rate  $\underline{\lambda}^\theta(\mathcal{I})$ ; thus, at sufficiently large  $t$ , the ex ante probability of inefficiency is driven solely by the state  $\theta$  with the slowest rate of convergence.

As the following example illustrates, this argument relies crucially on our findings that the efficiency index  $\underline{\lambda}(\mathcal{I})$  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

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

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  $k \in \{0, 1, \dots, |N| - 1\}$  other agents also attack. Under utilitarian welfare,  $W = (1/|N|) \sum_{i \in N} u_i$ , the efficient action profiles are  $a^\theta = (1, \dots, 1)$  and  $a^\theta = (0, \dots, 0)$ . Note

that assumptions 1 and 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  $\underline{\lambda}(\mathcal{I})$ , note that common  $p$ -belief of state  $\theta$  is sufficient for coordination on the efficient outcome  $a^\theta$ , as  $a^\theta$  is a strict Nash equilibrium under common knowledge of  $\theta$ . More precisely, if  $p \in (0, 1)$  is sufficiently large, then for every  $t$ , there is a BNE  $\sigma_i^*$  under which  $a^\theta$  is played, conditional on event  $C_i^p(\theta)$ . Under sequence  $(\sigma_i^*)$ , inefficiency vanishes at least as fast as the (ex ante expected) rate of common learning, which is  $\underline{\lambda}(\mathcal{I})$  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.<sup>19</sup> 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 be played only 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  $\underline{\lambda}(\mathcal{I})$  of common learning coincides with the rate of individual learning.<sup>20</sup>

By focusing on data-rich settings, theorem 2 yields some robust implications for information design in coordination games (and other environments satisfying assumptions 1 and 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  $\underline{\lambda}(\mathcal{I})$  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  $W_i(\mathcal{G}, \mathcal{I})$  in (10) considers the best BNE. Thus, 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. While the assumption

<sup>19</sup> See, e.g., Oyama and Takahashi (2020) for systematic analysis of this issue.

<sup>20</sup> When 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 app. 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  $i$ .

of designer-preferred equilibrium selection is also common in the literature on information design (“Related literature” in sec. I), some notable exceptions (e.g., Morris, Oyama, and Takahashi 2020) assume adversarial equilibrium selection, and the latter assumption could reverse the ranking in theorem 2.<sup>21</sup> At the same time, appendix G (apps. G–I are available online) 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 that agents observe one signal draw from each information structure and show that a generalization of Blackwell’s single-agent garbling condition characterizes when  $W_t(\mathcal{G}, \mathcal{I})$  exceeds  $W_t(\mathcal{G}, \tilde{\mathcal{I}})$  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.<sup>22</sup> Both points i and ii rely on agents observing sufficiently many signal draws: when  $t = 1$ , many information structures are incomparable even when focusing on 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 and 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  $\mathcal{I}$  versus  $\tilde{\mathcal{I}}$  can in general still be nonnegligible at  $T$ .<sup>23</sup>

## V. Discussion

### A. 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

<sup>21</sup> For example, if (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), then 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  $t$ .

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

<sup>23</sup> Relatedly, 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  $T$ .

the ordinal implications highlighted following theorem 2, here the cardinal value that  $\underline{\lambda}(\mathcal{I})$  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:

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

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 and 2 and any finite set  $\mathbb{I}$  of information structures. Whenever the budget  $\kappa$  is sufficiently large (i.e., information is sufficiently cheap), the designer's problem (13) simplifies to

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

Thus, the optimal information structure can be determined solely on the basis of the learning efficiency index and per-sample cost, and the solution is robust across all games and objectives satisfying assumptions 1 and 2. On the basis of this observation, one can explore properties of the optimal information structure, depending on the nature of the cost function  $c$ .

### B. Convergence of Belief Hierarchies

Since Rubinstein (1989), there is a discussion in the literature about which topologies over belief hierarchies are appropriate for measuring proximity to common knowledge. 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  $\underline{\lambda}^\theta(\mathcal{I})$  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  $\tau_i^k$  for different  $k$  (e.g., Brandenburger and Dekel 1993).<sup>24</sup> 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

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

$\tau_i(\theta) \in Z_i$  denote  $i$ 's belief hierarchy when there is common certainty of state  $\theta$ .

Let  $\rho_i^{\text{product}}$  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  $\rho^k$  denotes the Prokhorov metric over  $k$ th order beliefs (we endow  $\theta$  with the discrete metric). 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 et al. 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  $\theta$  is given by  $\underline{\lambda}^\theta(\mathcal{I})$ : for all sufficiently small  $\varepsilon > 0$ , we have

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

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

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. Indeed, a key difference between these topologies is that convergence to common certainty under the uniform weak topology requires common  $p$ -belief (with  $p \rightarrow 1$ ), whereas this is not the case under the product topology.<sup>25</sup> However, we have seen that in our setting, the rate of convergence to common  $p$ -belief coincides with the rate of individual learning. Thus, corollary 3 is a result of the fact that both these topologies agree on the metric over first-order beliefs.<sup>26</sup>

### C. Higher-Order Expectations

Beyond its use in this 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).

<sup>25</sup> One manifestation of this difference is that Rubinstein's (1989) email game generates approximate common knowledge in the sense of  $\rho^{\text{product}}$  but not in the sense of  $\rho^{\text{uniform}}$ .

<sup>26</sup> The result extends to other topologies considered in the literature, e.g., the strategic topology (Dekel, Fudenberg, and Morris 2006), which is in between the product and uniform weak topologies.

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  $\pi(\cdot | t_i) \in \Delta(T)$  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  $\pi(\cdot | t_i)$  a sequence of higher-order expectations about other agents' types. In particular,  $\mathbb{E}_{t_i}[t_j] := \sum_{t_j \in T_j} \pi(t_j | t_i) \delta_{t_j} \in \Delta(T_j)$  is  $t_i$ 's expectation of  $j$ 's type,  $\mathbb{E}_{t_i}[\mathbb{E}_{t_j}[t_k]] := \sum_{t_j \in T_j, t_k \in T_k} \pi(t_j | t_i) \pi(t_k | t_j) \delta_{t_k} \in \Delta(T_k)$  is  $t_i$ '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 that<sup>27</sup>

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

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

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

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 nonasymptotic finding may be useful for conducting comparative statics with respect to coordination motives away from the limit.

## VI. Conclusion

This paper conducted a comparison of multiagent information structures in a learning setting where players have access to rich data. We showed that

<sup>27</sup> See the proof of his proposition 6. A precursor in the mathematics literature is Amemiya and Andó (1965).

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 and state spaces are finite and signals are i.i.d. across draws. This allowed us to build on Cripps et al.'s (2008) result that this setting always gives rise to common learning. With infinite signals, Cripps et al. (2008) exhibit an example in which common learning fails even though individual learning is successful. At the same time, there are other natural infinite-signal (finite-state) settings—in particular, Gaussian signal structures—that do give rise to common learning (for general sufficient conditions for common learning under infinite signals, see Faingold and Tamuz 2022). Appendix H analyzes such Gaussian environments and shows that common and individual learning again occur at the same exponential rate. In contrast, it is known that with continuous states, common learning fails under a broad class of signal distributions.<sup>28</sup>

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. Appendix I analyzes when the learning efficiency index  $\underline{\lambda}(\mathcal{I} \times \tilde{\mathcal{I}})$  of alternating draws from  $\mathcal{I}$  and  $\tilde{\mathcal{I}}$  is greater or less than the sum  $\underline{\lambda}(\mathcal{I}) + \underline{\lambda}(\tilde{\mathcal{I}})$  of their separate indexes, 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 et al. 2013). We leave the analysis of such settings for future work, in particular, the question of 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 et al. (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.

<sup>28</sup> As noted by Cripps et al. (2008), if both states and signals follow a Gaussian distribution, then common learning fails by an analogous logic as in the global games literature. Dogan (2018) shows that with continuous states, common learning fails under mild conditions even when signals are finite.



## Appendix A

## Preliminaries

## A1. Preliminary Definitions

The following will be used throughout the appendix. As in Cripps et al. (2008), 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

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

As Cripps et al. (2008) observed, if agent  $i$ 's empirical signal distribution at  $t$  is  $\nu_{it}$ , then conditional on state  $\theta$ ,  $i$ '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  $d < \underline{\lambda}^\theta(\mathcal{I})$  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  $\mu^\theta$  has full support on  $X$  in all states  $\theta$ . We call  $\mathcal{I}$  *public* if signals are perfectly correlated across agents.<sup>29</sup>

## A2. 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  $\nu_i, \mu_i^\theta$ , and the marginals on  $X_j$  are  $\nu_i M_{ij}^\theta, \mu_j^\theta$ , respectively.<sup>30</sup>

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

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

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

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

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

<sup>29</sup> That is,  $X_i = X_j$  for all  $i, j$ , and for each  $x \in X$  and  $\theta$ ,

$$\mu^\theta(x) = \begin{cases} \mu_i^\theta(x_i) & \text{if } x_i = x_j \text{ for all } i, j, \\ 0 & \text{otherwise.} \end{cases}$$

<sup>30</sup> Note that  $m'$  is the marginal of  $\mu^\theta$  over  $X_i \times X_j$ .

To show the strict inequality, suppose that  $\nu_i \neq \mu_i^\theta$  and  $\mu^\theta$  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 assumption on  $\mu^\theta$ . By Gibbs's 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  $\text{KL}(\nu_i, \mu_i^\theta) > \text{KL}(\nu_i M_{ij}^\theta, \mu_j^\theta)$ . QED

### A3. Other Preliminary Lemmas

Let  $\|\cdot\|$  denote the sup norm for finite-dimensional real vectors. The following result is proved by Cripps et al. (2008, their lemma 3) using a concentration inequality argument:

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  $x'_i$ ,

$$\mathbb{P}_t^T(\{\|\nu_{it} M_{ij}^\theta - \nu_{jt}\| < \varepsilon, \forall j \neq i\} | x'_i, \theta) > q.$$

Let  $F_{-i}(\theta, d) := \cap_{j \neq i} F_{jt}(\theta, d)$ . 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  $\Theta$ ,  $\theta \in \Theta$ ,  $p \in (0, 1)$ , and  $d \in (0, \min_{i \in N, \theta \in \Pi_i(\theta)} d(\mu_i^\theta, \mu_i^\theta))$ . Assume that  $\mu^\theta$  has full support. There exists  $T$  such that for all  $i \in N$  and  $t \geq T$ ,

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

*Proof.* The proof proceeds through two claims.

#### A3.1. Claim 1

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

$$\text{KL}(\nu_{it}, \mu_i^\theta) \leq d + \kappa \implies \mathbb{P}_t^T(F_{-i}(\theta', d) | x'_i, \theta') \geq \sqrt{p}.$$

*Proof.* Lemma 1 implies that for all  $j \neq i$ ,  $\nu_i \in \Delta(X_j)$ , 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 right-hand side is strict when  $\nu_i \neq \mu_i^\theta$  (by lemma 1), and the second inequality on the right-hand side 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  $j \neq i$ ,  $\nu_i \in \Delta(X_j)$ , and  $\theta' \in \Theta$ ,

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

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  $j \neq i$ ,  $v_i \in \Delta(X_i)$ , and  $\theta' \in \Theta$ ,

$$\text{KL}(v_i, \mu_i^{\theta'}) \leq d + \kappa \implies \text{KL}(v_i M_{ij}^{\theta'}, \mu_j^{\theta'}) \leq d - \frac{\eta}{2}.$$

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

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

Combined with lemma A.1, this yields the desired conclusion. QED

### A3.2. Claim 2

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

$$\text{KL}(v_{it}, \mu_{it}^{\theta'}) \leq d \implies \mathbb{P}_i^T(\{\theta' \in \Pi_i(\theta) : \text{KL}(v_{it}, \mu_{it}^{\theta'}) \leq d + \kappa\} \mid x_i^t) \geq \sqrt{p}.$$

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

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

$$\begin{aligned} \log \mathbb{P}_i^T(\theta' \mid x_i^t) &\leq \log \frac{\mathbb{P}_i^T(\theta' \mid x_i^t)}{\mathbb{P}_i^T(\theta \mid x_i^t)} = \log \frac{p_0(\theta')}{p_0(\theta)} + t \sum_{x \in X_i} v_{it}(x_i) \log \frac{\mu_{it}^{\theta'}(x_i)}{\mu_{it}^{\theta}(x_i)} \\ &= \log \frac{p_0(\theta')}{p_0(\theta)} + t(\text{KL}(v_{it}, \mu_{it}^{\theta}) - \text{KL}(v_{it}, \mu_{it}^{\theta'})) \\ &\leq \log \frac{p_0(\theta')}{p_0(\theta)} - t\kappa. \end{aligned}$$

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

$$\text{KL}(v_{it}, \mu_{it}^{\theta'}) \leq d \implies \mathbb{P}_i^T(\theta' \mid x_i^t) < \frac{1 - \sqrt{p}}{|\Theta|},$$

proving claim 2. QED

Finally, to prove lemma A.2, let  $T = \max\{T', T''\}$ , with  $T'$  and  $T''$  as found in claims 1 and 2. Then, whenever  $t \geq T$  and  $\text{KL}(v_{it}, \mu_{it}^{\theta}) \leq d$ , we have

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

where the second inequality uses claim 1 and the last inequality uses claim 2. QED

**Appendix 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 app. sec. A1). Appendix E extends the proof to general information structures.

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

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

Take any  $d \in (0, \underline{\lambda}^\theta(\mathcal{I}))$ . 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)  $F_t(\theta, d) \subseteq B_t^p(F(\theta, d))$ , that is,  $F_t(\theta, d)$  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,

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

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

We next establish that

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

Take  $i \in N$  and  $\theta' \neq \theta$  such that  $d(\mu_i^\theta, \mu_i^{\theta'}) = \underline{\lambda}^\theta(\mathcal{I})$ . Take any  $d > d(\mu_i^\theta, \mu_i^{\theta'})$ . Then there is  $v_i \in \Delta(X_i)$  such that  $\text{KL}(v_i, \mu_i^\theta) = \text{KL}(v_i, \mu_i^{\theta'}) < d$ . Hence, for some  $v'_i$  close to  $v_i$ ,

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

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

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

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

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

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

## Appendix C

### Ranking of Information Structures without Assumption 2

This appendix generalizes the ranking in theorem 2 when assumption 2 is dropped; that is, 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 indexes that account for the presence of equivalent states for some players.

Formally, given any objective function  $W$ , define a partition  $\Pi_i^W$  over  $\Theta$  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,  $\Pi_i^W$  divides  $\Theta$  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  $\Theta$ , we define the learning efficiency index<sup>31</sup>

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

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 app. sec. A1).

**THEOREM C.1.** Fix any collection  $\Pi = (\Pi_i)_{i \in N}$  of partitions over  $\Theta$ . Take any information structures  $\mathcal{I}$  and  $\tilde{\mathcal{I}}$ , each of which is either fully private or public, with  $\underline{\lambda}(\mathcal{I}, \Pi) > \underline{\lambda}(\tilde{\mathcal{I}}, \Pi)$ . 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. Using the generalized learning efficiency indexes  $\underline{\lambda}(\cdot, \Pi)$ , 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  $\Pi$  of equivalent states. The proof of theorem C.1 is in appendix D.

<sup>31</sup> Slightly abusing notation, we set the index to be  $\infty$  when  $\Pi$  is degenerate (i.e.,  $\Pi_i(\theta) = \Theta$  for all  $i$ ).

**Appendix 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  $\theta$  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.

*D1. 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  $a^\theta$  in state  $\theta$ ) as  $t$  grows large. The following result provides a lower bound on this probability for arbitrary sequences of strategy profiles  $(\sigma_t)$ :

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

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

*Proof.* Pick  $i$ ,  $\theta$ , 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  $(\sigma_t)$  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  $\tilde{\sigma}_{it}(a_i^{\theta'} | x_i^t) \geq \sigma_{it}(a_i^{\theta'} | x_i^t)$ ; and
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  $(\sigma_{it})$  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  $(\sigma_{it}^*)$  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  $\nu_{it}$  is the empirical signal distribution associated with  $x_i^t$ . Note that  $\sigma_{it}^*$  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  $t$ ,

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

Hence,

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

where the first inequality follows from the construction of  $(\tilde{\sigma}_{ii})$  and the second inequality uses (17). The last line is equal to  $-d(\mu_i^\theta, \mu_i^{\theta'}) = -\underline{\lambda}(\mathcal{I}, \Pi^W)$ , 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),<sup>32</sup> that is,

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

This implies that

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

as claimed. QED

Under assumption 1, the following result provides an upper bound on the probability of inefficient play under some equilibrium sequence  $(\sigma_t)$ :

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

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \left( 1 - \sum_{x_i^t \in X_i^t} \mathbb{P}_t^{\mathcal{I}}(x_i^t \mid \theta) \sigma_t(a_i^\theta \mid x_i^t) \right) \leq -\underline{\lambda}(\mathcal{I}, \Pi^W).$$

*Proof.* Take  $p \in (0, 1)$  sufficiently close to 1 such that for all  $i$  and  $\theta$ , choosing  $a_i^\theta$  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,  $a_i^\theta$  is the unique maximizer of  $u_i(\cdot, a_{-i}^{\theta'}, \theta')$  for each  $\theta' \in \Pi_i^W(\theta)$ .

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

<sup>32</sup> This in turn follows from a simple application of Sanov's theorem.

We show that there exists  $T$  such that for any  $t > T$ , there is a BNE  $\sigma_t$  of  $\mathcal{G}_t(\mathcal{I})$ , with  $\sigma_{ii} \in \Sigma_{ii}(d)$  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  $i, \theta$ , and  $t \geq T$ . Thus, for all  $t \geq T$ , each agent  $i$ 's best response against any strategy profile in  $\prod_{j \neq i} \Sigma_{ji}(d)$  must be in  $\Sigma_{ii}(d)$ , because whenever  $\text{KL}(v_{ii}, \mu_i^\theta) \leq d$ , 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  $\prod_i \Sigma_{ii}(d)$  (which is convex), we obtain a BNE  $\sigma_t$  of  $\mathcal{G}_t(\mathcal{I})$  such that  $\sigma_{ii} \in \Sigma_{ii}(d)$  for every  $i$ . Next, suppose that  $\mathcal{I}$  is public. In this case, all players' posteriors coincide, that is,  $\mathbb{P}_i^{\mathcal{I}}(\cdot | x^t) = \mathbb{P}_j^{\mathcal{I}}(\cdot | x^t)$  for all  $i, j$ , and  $t$ . Moreover,  $\text{KL}(v_{ii}, \mu_i^\theta) \leq d \Leftrightarrow \text{KL}(v_{jj}, \mu_j^\theta) \leq d$  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

$$\text{KL}(v_{ii}, \mu_i^\theta) \leq d \implies \mathbb{P}_i^{\mathcal{I}}(\{\theta' \in \cap_j \Pi_j(\theta)\} | x^t) \geq p$$

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

The above implies that there is a sequence of BNE  $(\sigma_t)$  such that for all  $\theta$ , we have that as  $t \rightarrow \infty$ ,

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

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

D2. 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  $\underline{\lambda}(\mathcal{I}, \Pi) > \underline{\lambda}(\tilde{\mathcal{I}}, \Pi)$ . 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  $\sigma_t$  of  $\mathcal{G}_t(\mathcal{I})$  and  $\tilde{\sigma}_t$  of  $\mathcal{G}_t(\tilde{\mathcal{I}})$ , and all  $\theta \in \Theta$ ,

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

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

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

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

which by (18) implies



$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \sum_{\theta} p_{\theta}(\theta) \left( W(a^{\theta}, \theta) - \sum_{x'} \mathbb{P}_t^{\mathcal{I}}(x' | \theta) \sigma_t(a^{\theta} | x') W(a, \theta) \right) \leq -\underline{\lambda}(\mathcal{I}, \Pi). \quad (20)$$

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

$$\begin{aligned} -\underline{\lambda}(\tilde{\mathcal{I}}, \Pi) &\leq \liminf_{t \rightarrow \infty} \max_{\theta} \frac{1}{t} \log \left( 1 - \sum_{\tilde{x}'} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}' | \theta) \tilde{\sigma}_t(a^{\theta} | \tilde{x}') \right) \\ &\leq \liminf_{t \rightarrow \infty} \frac{1}{t} \log \sum_{\theta} p_{\theta}(\theta) \left( 1 - \sum_{\tilde{x}'} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}' | \theta) \tilde{\sigma}_t(a^{\theta} | \tilde{x}') \right), \end{aligned}$$

which by (19) implies

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \sum_{\theta} p_{\theta}(\theta) \left( W(a^{\theta}, \theta) - \sum_{\tilde{x}'} \mathbb{P}_t^{\tilde{\mathcal{I}}}(\tilde{x}' | \theta) \tilde{\sigma}_t(a^{\theta} | \tilde{x}') W(a, \theta) \right) \geq -\underline{\lambda}(\tilde{\mathcal{I}}, \Pi). \quad (21)$$

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

## Appendix E

### Proofs of Theorems 1 and 2 (General Case)

#### E1. Overview

In this section, we extend the proofs of theorems 1 and 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  $\mu^{\theta}$  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  $\theta$ . Let  $X^{\theta} \subseteq X$  denote the support of  $\mu^{\theta}$ . Conditional on state  $\theta$ , define  $H_i^{\theta} = (h_i^{\theta}(x))_{x \in X^{\theta}}$  to be agent  $i$ 's information partition of  $X^{\theta}$  based on observing her own private signal; that is,

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

For any distribution  $\nu \in \Delta(X^{\theta})$  and any partition  $H$  of  $X^{\theta}$ , 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  $(\nu_t)_{H_t^i}$  can be identified with  $i$ 's empirical distribution  $\nu_{it}$ . 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  $H_i^{\theta}$  such that  $i \in S$ . For any joint empirical signal distribution  $\nu_t$ , 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:

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

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  $\mu^{\theta}$  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 and 2 is the following result, which we prove in appendix E.1.

PROPOSITION E.1. Take any  $d \in (0, \underline{\lambda}^\theta(\mathcal{I}))$  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

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

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

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  $d \in (0, \underline{\lambda}^\theta(\mathcal{I}))$  and  $\varepsilon \in (0, d)$ . Then for all  $p \in (0, 1)$  and large enough  $t$ , the events  $F_i(\theta, d, \varepsilon_1, \dots, \varepsilon_{|N|})$  constructed in proposition E.1 satisfy

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

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

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

Since this holds for all  $d < \underline{\lambda}^\theta(\mathcal{I})$ , (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  $d \in (0, \underline{\lambda}(\mathcal{I}))$  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  $(\sigma_i)$  such that for all large enough  $t$  and each  $\theta$ , we have  $\sigma_i(a^\theta \mid x^t) = 1$  at all signal sequences  $x^t \in F_i(\theta, d, \varepsilon_1, \dots, \varepsilon_{|N|})$ . Thus,

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

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

## E2. Proof of Proposition E.1

### E2.1. 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  $\nu_i = \text{marg}_X \nu$ , define distribution  $\nu M_i^\theta \in \Delta(X^\theta)$  by

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

When the joint empirical signal distribution is  $\nu$ , then  $\nu_i M_i^\theta$  is  $i$ 's expectation of this joint distribution conditional on state  $\theta$  and on observing  $\nu_i$ .

LEMMA E.1. Take any  $\nu \in \Delta(X^\theta)$ ,  $i \in N$ , and  $S \subseteq N$ . Then  $\text{KL}((\nu M_i^\theta)_{H_i^\theta}, \mu_{H_i^\theta}^\theta) \leq \text{KL}(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta)$ . Moreover, the inequality is an equality only if  $\nu_{H_i^\theta}(\cdot|h) = \mu_{H_i^\theta}^\theta(\cdot|h)$  for every  $h \in H_i^\theta \wedge H_S^\theta$  such that  $\nu_{H_i^\theta \wedge H_S^\theta}(h) > 0$ .

*Proof.* To show the inequality, first note that

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

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

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

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

For the ‘‘moreover’’ part, suppose that  $\text{KL}((\nu M_i^\theta)_{H_i^\theta}, \mu_{H_i^\theta}^\theta) = \text{KL}(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta)$ . Then, by (23) and (24), for every  $h \in H_S^\theta$  such that  $(\nu M_i^\theta)_{H_i^\theta}(h) > 0$ , 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  $h \in H_i^\theta \wedge H_S^\theta$  such that  $\nu_{H_i^\theta \wedge H_S^\theta}(h) > 0$ . QED

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

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

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

*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  $\text{KL}((\nu M_i^\theta)_{H_i^\theta}, \mu_{H_i^\theta}^\theta) < d$ , as the left-hand side of this inequality is continuous in  $\nu$  and  $\Delta(X^\theta)$  is compact.

To show the latter inequality, note that lemma E.1 implies  $\text{KL}((\nu M_i^\theta)_{H_i^\theta}, \mu_{H_i^\theta}^\theta) \leq \text{KL}(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta) \leq d$ . Thus, we can focus on the case in which  $\text{KL}((\nu M_i^\theta)_{H_i^\theta}, \mu_{H_i^\theta}^\theta) = \text{KL}(\nu_{H_i^\theta}, \mu_{H_i^\theta}^\theta)$ . In this case,

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

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

## E.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

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

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

Next, choose some  $\varepsilon_{|N|-1} \in (0, \rho_{|N|-1})$  and proceed inductively in the same manner. In particular, once we have constructed  $\varepsilon_{k+1}$ , 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

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

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

This yields a sequence

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

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

We now show that this sequence is as required by proposition E.1. By the same argument as in the proof of lemma A.2, for any  $p \in (0, 1)$  and sufficiently large  $t$ , we have  $F_t(\theta, d, \varepsilon_1, \dots, \varepsilon_{|N|}) \subseteq B_i^\theta(\theta)$ . Thus, it suffices to show that for any  $p \in (0, 1)$ ,  $i \in N$ , and  $S \subseteq N$ , there exists  $T$  such that

$$\mathbb{P}_t^{\mathcal{X}}(\{\text{KL}((\nu_t)_{H_i^\theta}, \mu_{H_i^\theta}^\theta) \leq d - \varepsilon_{|S|}\} \mid x_t^i, \theta) \geq p \quad (25)$$

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  $H_S^\theta$  is coarser than  $H_i^\theta$ . Hence, for any  $t \geq 1$  and signal sequence  $x^t$  with corresponding empirical distribution  $\tilde{\nu}_t \in \Delta(X^\theta)$ , we have

$$\mathbb{P}_t^{\mathcal{X}}(\{\nu_t \in \Delta(X^\theta) : (\nu_t)_{H_S^\theta} = (\tilde{\nu}_t)_{H_S^\theta}\} \mid x_t^i, \theta) = 1.$$

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

$$\mathbb{P}_t^{\mathcal{X}}(\{\text{KL}((\nu_t)_{H_i^\theta}, \mu_{H_i^\theta}^\theta) \leq d - \varepsilon_{|S|}\} \mid x_t^i, \theta) = 1 > p,$$

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

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

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

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

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  $\tilde{v}_t$ , we have

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

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

## Appendix F

### Proofs for Section V

#### F1. Proof of Corollary 2

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

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

and that (28) holds with equality for some BNE sequence  $(\sigma_\kappa^*)$ . 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 on the right-hand side 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  $\sum_\theta p_\theta(\theta) W(a^{\theta, W}, \theta)$  and the designer's value under each information structure  $\mathcal{I}$  takes the form  $\exp[-\kappa(\underline{\lambda}(\mathcal{I})/c(\mathcal{I})) + o(\kappa)]$ .

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

#### F2. Proof of Corollary 3

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

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

The convergence under the uniform weak topology cannot be slower than  $\underline{\lambda}^\theta(\mathcal{I})$ . To see this, fix any  $\varepsilon > 0$ . Note that the proof of proposition 6 in Chen et al. (2010) implies that the  $\varepsilon$ -ball around  $\tau_i(\theta)$  consists of all belief hierarchies for player  $i$  that have common  $(1 - \varepsilon)$ -belief on  $\theta$ . Thus,

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

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

## References

- Acemoglu, D., V. Chernozhukov, and M. Woldic. 2016. "Fragility of Asymptotic Agreement under Bayesian Learning." *Theoretical Econ.* 11 (1): 187–225.
- Amemiya, I., and T. Andô. 1965. "Convergence of Random Products of Contractions in Hilbert Space." *Acta Sci. Math. (Szeged)* 26 (3–4): 239–44.
- 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): 415–23.
- Basu, P., K. Chatterjee, T. Hoshino, and O. Tamuz. 2020. "Repeated Coordination with Private Learning." *J. Econ. Theory* 190:105106.
- Bergemann, D., and S. Morris. 2016. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Econ.* 11 (2): 487–522.
- . 2019. "Information Design: A Unified Perspective." *J. Econ. Literature* 57 (1): 44–95.
- Blackwell, D. 1951. "Comparison of Experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, edited by J. Neyman, 93–102. Berkeley: Univ. California Press.
- Brandenburger, A., and E. Dekel. 1993. "Hierarchies of Beliefs and Common Knowledge." *J. Econ. Theory* 59 (1): 189–98.
- Brooks, B., A. Frankel, and E. Kamenica. 2021. "Information Hierarchies." *Econometrica* 90 (5): 2187–214.
- Carlsson, H., and E. Van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica* 61 (5): 989–1018.
- Chen, Y.-C., A. Di Tillio, E. Faingold, and S. Xiong. 2010. "Uniform Topologies on Types." *Theoretical Econ.* 5 (3): 445–78.
- Chernoff, H. 1952. "A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the Sum of Observations." *Ann. Math. Statist.* 23 (4): 493–507.
- Cover, T. M., and J. A. Thomas. 1999. *Elements of Information Theory*. Hoboken, NJ: Wiley.
- Cripps, M. W., J. C. Ely, G. J. Mailath, and L. Samuelson. 2008. "Common Learning." *Econometrica* 76 (4): 909–33.
- . 2013. "Common Learning with Intertemporal Dependence." *Internat. J. 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 Econ.* 1:275–309.
- Dembo, A., and O. Zeitouni. 2010. *Large Deviations Techniques and Applications*. Berlin: Springer.
- Dogan, E. 2018. "Common Learning on a Continuum State Space." Working paper.
- Duffie, D., and G. Manso. 2007. "Information Percolation in Large Markets." *A.E.R.* 97 (2): 203–9.
- Dupuis, P., and R. S. Ellis. 2011. *A Weak Convergence Approach to the Theory of Large Deviations*, vol. 902. New York: Wiley.
- 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 Econ. Behavior* 30 (1): 44–63.

- Hann-Caruthers, W., V. V. Martynov, and O. Tamuz. 2018. "The Speed of Sequential Asymptotic Learning." *J. Econ. Theory* 173:383–409.
- Harel, M., E. Mossel, P. Strack, and O. Tamuz. 2021. "Rational Groupthink." *Q.J.E.* 136 (1): 621–68.
- Hirshleifer, J. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *A.E.R.* 61 (4): 561–74.
- Kajii, A., and S. Morris. 1997. "The Robustness of Equilibria to Incomplete Information." *Econometrica* 65 (6): 1283–309.
- Kamenica, E. 2019. "Bayesian Persuasion and Information Design." *Ann. Rev. Econ.* 11:249–72.
- Lehmann, E. L., and J. P. Romano. 2006. *Testing Statistical Hypotheses*. New York: Springer Science and Business Media.
- Lehrer, E., D. Rosenberg, and E. Shmaya. 2010. "Signaling and Mediation in Games with Common Interests." *Games and Econ. Behavior* 68 (2): 670–82.
- Liang, A. 2019. "Games of Incomplete Information Played by Statisticians." Working paper.
- Liang, A., and X. Mu. 2020. "Complementary Information and Learning Traps." *Q.J.E.* 135 (1): 389–448.
- Lipman, B. L. 2003. "Finite Order Implications of Common Priors." *Econometrica* 71 (4): 1255–67.
- Monderer, D., and D. Samet. 1989. "Approximating Common Knowledge with Common Beliefs." *Games and Econ. Behavior* 1 (2): 170–90.
- 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." *A.E.R.* 88 (3): 587–97.
- . 2002. "Social Value of Public Information." *A.E.R.* 92 (5): 1521–34.
- Moscarini, G., and L. Smith. 2002. "The Law of Large Demand for Information." *Econometrica* 70 (6): 2351–66.
- 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.
- Peşki, M. 2008. "Comparison of Information Structures in Zero-Sum Games." *Games and Econ. Behavior* 62 (2): 732–35.
- Rosenberg, D., and N. Vieille. 2019. "On the Efficiency of Social Learning." *Econometrica* 87 (6): 2141–68.
- Rubinstein, A. 1989. "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge.'" *A.E.R.* 79 (3): 385–91.
- Samet, D. 1998. "Iterated Expectations and Common Priors." *Games and Econ. Behavior* 24 (1–2): 131–41.
- Steiner, J., and C. Stewart. 2011. "Communication, Timing, and Common Learning." *J. Econ. Theory* 146 (1): 230–47.
- Sugaya, T., and Y. Yamamoto. 2020. "Common Learning and Cooperation in Repeated Games." *Theoretical Econ.* 15 (3): 1175–219.
- Vives, X. 1993. "How Fast Do Rational Agents Learn?" *Rev. Econ. Studies* 60 (2): 329–47.
- Weinstein, J., and M. Yildiz. 2007. "A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements." *Econometrica* 75 (2): 365–400.