

# Games of Incomplete Information Played by Statisticians\*

Annie Liang<sup>†</sup>

January 1, 2016

## Abstract

The common prior assumption is a convenient restriction on beliefs in games of incomplete information, but conflicts with evidence that agents publicly disagree in many economic environments. This paper proposes a foundation for heterogeneous beliefs in games, in which disagreement arises not from different information, but from different interpretations of common information. I model players as statisticians who infer an unknown parameter from data. Players know that they may use different inference rules (and, therefore, may disagree about the distribution of payoffs), but have common certainty in the predictions of a class of inference rules. Using this framework, I study the robustness of solutions to a relaxation of the common prior assumption. The main results characterize which rationalizable actions and which Nash equilibria persist given finite quantities of data, and provide a lower bound on the quantity of data needed to learn these solutions. I suggest a new criterion for equilibrium selection based on statistical complexity—solutions that are “hard to learn” are selected away.

## 1 Introduction

In games with a payoff-relevant parameter, players’ beliefs about this parameter, as well as their beliefs about opponent beliefs about this parameter, are important for predictions of play. The standard approach to restricting the space of beliefs

---

\*See current version at: <http://scholar.harvard.edu/aliang/publications/games-incomplete-information-played-statisticians>.

<sup>†</sup>Department of Economics, Harvard University. Email: [annieliang@fas.harvard.edu](mailto:annieliang@fas.harvard.edu). I am deeply grateful to Drew Fudenberg, Jerry Green, David Laibson, and Sendhil Mullainathan for their continual advice and support. Drew Fudenberg’s extensive guidance was essential in the development of this paper. Comments from Jetlir Duraj, Siddharth George, Ben Golub, Scott Kominers, Jonathan Libgober, Harry Pei, Andrei Shleifer, Dov Samet, Tomasz Strzalecki, Anton Tsoy, and Muhamet Yildiz improved this paper and are gratefully acknowledged.

assumes that players share a common prior distribution.<sup>1</sup> This assumption is known to have strong implications, including that beliefs that are commonly known must be identical (Aumann 1976), and repeated communication of beliefs will eventually lead to agreement (Geanakoplos & Polemarchakis 1982). These properties conflict not only with considerable empirical evidence of public and persistent disagreement,<sup>2</sup> but also with the more basic, day-to-day, experience that people sometimes come to different conclusions given the same information.

As a consequence, the following questions arise: When is disagreement a feature of agent’s beliefs, and how can this disagreement be predicted from the primitives of the economic environment? Can we relax the common prior assumption to accommodate (commonly known) disagreement in a structured way? Finally, when are strategic predictions robust to relaxations of the common prior assumption?

Towards the first questions of modeling and predicting disagreement, I propose a reformulation of incomplete information in which agents form beliefs by learning from data. I take *data* be a random sequence of observations, drawn i.i.d. from an exogenous distribution  $P$ , and define an *inference rule* to be any map from possible datasets into distributions over the parameter space. (For example, we can think of data as historical stock prices, and inference rules as maps from possible time-series of stock returns to distributions over returns next period.)

This perspective on beliefs provides a way to rationalize disagreement—in the absence of a “privileged” or “correct” inference rule, different interpretations of common data is not only possible, but even natural.<sup>3</sup> The key restriction I impose to structure this approach is that while agents may learn from data using different inference rules, they have common certainty in the predictions of a *family* of plausible inference rules.<sup>4</sup> This assumption is referred to as *common inference*. In the main part of the paper, I additionally assume a condition on the family of inference rules (uniform consistency<sup>5</sup>) that implies that agents commonly learn the true parameter

---

<sup>1</sup>The related, stronger, notion of rational expectations assumes moreover that this common prior distribution is in fact the “true” distribution shared by the modeler.

<sup>2</sup>In financial markets, agents publicly disagree in their interpretations of earnings announcements (Kandel & Pearson 1995), valuations of financial assets (Carlin, Kogan & Lowery 2013), forecasts for inflation (Mankiw, Reis & Wolfers 2004), forecasts for stock movements (Yu 2011), and forecasts for mortgage loan prepayment speeds (Carlin, Longstaff & Matoba 2014). Agents publicly disagree also in matters of politics (Wiegel 2009) and climate change (Marlon, Leiserowitz & Feinberg 2013).

<sup>3</sup>Indeed, this perspective has been taken in work by Al-Najjar (2009), Gilboa, Samuelson & Schmeidler (2013), and Acemoglu, Chernozhukov & Wold (2015), among others, in various non-strategic settings (see Section 9.3 for an extended review). I embed these ideas into an incomplete information game, and study their implications for strategic behavior.

<sup>4</sup>Reflecting, for example, common cultural experiences or industry-specific norms.

<sup>5</sup>The property of uniform consistency is satisfied by many families of inference rules, including any finite inference rule class, as well as certain classes of kernel density estimators with variable bandwidths, and certain classes of Bayes estimators with heterogeneous priors.

(see Proposition 1).<sup>6</sup> In this framework, complete information is interpreted as a reduced form for agents having beliefs coordinated by an infinite quantity of data.<sup>7</sup>

Towards the second question of robustness to the common prior assumption, I propose a new robustness criterion for strategic predictions based in the quantity of data that agents need to see. I define a sequence of incomplete information games, called *inference* games, which are indexed by a quantity of (public) observations  $n < \infty$ . In each of these games, agents observe  $n$  random observations, and form beliefs satisfying common inference. As the quantity of data  $n$  tends to infinity, this sequence of games (almost surely) converges to the game in which agents have common certainty of the true parameter value. But for any  $n < \infty$ , agents have different beliefs.

The main part of the paper (Sections 5 and 6) asks: Which solutions of the limit complete information game persist (with high probability) in these finite-data inference games? The key object of study is  $p_n(a)$ , the probability that an action profile  $a$  is a solution given  $n$  observations. Section 5 characterizes which solutions have the property that  $p_n(a) \rightarrow 1$  as  $n$  tends to infinity; these solutions are said to be *robust to inference*. I find that Nash equilibria are robust to inference if and only if they are strict (Theorem 1), and that the robustness of rationalizable actions can be characterized using procedures of iterated elimination of strategies that are never a strict best reply (Theorem 2).

In practice, agents only observe restricted amounts of data. Thus, strategic behavior in the limit (as the quantity of data grows arbitrarily large) may not be the most interesting criterion for predictions in real economic environments. I suggest next that we can provide a measure for *how* robust a solution is by looking at how much data is required to support the solution. In Section 6, I provide lower bounds on  $p_n(a)$  for Nash equilibria (Proposition 2) and rationalizable actions (Proposition 3). For both solution concepts, the quantity of data required depends on several features:

First, it depends on a cardinal measure of strictness of the solution. Say that an action profile is a  $\delta$ -strict NE if each agent's prescribed action is at least  $\delta$  better than his next best action; and say that an action profile is  $\delta$ -strict rationalizable if it can be rationalized by a chain of best responses, in which each action yields at least

---

<sup>6</sup>This assumes implicitly that the unknown parameter can be identified in the data. In the proposed framework, disagreement may persist even given infinite quantities of data if the parameter is not identified.

<sup>7</sup>Recent papers have argued that agreement need not occur even in an infinite data limit. For example, Acemoglu, Chernozhukov & Werning (2015) show that asymptotic beliefs need not agree when individuals are uncertain about signal distributions. I assume agreement given infinite data in the main part of the paper to emphasize the question of when (sufficient) agreement occurs given finite data. In Section 7.1, I show that this is a stronger assumption than is necessary for the main results.

$\delta$  over the next best alternative. This parameter  $\delta$  turns out to determine how much estimation error the solution can withstand—the higher the degree of strictness (the larger the parameter  $\delta$ ), the less data agents need to see.

Second, the quantity of data required depends on the “diversity” of the inference rules. When agents have common knowledge of a smaller set of inference rules, or when these inference rules diverge less in their predictions given common data, then fewer observations are needed to coordinate beliefs. Conversely, lack of common knowledge over what constitutes a reasonable interpretation of data serves to prolong disagreement. Thus, this criterion provides a formal sense in which the common prior assumption is less appropriate for predicting strategic interactions across cultures, nations, and industries.

Finally, the quantity of data required depends on the “complexity” of the learning problem. I do not provide a universal notion of complexity; instead, the relevant determinants are seen to vary with the choice of inference rules. For many classes of inference rules, an important determinant is dimensionality, and I provide several concrete examples to illustrate this. In these cases, predictions in the limit complete information game are less robust when payoffs are a function of a greater number of covariates.

These comparative statics are, in my view, a key advantage to modeling beliefs using the proposed framework. When agents learn from data, possibly using different inference rules, then channels for disagreement emerge that are complementary to (and distinct from) the traditional channel of differential information. In particular, the amount of common knowledge over how to interpret data, and the “dimensionality” or “complexity” of the unknown parameter, are both crucial to determining dispersion in beliefs. These sources for disagreement have potentially new implications for policy design and informational design: for example, summary statistics may facilitate coordination by reducing the complexity of information (and thus, coordinating beliefs).

The final sections proceed as follows: Section 7 examines several modeling choices made in the main text, and discusses the extent to which the main results rely on these choices. In particular, I look at relaxations of uniform consistency (Section 7.1), the introduction of private data (Section 7.2), and the introduction of limit uncertainty (Section 7.3).

Section 8 surveys the related literature. This paper builds a bridge between the body of work that studies the robustness of equilibrium and equilibrium refinements (Fudenberg, Kreps & Levine 1988, Carlsson & van Damme 1993, Kajii & Morris 1997, Weinstein & Yildiz 2007), and the body of work that studies the asymptotic properties of learning from data (Cripps, Ely, Mailath & Samuelson 2008, Al-Najjar 2009, Acemoglu, Chernozhukov & Yildiz 2015).

Section 9 concludes.

## 2 Example

I begin by illustrating ideas with a simple coordination game, in which two farmers decide simultaneously whether or not to adopt a new agricultural technology—for example, a new pest-resistant grain. Continued production of the existing grain results in a payoff of  $\frac{1}{2}$ . Transitioning alone to the new grain results in a payoff of 0, since neither farmer can individually finance the distribution and transportation costs of this new grain. Finally, coordinated transition results in an unknown payoff of  $\theta$ , which I will assume for simplicity takes the value  $\theta = 1$  if crop yield is high, and  $\theta = -1$  if crop yield is low. The payoffs to this game are summarized below:

	Adopt	Not Adopt
Adopt	$\theta, \theta$	$0, \frac{1}{2}$
Not Adopt	$\frac{1}{2}, 0$	$\frac{1}{2}, \frac{1}{2}$

When should we predict coordination on adoption of the new grain?

In the standard approach, all uncertainty about the new grain is described by a state space  $\Omega$ , and we assume that agents share a common prior over  $\Omega$ . In the absence of any private information about the new grain, this approach implies that the two farmers have an identical belief over its future yield. But predicting yields of a new kind of crop is not easy: crop yield is a function of many environmental conditions—the soil structure, light exposure, quantity of rain, etc. I propose an alternative perspective for modeling their beliefs to capture the role that this complexity may play in determining disagreement between agents.

*Learning from Data.* In the proposed model, farmers predict the future yield of the new crop based on how it previously fared in other environments. There are  $r < \infty$  relevant environmental conditions (soil structure, light, ...). For this example, let us assume that each condition takes a value in the interval  $[-c, c]$ , and that the true relationship between environmental conditions in  $[-c, c]^r$  and crop yields (high or low) is given by the following deterministic function:

$$\pi(\mathbf{x}) = \begin{cases} \text{High} & \text{if } \mathbf{x} \in [-c', c']^r \\ \text{Low} & \text{otherwise} \end{cases} \quad \forall \mathbf{x} \in [-c, c]^r$$

where  $c' \in (0, c)$ . That is, crop yields are high under conditions in  $[-c', c']^r$ , and low otherwise. (See Figure 1 below for an illustration of this relationship with  $r = 2$ .)

It is common knowledge that there is a (hyper-)rectangular region of favorable environmental conditions (high yield), and a remaining region of unfavorable conditions. The farmers do not, however, know the exact regions. Instead, they observe the common data

$$(\mathbf{x}_1, \pi(\mathbf{x}_1)), \dots, (\mathbf{x}_n, \pi(\mathbf{x}_n)),$$

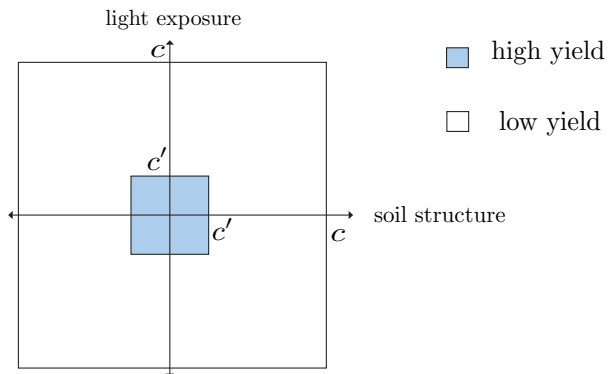


Figure 1: Two relevant attributes ( $r = 2$ ). Yield is high under environmental conditions in  $[-c', c']^2$ , and low otherwise. Farmers do not know the high yield region (shaded).

where  $\mathbf{x}_i$  are identically and independently drawn from a uniform distribution on  $[-c, c]^r$ . That is, farmers observe crop yields in  $n$  different sampled environments. From this data, farmers infer a partitioning  $\hat{\pi}$  that correctly classifies every observation, and use this inferred relationship to predict whether the yield will be low or high in their region. For simplicity, let us take this region to be described by the origin.

The key observation is that many rectangular partitionings will perfectly fit the data; some of which may have different predictions at the origin. This creates room for potential (rational) disagreement. (Figure 2 illustrates two such partitionings based on an example dataset.) Say that a strategic prediction is robust if it holds without further assumption regarding which partitioning either farmer infers, or which partitioning he believes the other farmer to infer. When is coordinated adoption robust?

*Robustness.* Let us first clarify this criterion as follows. Every realization of the data pins down a set of predictions, each of which is consistent with some rectangular partitioning that exactly fits the data. Now, suppose *only* that this set of predictions is common certainty—that is, both farmers put probability 1 on this set of predictions, believe that the other puts probability 1 on this set of predictions, and so forth.<sup>8</sup> This defines a set of possible (hierarchies of) beliefs that either farmer

<sup>8</sup> Formally, define  $\Pi$  to be the family of functions

$$\hat{\pi}(\mathbf{x}) = \hat{\pi}(x^1, \dots, x^r) = \begin{cases} 1 & \text{if } x^k \in [\underline{c}_k, \bar{c}_k] \text{ for every } k = 1, \dots, r \\ 0 & \text{otherwise} \end{cases} \quad \forall \mathbf{x} \in \mathcal{X},$$

parametrized by the tuple  $(\underline{c}_1, \bar{c}_1, \dots, \underline{c}_r, \bar{c}_r) \in \mathbb{R}^{2r}$ . This defines the class of all axis-aligned hyper-




Figure 2: Circles indicate low yields, and squares indicate high yields. The two rectangles identify partitionings (predict high yield if  $\mathbf{x}$  is contained within the rectangle, and low yield otherwise) that exactly fit the data.

could hold.

The key object of interest will be the probability that data is realized such that Adopt is rationalizable given any belief in this set. This probability is a function of the quantity of data  $n$  and of the number of conditions  $r$ ; I will write it as  $p(n, r)$ , and refer to it as the plausibility of coordination on adoption.

**Claim 1.** *For every quantity of data  $n \geq 1$ , number of environmental conditions  $r \geq 1$ , and constants  $c, c' \in \mathbb{R}_+$ ,*

$$p(n, r) = \left( 1 - \left[ 2 \left( \frac{2c - c'}{2c} \right)^n - \left( \frac{c - c'}{c} \right)^n \right] \right)^r.$$

*Proof.* See appendix. □

This claim has several implications.

**Observation 1.** *Coordinated adoption is more plausible when the quantity of data  $n$  is larger.*

From Claim 1, we see that  $p(n, r)$  is increasing in  $n$ . Indeed,  $p(n, r) \rightarrow 1$  as the quantity of data  $n$  tends to infinity (for fixed  $r$ ). Thus, if farmers observe crop yields in sufficiently many different environments, then coordinated adoption is arbitrarily plausible.

---

rectangles. Agents have common certainty in the set


$$\{\hat{\pi}(\mathbf{0}) : \hat{\pi} \in \Pi \text{ and } \hat{\pi}(\mathbf{x}_k) = \pi(\mathbf{x}_k) \quad \forall k = 1, \dots, n\}.$$

**Observation 2.** *Coordinated adoption is less plausible when the number of environmental conditions  $r$  is larger.*

From Claim 1, we see that  $p(n, r)$  is decreasing in  $r$  when  $n$  is sufficiently large.<sup>9</sup> In fact,  $p(n, r) \rightarrow 0$  as the number of environmental conditions  $r$  tends to infinity (for fixed  $n$ ). This suggests that coordinated adoption is less plausible when crop yield depends purely on a single environmental condition, than when it depends on a high-dimensional set of covariates.

**Observation 3.** *Coordinated adoption is less plausible when the set of inference rules is larger.*

The probability  $p(n, r)$  weakly decreases for every  $n$  and  $r$  as we expand the set of possible interpretations of the data. For example, suppose we assume that it is common knowledge that the region of high crop yields is described by a *rotated* rectangle, instead of axis-aligned rectangles as assumed above. This weakly expands the room for possible disagreement, since there are datasets such as the one in the figure below



where every axis-aligned rectangle partitioning predicts high yield at the origin, but some rotated rectangle partitioning predicts low yield.<sup>10</sup> This suggests that coordinated adoption is more plausible when extrapolation from past crop yields is coordinated by external means—for example, a common culture, or a common set of heuristics.

<sup>9</sup>A sufficient condition is

$$2 \left( \frac{2c - c'}{2c} \right)^n - \left( \frac{c - c'}{c} \right)^n < \frac{1}{r}.$$

<sup>10</sup>Since the set of rotated rectangle partitionings includes the set of axis-aligned rectangle partitionings, clearly if every partitioning in the former set predicts that rebellion will be successful, then every partitioning in the latter set will as well.



*Takeaways.* Under the proposed approach, prediction of coordinated adoption of the new grain is more plausible when agents have previously observed few trial instances of the new crop, when the determinants of crop yield are high-dimensional, and when there is not a common approach to extrapolating from past yields. In the main body of the paper, I generalize the ideas in this example, proposing a model in which agents have common certainty in the predictions of an arbitrary class of inference rules, and a robustness criterion for equilibria and rationalizable actions in all finite normal-form games.

### 3 Preliminaries and Notation

#### 3.1 The game

Consider a set  $\mathcal{I}$  of  $I < \infty$  agents and finite action sets  $(A_i)_{i \in \mathcal{I}}$ . As usual, let  $A = \prod_{i \in \mathcal{I}} A_i$ . The set of complete information (normal-form) games defined on these primitives is the set of payoff matrices in  $U := \mathbb{R}^{|\mathcal{I}| \times |A|}$ . Let  $\Theta \subseteq \mathbb{R}^k$  be a compact set of parameters and fix a metric  $d^0$  on  $\Theta$  such that  $(\Theta, d^0)$  is complete and separable. I will identify these parameters with payoff matrices under a map  $g$  satisfying:

**Assumption 1.**  $g : \Theta \rightarrow U$  is a bounded and Lipschitz continuous embedding.<sup>11</sup>

This map  $g$  can be interpreted as capturing the known information about the structure of payoffs. For example, in the game presented in Section 2, players know that payoffs belonged to the parametric family of payoffs

$$\begin{array}{cc} & \begin{array}{cc} a_1 & a_2 \end{array} \\ \begin{array}{c} a_1 \\ a_2 \end{array} & \begin{array}{cc} \theta, \theta & 0, \frac{1}{2} \\ \frac{1}{2}, 0 & \frac{1}{2}, \frac{1}{2} \end{array} \end{array}$$

but do not know the value of  $\theta$ . Notice that identifying payoffs with parameters in this way is without loss of generality, since we can always take  $\Theta := U$  and set  $g$  to be the identity map. For clarity of exposition, I will sometimes write  $u(a, \theta)$  for  $g(\theta)(a)$ , or  $u_i(a, \theta)$  for the payoffs to agent  $i$ . Finally, denote the true value of the parameter by  $\theta^*$ , and suppose that it is unknown.

*Remark 1.* It is also possible to interpret  $\theta$  as indexing a family of distributions over payoffs; for example,  $\theta$  may be the mean of a normal distribution with a fixed variance. In this case,  $g$  maps parameters in  $\Theta$  to *expected* payoffs under the distribution determined by  $\theta$ .

---

<sup>11</sup>A map is an embedding if it is a homeomorphism onto its image.

### 3.2 Beliefs

Now let us define beliefs on  $\Theta$ .

*Type space.* For notational simplicity, consider first  $I = 2$ . Following Brandenburger & Dekel (1993), recursively define

$$\begin{aligned} X_0 &= \Theta \\ X_1 &= X_0 \times (\Delta(X_0)) \\ &\vdots \\ X_n &= X_{n-1} \times (\Delta(X_{n-1})) \end{aligned}$$

and take  $T_0 = \prod_{n=0}^{\infty} \Delta(X_n)$ . An element  $(t^1, t^2, \dots) \in T_0$  is a complete description of beliefs over  $\Theta$  (describing the agent's uncertainty over  $\Theta$ , his uncertainty over his opponents' uncertainty over  $\Theta$ , and so forth), and is referred to as a *type*.

This approach can be generalized for  $I$  agents, taking  $X_0 = \Theta$ ,  $X_1 = X_0 \times (\Delta(X_0))^{I-1}$ , and building up in this way. Mertens & Zamir (1985) have shown that for every agent  $i$ , there is a subset of types  $T_i^*$  (that satisfy the property of *coherency*<sup>12</sup>) and a function  $\kappa_i^* : T_i^* \rightarrow \Delta(\Theta \times T_{-i}^*)$  such that  $\kappa_i(t_i)$  preserves the beliefs in  $t_i$ ; that is,  $\text{marg}_{X_{n-1}} \kappa_i(t_i) = t_i^n$  for every  $n$ . Notice that  $T_{-i}^*$  is used here to denote the set of profiles of opponent types.

The tuple  $(T_i^*, \kappa_i^*)_{i \in \mathcal{I}}$  is known as the *universal type space*. Other tuples  $(T_i, \kappa_i)_{i \in \mathcal{I}}$  with  $T_i \subseteq T_i^*$  for every  $i$ , and  $\kappa_i : T_i \rightarrow \Delta(\Theta \times T_{-i})$  represent alternative (smaller) *type spaces*. Finally, let  $T^* = T_1^* \times \dots \times T_I^*$  denote the set of all type profiles, with typical element  $t = (t_1, \dots, t_I)$ .

*Remark 2.* Types are sometimes inference ruled as encompassing *all* uncertainty in the game. In this paper, I separate strategic uncertainty over opponent actions from structural uncertainty over  $\Theta$ .

*Topology on types.* Let  $T_i^k = \Delta(X_{k-1}) = \Delta(\Theta \times T_{-i}^{k-1})$  denote the set of possible  $k$ -th order beliefs for agent  $i$ .<sup>13</sup> The *uniform-weak topology* on  $T_i^*$ , proposed in Chen, di Tillio, Faingold & Xiong (2010), is the metric topology generated by the distance

$$d_i^{UW}(t_i, t'_i) = \sup_{k \geq 1} d^k(t_i, t'_i) \quad \forall t_i, t'_i \in T_i^*,$$

where  $d^0$  is the metric defined on  $\Theta$  (see Section 2.1)<sup>14</sup> and recursively for  $k \geq 1$ ,  $d^k$

<sup>12</sup> $\text{marg}_{X_{n-2}} t^n = t^{n-1}$ , so that  $(t^1, t^2, \dots)$  is a consistent stochastic process.

<sup>13</sup>Working only with types in the universal type space, it is possible to identify each  $X_k$  with its first and last coordinates, since all intermediate information is redundant.

<sup>14</sup>In Chen et al. (2010),  $\Theta$  is finite and  $d^0$  is the discrete metric, but this construction extends to all complete and separable  $(\Theta, d^0)$ .

is the Prokhorov distance<sup>15</sup> on  $\Delta\left(\Theta \times T_{-i}^{k-1}\right)$  induced by the metric  $\max\{d^0, d^{k-1}\}$  on  $\Theta \times T_i^{k-1}$ .

*Common  $p$ -belief.* Define  $\Omega = \Theta \times T^*$  to be the set of all “states of the world,” such that every element in  $\Omega$  corresponds to a complete description of uncertainty. Following Monderer & Samet (1989), for every  $E \subseteq \Omega$ , let

$$\mathcal{B}^p(E) := \{(\theta, t) : \kappa_i(t_i)(E) \geq p \text{ for every } i\}, \quad (1)$$

describe the event in which every agent believes  $E \subseteq \Omega$  with probability at least  $p$ . Common  $p$ -belief in the set  $E$  is given by

$$\mathcal{C}^p(E) := \bigcap_{k \geq 1} [\mathcal{B}^p]^k(E).$$

The special case of common 1-belief is referred to in this paper as *common certainty*.

I use in particular the concept of common certainty in a set of first-order beliefs. For any  $F \subseteq \Delta(\Theta)$ , define the event

$$E_F := \{(\theta, t) : \text{marg}_{\Theta} t_i \in F \text{ for every } i\}, \quad (2)$$

in which every agent’s first-order belief is in  $F$ . Then,  $\mathcal{C}^1(E_F)$  is the event in which it is common certainty that every agent has a first-order belief in  $F$ . The set of types  $t_i$  given which agent  $i$  believes that  $F$  is common certainty is the projection of  $\mathcal{C}^1(E_F)$  onto  $T_i^*$ .<sup>16</sup> Since this set is identical across agents, I will refer to this simply as the set of types with common certainty in  $F$ .

### 3.3 Solution concepts

Two solution concepts for incomplete information games are used in this paper.

*Interim Correlated Rationalizability* (Dekel, Fudenberg & Morris 2007). For every agent  $i$  and type  $t_i$ , set  $S_i^0[t_i] = A_i$ , and define  $S_i^k[t_i]$  for  $k \geq 1$  such that  $a_i \in S_i^k[t_i]$  if and only if  $a_i \in BR_i\left(\text{marg}_{\Theta \times A_{-i}} \pi\right)$  for some  $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$  satisfying (1)  $\text{marg}_{\Theta \times T_{-i}} \pi = \kappa_i(t_i)$  and (2)  $\pi\left(a_{-i} \in S_{-i}^{k-1}[t_{-i}]\right) = 1$ , where

<sup>15</sup>Recall that the *Levy-Prokhorov* distance  $\rho$  between measures on metric space  $(X, d)$  is defined

$$\rho(\mu, \mu') = \inf \left\{ \delta > 0 : \mu(E) \leq \mu'(E^\delta) + \delta \text{ for each measurable } E \subseteq X \right\}$$

for all  $\mu, \mu' \in \Delta(X)$ , where  $E^\delta = \{x \in X : \inf_{x' \in E} d(x, x') < \delta\}$ .

<sup>16</sup>Notice that when beliefs are allowed to be wrong (as they are in this approach), individual *perception* of common certainty is the relevant object of study. That is, agent  $i$  can believe that a set of first-order beliefs is common certainty, even if no other agent in fact has a first-order belief in this set. Conversely, even if every agent indeed has a first-order belief in  $F$ , agent  $i$  may believe that no other agent has a first-order belief in this set.

$S_{-i}^{k-1}[t_{-i}] = \prod_{j \neq i} S_j^{k-1}[t_{-j}]$ . We can interpret  $\pi$  to be an extension of belief  $\kappa_i(t_i)$  onto the space  $\Delta(\Theta \times T_{-i} \times A_{-i})$ , with support in the set of actions that survive  $k-1$  rounds of iterated elimination of strictly dominated strategies for types in  $T_{-i}$ . For every  $i$ , define

$$S_i^\infty[t_i] = \bigcap_{k=0}^{\infty} S_i^k[t_i]$$

to be the set of actions that are interim correlated rationalizable for agent  $i$  of type  $t_i$ , or (henceforth) simply *rationalizable*.

*Interim Bayesian Nash equilibrium.* Fix any type space  $(T_i, \kappa_i)_{i \in \mathcal{I}}$ . A *strategy* for player  $i$  is a measurable function  $\sigma_i : T_i \rightarrow A_i$ . The strategy profile  $(\sigma_1, \dots, \sigma_I)$  is a Bayesian Nash equilibrium if

$$\sigma_i(t_i) \in \operatorname{argmax}_{a \in A_i} \int_{\Theta \times T_{-i}} u_i(a_i, \sigma_{-i}(t_{-i}), \theta) d\kappa_i(t_i) \quad \text{for every } i \in \mathcal{I} \text{ and } t_i \in T_i.$$

In a slight abuse of terminology, I will say throughout that action profile  $a$  is an (interim) Bayesian Nash equilibrium if the strategy  $\sigma$  with  $\sigma_i(t_i) = a_i$  for every  $t_i \in T_i$  is a Bayesian Nash equilibrium.

## 4 Learning from Data

What are agent beliefs over the unknown parameter (and over opponent beliefs over the unknown parameter), and how are they formed? In this section, I describe a framework in which beliefs are formed by learning from data.

Say that a *dataset* is any sequence of  $n$  observations  $z_1, \dots, z_n$ , sampled i.i.d. from a set  $\mathcal{Z}$  according to an exogenous distribution  $P$ . Throughout, I use  $Z_n$  to denote the random sequence corresponding to  $n$  observations, and  $\mathbf{z}_n$  to denote a typical realization (or simply  $Z$  and  $\mathbf{z}$  when the number of observations is not important).

The key assumption of my approach is a restriction on the possible types that agents can have following rationalization of the data. I begin by introducing a few relevant concepts. Define a *inference rule* to be any map  $\mu : \mathbf{z} \mapsto \mu_{\mathbf{z}}$  from the set of possible datasets<sup>17</sup> to  $\Delta(\Theta)$ , the set of (Borel) probability measures on  $\Theta$ . Fix a family of inference rules  $M$ .

**Definition 1.** For every dataset  $\mathbf{z}$ , say that

$$F_{\mathbf{z}} = \{\mu_{\mathbf{z}} : \mu \in M\} \subseteq \Delta(\Theta)$$

is the set of plausible *first-order beliefs*.

<sup>17</sup> $\cup_{n \geq 1} \mathcal{Z}^n$ , where  $\mathcal{Z}^n$  denotes the  $n$ -fold Cartesian product of the set  $\mathcal{Z}$ .

This is the set of all distributions over  $\Theta$  that emerge from evaluating the dataset  $\mathbf{z}$  using an inference rule in  $M$ . For every dataset  $\mathbf{z}$ , define  $T_{\mathbf{z}}$  to be the set of all (interim) types for whom  $F_{\mathbf{z}}$  is common certainty.<sup>18</sup> That is, every type in  $T_{\mathbf{z}}$  has a plausible first-order belief, puts probability 1 on every other agent having a plausible first-order belief, and so forth. The main restriction below, which I will refer to from now on as *common inference*, assumes that following realization of data  $\mathbf{z}$ , every agent has a type in  $T_{\mathbf{z}}$ .<sup>19</sup>

**Assumption 2** (Common inference). *Given any dataset  $\mathbf{z}$ , every agent  $i$  has an (interim) type  $t_i \in T_{\mathbf{z}}$ .*<sup>20</sup>

Several special examples for the set of inference rules  $M$  are collected below.

*Example 1* (Bayesian updating with a common prior). Define  $\mu$  to be the map that takes any dataset  $\mathbf{z}$  into the Bayesian posterior induced from the common prior and a common likelihood function. Let  $M = \{\mu\}$ . Then, for every dataset  $\mathbf{z}$ , the set  $F_{\mathbf{z}}$  consists of the singleton Bayesian posterior induced from the common prior, and the set  $T_{\mathbf{z}}$  consists of the singleton type with common certainty in this Bayesian posterior.<sup>21</sup>

*Example 2* (Bayesian updating with uncommon priors). Fix a set of prior distributions on  $\Theta$  and a common likelihood function. Every inference rule  $\mu \in M$  is identified with a prior distribution in this set, and maps the observed data to the Bayesian posterior induced from this prior and the common likelihood.

*Example 3* (Kernel regression with different bandwidth sequences). Let  $\mathcal{X} \subseteq \mathbb{R}$  be a set of attributes, which are related to outcomes in  $\Theta$  under the unknown map  $f : \mathcal{X} \rightarrow \Theta$ . Data is of form

$$(x_1, y_1), \dots, (x_n, y_n),$$

where every  $x_k \in \mathcal{X}$  and every  $y_k = f(x_k)$ . Suppose that the unknown parameter  $\theta^*$  is the value of the function  $f$  evaluated at input  $x_0$ .

Inference rules in  $M$  map the data to an estimate for  $\theta^*$  by first producing an estimated function  $\hat{f}$ , and then evaluating this function at  $x_0$ . The approach for estimating  $f$  is as follows: Fix a kernel function<sup>22</sup>  $K : \mathbb{R}^d \rightarrow \mathbb{R}$ , and let  $h_n \rightarrow 0$  be

<sup>18</sup>See the end of Section 3.2 for a more formal exposition.

<sup>19</sup>The results in this paper follow without modification if we relax this assumption to common certainty in the convex hull of distributions in  $F_{\mathbf{z}}$ . See Lemma 5.

<sup>20</sup>Notice that this paper takes an unusual interpretation of the ex-ante/interim distinction, which does not explicitly invoke a Bayesian perspective. In this paper, the role of the prior is replaced by a data-generating process.

<sup>21</sup>That is, his first-order beliefs are given by this posterior, and he believes with probability 1 that his opponents' first-order beliefs are given by this posterior, and so forth.

<sup>22</sup> $K$  is measurable and satisfies the conditions

$$\int_{\mathbb{R}^d} K(x) dx = 1$$

a sequence of constants tending to zero. Define  $\hat{f}_{n,h} : \mathcal{X} \rightarrow \Theta$  to be the Nadaraya-Watson estimator

$$\hat{f}_{n,h}(x) = \frac{(nh_n)^{-1} \sum_{k=1}^n y_k K\left(\frac{x - x_k}{h_n^{1/d}}\right)}{(nh_n)^{-1} \sum_{k=1}^n K\left(\frac{x - x_k}{h_n^{1/d}}\right)},$$

which produces estimates by taking a weighted average of nearby observations. This describes an individual inference rule  $\mu$ .

Now let  $H$  be a set of (bandwidth) sequences, each of which determines a different level of “smoothing” applied to the data. Every inference rule  $\mu \in M$  is identified with a sequence  $h_n \in H$ . Thus,  $M$  is a set of kernel regression estimators with different bandwidth sequences.

*Remark 3.* Common inference does not impose an explicit relationship between agent beliefs and estimators. For example, all of the following are consistent with common inference:

- Every agent  $i$  is identified with an inference rule  $\mu_i \in M$ , and the sequence of inference rules  $(\mu_i)_{i \in \mathcal{I}}$  is common knowledge.
- Every agent  $i$  is identified with an inference rule  $\mu_i \in M$ . Agent  $i$  knows his own inference rule  $\mu_i$ , but has a nondegenerate belief distribution over the inference rules of other agents.
- Every agent  $i$  is identified with a distribution  $P_i$  on  $M$ , and draws an inference rule at random from  $M$  under this distribution.

In the main part of this paper, I assume common inference (Assumption 2), and ask what properties of beliefs and strategic behavior can be deduced from this assumption alone.

#### 4.1 When do agents commonly learn?

Let us begin by considering the property of common learning. Say that agents *commonly learn* the true parameter  $\theta^*$  if, as the quantity of data increases, every agent believes that it is approximate common certainty that the value of the parameter is close to  $\theta^*$ . It will be useful in this section to assign to every agent  $i$  a map  $t^i : \mathbf{z} \mapsto t_{\mathbf{z}}^i$ , which takes the realized data into a type in  $T_{\mathbf{z}}$ .

**Definition 2** (Common Learning). *Agents commonly learn  $\theta^*$  if*

$$\lim_{n \rightarrow \infty} P^n \left( \left\{ \mathbf{z}_n : t_{\mathbf{z}_n}^i \in \mathcal{C}^p(B_\epsilon(\theta^*)) \right\} \right) = 1 \quad \forall i$$

---


$$\sup_{x \in \mathbb{R}^d} \|K(x)\| = \kappa < \infty$$

for every  $p \in [0, 1)$  and  $\epsilon > 0$ .

That is, for every level of confidence  $p$  and precision  $\epsilon$ , every agent eventually believes that the  $\epsilon$ -ball around the true parameter  $\theta^*$  is common  $p$ -belief. This definition is modified from Cripps et al. (2008).<sup>23</sup> When does common inference imply that agents commonly learn the true parameter  $\theta^*$ ?

The following property of families of inference rules  $M$  will be useful:

**Definition 3** (Uniform consistency.). *The family of inference rules  $M$  is  $\theta^*$ -uniformly consistent if*

$$\sup_{\mu \in M} d_P(\mu_{Z_n}, \delta_{\theta^*}) \rightarrow 0 \text{ a.s.}$$

where  $d_P$  is the Prokhorov metric on  $\Delta(\Theta)$ .

*Remark 4.* Say that an individual inference rule  $\mu$  is  $\theta^*$ -consistent if  $d_P(\mu_{Z_n}, \delta_{\theta^*}) \rightarrow 0$  a.s. Uniform consistency is immediately satisfied by any finite family of  $\theta^*$ -consistent inference rules.

Recalling that  $d_P$  metrizes the topology of weak convergence of measures, this property says that for every  $\mu \in M$ , the distribution  $\mu_{Z_n}$  (almost surely) weakly converges to a degenerate distribution on  $\theta^*$ . Moreover, this convergence is uniform in inference rules.

**Proposition 1.** *Every agent commonly learns the true parameter  $\theta^*$  if and only if  $M$  is  $\theta^*$ -uniformly consistent.*

The structure of the argument is as follows. From Chen et al. (2010), we know that convergence in the uniform-weak topology is equivalent to approximate common certainty in the true parameter. Under  $\theta^*$ -uniform consistency, it can be shown that with probability 1, every sequence of types from  $\{T_{Z_n}\}_{n \geq 1}$  converges in the uniform-weak topology to the type with common certainty in  $\theta^*$ . This is, loosely, because possible  $k$ -th order beliefs are restricted to have support in the possible  $(k - 1)$ -th order beliefs, so that in fact the rate of convergence of first-order beliefs is a uniform upper-bound on the rate of convergence of beliefs at every order. The details of this proof can be found in the appendix.

I assume in the remainder of the paper that  $M$  is  $\theta^*$ -uniformly consistent, so that beliefs converge as the quantity of data tends to infinity. The next part of the paper transitions the focus to the stronger property of convergence of solution sets.

---

<sup>23</sup>I take  $\epsilon > 0$ , so that agents believe it is approximate common certainty that the parameter is close to  $\theta^*$ ; in Cripps et al. (2008),  $\Theta$  is finite, so agents believe it is approximate common certainty that the parameter is exactly  $\theta^*$ .

## 5 Robustness to Estimation

Suppose that action  $a_i$  is rationalizable for agent  $i$  (or, action profile  $a$  is an equilibrium) in the complete information game in which the true parameter  $\theta^*$  is common certainty. Can we guarantee that action  $a_i$  (action profile  $a$ ) remains rationalizable (an equilibrium) when payoffs are inferred from data, so long as agents observe a sufficiently large number of observations?

### 5.1 Concepts

I will first introduce the idea of an inference game. For any dataset  $\mathbf{z}$ , define  $G(\mathbf{z})$  to be the incomplete information game with primitives  $\mathcal{I}, (A_i)_{i \in \mathcal{I}}, \Theta, g$ , and type space

$$\mathcal{T}_{\mathbf{z}} = (T_i^{\mathbf{z}}, \kappa_i^{\mathbf{z}})_{i \in \mathcal{I}},$$

where  $T_i^{\mathbf{z}} = T_{\mathbf{z}}$  for every  $i$ , and  $\kappa_i^{\mathbf{z}}$  is the restriction of  $\kappa_i$  (as defined in Section 3.2) to  $T_i^{\mathbf{z}}$ .<sup>24</sup> Notice that if  $T_{\mathbf{z}}$  consists only of the type with common certainty of  $\theta^*$ , then this game reduces to the complete information game with payoffs given by  $g(\theta^*)$ .

We can interpret inference games as follows. Suppose the analyst knows *only* that agents have observed data  $\mathbf{z}$ , and that Assumption 2 (Common Inference) holds. Then, the set of types that any player  $i$  may have is given by  $T_{\mathbf{z}}$ . Recall that as the quantity of data tends to infinity, this set  $T_{\mathbf{z}}$  converges almost surely to the singleton type with common certainty in  $\theta^*$ . So, for large quantities of data, inference games approximate the (true) complete information game. The question of interest is with what probability solutions in this limit game persist in finite-data inference games. This question is made precise for the solution concepts of Nash equilibrium and rationalizability in the following way.

For any Nash equilibrium  $a$  of the limit complete information game, define  $p_n^{NE}(a)$  to be the probability that data  $\mathbf{z}_n$  is realized, such that the strategy profile

$$(\sigma_i)_{i \in \mathcal{I}}, \quad \text{with } \sigma_i(t_i) = a_i \quad \forall i \in \mathcal{I}, t_i \in T_i^{\mathbf{z}_n}$$

is a Bayesian Nash equilibrium. Analogously, define  $p_n^R(i, a_i)$  to be the probability that data  $\mathbf{z}_n$  is realized, such that

$$a_i \in S_i^\infty[t_i] \quad \forall t_i \in T_i^{\mathbf{z}_n};$$

that is,  $a_i$  is rationalizable for agent  $i$  given any type in  $T_i^{\mathbf{z}_n}$ .

**Definition 4.** *The rationalizability of action  $a_i$  for player  $i$  is robust to inference if*

$$p_n^R(i, a_i) \rightarrow 1 \text{ as } n \rightarrow \infty.$$

---

<sup>24</sup>Notice that  $\kappa_i^{\mathbf{z}}(T_i^{\mathbf{z}}) = T_{\mathbf{z}}$  for every agent  $i$ , so this is a belief-closed type space.



The equilibrium property of action profile  $a$  is robust to inference if

$$p_n^{NE}(a) \rightarrow 1 \text{ as } n \rightarrow \infty.$$

What is the significance of robustness to inference? Suppose that action  $a_i$  is rationalizable when the true parameter is common certainty, and suppose moreover that this property of  $a_i$  is robust to inference. Then, the analyst believes with high probability that  $a_i$  is rationalizable for all types in the realized inference game, so long as the quantity of observed data is sufficiently large. Conversely, suppose that  $a_i$  is rationalizable when the true parameter is common certainty, but that this property of  $a_i$  is not robust to inference. Then, there exists a constant  $\delta > 0$  such that for *any* finite quantity of data, the probability that  $a_i$  fails to be rationalizable for some type in the realized inference game is at least  $\delta$ . In this way, robustness to inference is a minimal requirement for the rationalizability of  $a_i$  to persist when agents infer their payoffs from data. Analogous statements apply when we replace rationalizability with equilibrium.

Let us first consider two trivial examples in which robustness to inference imposes no restrictions. Consider the game with payoff matrix

$$\begin{array}{cc} & a_1 & a_2 \\ a_1 & \theta^*, \theta^* & 0, 0 \\ a_2 & 0, 0 & \frac{1}{2}, \frac{1}{2} \end{array}$$

where  $\theta^* > 0$ . Is the equilibrium  $(a_1, a_1)$  robust to inference?

*Example 4* (Trivial inference.). Let  $M$  consist of the singleton inference rule  $\mu$  satisfying

$$\mu_{\mathbf{z}} = \delta_{\theta^*} \quad \forall \mathbf{z},$$

so that  $\mu_{\mathbf{z}}$  is always degenerate on the true value  $\theta^*$ . Then, the set of plausible first-order beliefs is  $F_{\mathbf{z}} = \{\delta_{\theta^*}\}$  for every  $\mathbf{z}$ , so that the true parameter  $\theta^*$  is common certainty with probability 1. Thus, the inference game  $G(\mathbf{z})$  reduces to a complete information game, and the equilibrium property of  $(a_1, a_1)$  is trivially robust to inference.

*Example 5* (Unnecessary inference.). Let  $\Theta := [0, \infty)$ . Then, action profile  $(a_1, a_1)$  is a Nash equilibrium for every possible value of  $\theta \in \Theta$ . Thus, the strategy profile that maps any type of either player to the action  $a_1$  is a Bayesian Nash equilibrium for any beliefs that players might hold over  $\Theta$ . In this way, the family of inference rules  $M$  is irrelevant, and  $(a_1, a_1)$  is again trivially robust to inference.

The two following conditions rule out these cases in which inference is either trivial or unnecessary.

**Assumption 3** (Nontrivial Inference.). *There exists a constant  $\gamma > 0$  such that*

$$P^n(\{\mathbf{z}_n : \delta_{\theta^*} \in \text{Int}(F_{\mathbf{z}_n})\}) > \gamma.$$

for every  $n$  sufficiently large.

This property says that for sufficient quantities of data, the probability that  $\delta_{\theta^*}$  is contained in the interior of the set of plausible first-order beliefs  $F_{\mathbf{z}_n}$  is bounded away from 0. Assumption 3 rules out the example of trivial inference, as well as related examples in which every inference rule in  $M$  overestimates, or every inference rule underestimates, the unknown parameter.<sup>25</sup>

To rule out the second example, I impose a richness condition on the image of  $g$ . For every agent  $i$  and action  $a_i \in A_i$ , define  $S(i, a_i)$  to be the set of complete information games in which  $a_i$  is a strictly dominant strategy for agent  $i$ ; that is,

$$S(i, a_i) := \{u' \in U : u'_i(a_i, a_{-i}) > u'_i(a'_i, a_{-i}) \quad \forall a'_i \neq a_i \text{ and } \forall a_{-i}\}.$$

**Assumption 4** (Richness.). *For every  $i \in \mathcal{I}$  and  $a_i \in A_i$ ,  $g(\Theta) \cap S(i, a_i) \neq \emptyset$ .*

Under this restriction, which is also assumed in Carlsson & van Damme (1993) and Weinstein & Yildiz (2007), every action is strictly dominant at some parameter value. This condition is trivially satisfied if  $\Theta = U$ .

In the subsequent analysis, I assume that the family of inference rules  $M$  satisfies nontrivial inference, and the map  $g$  satisfies richness. These conditions are abbreviated to NI and R, respectively.

## 5.2 Bayesian Nash Equilibrium

When is the equilibrium property of an action profile robust to inference? (From now on, I will abbreviate this to saying that the action profile is itself robust to inference.)

**Theorem 1.** *Assume NI and R. Then, the equilibrium property of action profile  $a^*$  is robust to inference if and only if it is a strict Nash equilibrium.*

The intuition for the proof is as follows. Define  $U_{a^*}^{NE}$  to be the set of all payoffs  $u$  such that  $a^*$  is a Nash equilibrium in the complete information game with payoffs  $u$ . The interior of  $U_{a^*}^{NE}$  is exactly the set of payoffs  $u$  with the property that  $a^*$  is a *strict* Nash equilibrium given these payoffs. I show that as the quantity of data tends to infinity, agents (almost surely) have common certainty in a shrinking neighborhood of the true payoffs, so it follows that  $a^*$  is robust to inference if and only if the true payoff function  $u^* = g(\theta^*)$  lies in the interior of  $U_{a^*}^{NE}$ .

<sup>25</sup>This does not rule out sets of *biased* estimators. It may be that in expectation, every inference rule in  $M$  overestimates the true parameter; Assumption 3 requires that underestimation occurs with probability bounded away from 0.

*Proof.* First, I show that the interior of the set  $U_{a^*}^{NE}$  is characterized by the set of complete information games in which  $a^*$  is a strict Nash equilibrium.

**Lemma 1.**  $u \in \text{Int}(U_{a^*}^{NE})$  if and only if action profile  $a^*$  is a strict Nash equilibrium in the complete information game with payoffs  $u$ .

*Proof.* Suppose  $a^*$  is not a strict Nash equilibrium in the complete information game with payoffs  $u$ . Then, there is some agent  $i$  and action  $a_i \neq a_i^*$  such that

$$u_i(a_i, a_{-i}^*) \geq u_i(a_i^*, a_{-i}^*).$$

Define  $u^\epsilon$  such that  $u_i^\epsilon(a_i, a_{-i}^*) = u_i(a_i, a_{-i}^*) + \epsilon$ , and otherwise  $u_i^\epsilon$  agrees with  $u_i$ . Then,  $u^\epsilon \in B_\epsilon(u)$  for every  $\epsilon > 0$ , but  $a_i$  is a strictly profitable deviation for agent  $i$  in response to  $a_{-i}^*$  in the game with payoffs  $u_i^\epsilon$ . So  $a^*$  is not an equilibrium in this game. Fix any sequence of positive constants  $\epsilon_n \rightarrow 0$ . Then,  $u^{\epsilon_n} \rightarrow u$  as  $n \rightarrow \infty$ , but  $u^{\epsilon_n} \notin U_{a^*}^{NE}$  for every  $n$ , so it follows that  $u \notin \text{Int}(U_{a^*}^{NE})$  as desired.

Now suppose that  $a^*$  is a strict Nash equilibrium in the complete information game with payoffs  $u$ . Then,

$$\epsilon^* := \inf_{i \in \mathcal{I}} \left( u_i(a_i^*, a_{-i}^*) - \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*) \right) > 0,$$

so  $u \in B_{\epsilon^*}(u) \subseteq U_{a^*}^{NE}$ , with  $B_{\epsilon^*}$  nonempty and open. It follows that  $u \in \text{Int}(U_{a^*}^{NE})$ , as desired.  $\square$

Next, I show that  $a^*$  is robust to inference if and only if the true payoff function is in the interior of the set  $U_{a^*}^{NE}$ .

**Lemma 2.** Let  $u^* = g(\theta^*)$ . The equilibrium property of action profile  $a^*$  is robust to inference if and only if  $u^* \in \text{Int}(U_{a^*}^{NE})$ .

*Proof.* Define  $h(\mu) = \int_{\Theta} g(\theta) d\mu$  to be the map from (first-order) beliefs  $\mu \in \Delta(\Theta)$  into the expected payoff function under  $\mu$ .




Figure 3: The map  $h$  takes first-order beliefs  $\mu$  into expected payoff functions.

Recall that every dataset  $\mathbf{z}$  induces a set of plausible first-order beliefs  $F_{\mathbf{z}}$ . The following claim says that the equilibrium property of  $a^*$  is robust to inference if and only if with high probability the set of expected payoffs  $h(F_{\mathbf{z}_n})$  is contained within  $U_{a^*}^{NE}$  as  $n \rightarrow \infty$ .

**Claim 2.** *The equilibrium property of  $a^*$  is robust to inference if and only if*

$$P^n(\{\mathbf{z}_n : h(F_{\mathbf{z}_n}) \subseteq U_{a^*}^{NE}\}) \rightarrow 1 \quad \text{as } n \rightarrow \infty.$$

*Proof.* I will show that the strategy profile  $(\sigma_i)_{i \in \mathcal{I}}$  with

$$\sigma_i(t_i) = a_i^* \quad \forall i \in \mathcal{I}, \forall t_i \in T_{\mathbf{z}}$$

is a Bayesian Nash equilibrium if and only if  $h(F_{\mathbf{z}}) \subseteq U_{a^*}^{NE}$ . From this, the above claim follows immediately.

Suppose  $h(F_{\mathbf{z}}) \subseteq U_{a^*}^{NE}$ . Then, for any payoff function  $u \in h(F_{\mathbf{z}})$ ,

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \forall i \in \mathcal{I} \text{ and } a_i \neq a_i^*. \quad (3)$$

Fix an arbitrary agent  $i$  and type  $t_i$  with common certainty in  $h(F_{\mathbf{z}})$ , and define  $\mu_i := \text{marg}_{\Theta} t_i$  to be his first-order belief. By construction,  $\mu_i$  assigns probability 1 to  $h(F_{\mathbf{z}})$ , so it follows from (3) that

$$\int_U u_i(a_i^*, a_{-i}^*) dg_*(\mu_i) \geq \int_U u_i(a_i, a_{-i}^*) dg_*(\mu_i) \quad \forall a_i \neq a_i^*,$$

where  $g_*(\mu)$  denotes the pushforward measure of  $\mu$  under mapping  $g$ . Repeating this argument for all agents and all types with common certainty in  $h(F_{\mathbf{z}})$ , it follows that  $(\sigma_i)_{i \in \mathcal{I}}$  is indeed a Bayesian Nash equilibrium.

Now suppose to the contrary that  $h(F_{\mathbf{z}}) \not\subseteq U_{a^*}^{NE}$  and consider any payoff function  $u$  that is in  $h(F_{\mathbf{z}})$  but not in  $U_{a^*}^{NE}$ . Then, there exists some agent  $i$  for whom

$$u_i(a_i^*, a_{-i}^*) - \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*) < 0.$$

Let  $t_i$  be the type with common certainty in  $g^{-1}(u)$ . Then, agent  $i$  of type  $t_i$  has a profitable deviation to some  $a_i \neq a_i^*$ , so  $(\sigma_i)_{i \in \mathcal{I}}$  is not a Bayesian Nash equilibrium.  $\square$

The final claim says that

$$P^n(\{\mathbf{z}_n : h(F_{\mathbf{z}_n}) \subseteq U_{a^*}^{NE}\}) \rightarrow 1 \quad \text{as } n \rightarrow \infty$$

if and only if  $u^*$  is in the interior of the set  $U_{a^*}^{NE}$ . This is, loosely, because  $h(F_{\mathbf{z}_n})$  converges to the singleton set  $\{u^*\}$ ; its proof is deferred to the appendix.

**Claim 3.**  $\lim_{n \rightarrow \infty} P^n(\{\mathbf{z}_n : h(F_{\mathbf{z}_n}) \subseteq U_{a^*}^{NE}\}) = 1$  if and only if  $u^* \in \text{Int}(U_{a^*}^{NE})$ .  $\square$

The theorem directly follows from Lemmas 1 and 2.  $\square$

### 5.3 Rationalizable Actions

When is the property of rationalizability of an action robust to inference? Theorem 1 suggests that the corresponding condition is strict rationalizability in the limit complete information game. This intuition is roughly correct, but subtleties in the procedure of elimination are relevant, and the theorem below will rely on two different such procedures.

First, recall the usual definition for strict rationalizability, introduced in Dekel, Fudenberg & Morris (2006). For every agent  $i$  and type  $t_i$ , set  $R_i^1(t_i) = A_i$ . Then, recursively define  $R_i^k(t_i)$ , for every  $k \geq 2$ , such that  $a_i \in R_i^k[t_i]$  if and only if

$$\int_{\Theta \times T_{-i} \times A_{-i}} (u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)) \, d\pi > 0 \quad \forall a'_i \neq a_i \quad (4)$$

for some distribution  $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$  satisfying (1)  $\text{marg}_{\Theta \times T_{-i}} \pi = \kappa_{t_i}$ , and (2)  $\pi(a_{-i} \in R_{-i}^{k-1}[t_{-i}]) = 1$ . That is, an action survives the  $k$ -th round of elimination only if it is a strict best response to some distribution over opponent strategies surviving the  $(k-1)$ -th round of elimination. Let

$$R_i^\infty[t_i] = \bigcap_{k=0}^{\infty} R_i^k[t_i]$$

be the set of player  $i$  actions that survive every round of elimination. Define  $t_{\theta^*}$  to be the type with common certainty in the true parameter  $\theta^*$ . I will say that action  $a_i$  is *strongly strict-rationalizable* if  $a_i \in R_i^\infty[t_{\theta^*}]$ , where *strongly* is used to contrast with the definition below.

Notice that in this definition, every action that is never a strict best response (to surviving opponent strategies) is eliminated at once. This choice has consequences for the surviving set, since elimination of strategies that are never a strict best response is an order-dependent process. Following, I introduce a new procedure, in which actions are eliminated (at most) one at a time.

For every agent  $i$ , let  $W_i^1 := A_i$ . Then, for  $k \geq 2$ , recursively remove (at most) one action in  $W_i^k$  that is not a strict best reply to any opponent strategy  $\alpha_{-i}$  with support in  $W_{-i}^{k-1}$ . That is, either the set difference  $W_i^k - W_i^{k+1}$  is empty, or it consists of a singleton action  $a_i$  where there does not exist any  $\alpha_{-i} \in \Delta(W_{-i}^{k-1})$  such that

$$u_i(a_i, \alpha_{-i}) > u_i(a'_i, \alpha_{-i}) \quad \forall a'_i \neq a_i.$$

That is,  $a_i$  is not a strict best reply to any distribution over surviving opponent actions. Let

$$W_i^\infty = \bigcap_{k \geq 1} W_i^k$$

be the set of player  $i$  actions that survive every round of elimination, and say that any set  $W_i^\infty$  constructed in this way survives an order of weak strict-rationalizability. Define  $\mathcal{W}_i^\infty$  to be the intersection of all sets  $W_i^\infty$  surviving an order of weak strict-rationalizability. I will say that an action  $a_i$  is *weakly strict-rationalizable* if  $a_i \in \mathcal{W}_i^\infty$ .<sup>26</sup>

**Theorem 2.** *Assume NI and R. Then, the rationalizability of action  $a_i^*$  for agent  $i$  is robust to inference if  $a_i^*$  is strongly strict-rationalizable, and only if  $a_i^*$  is weakly strict-rationalizable.*

*Remark 5.* If there are two players, then the theorem above can be strengthened as follows: Assume NI and R. Then, the rationalizability of action  $a_i^*$  for agent  $i$  is robust to inference if and only if  $a_i^*$  is weakly strict-rationalizable.

*Remark 6.* The existence of actions that are strongly strict-rationalizable, but not weakly strict-rationalizable, occurs only for a non-generic set of payoffs.<sup>27</sup> See the discussion preceding Figure 5.3 for a characterization of these intermediate cases.

*Remark 7.* Rationalizable actions that are robust to inference need not exist. For example, in the degenerate game

	$a_3$	$a_4$
$a_1$	0, 0	0, 0
$a_2$	0, 0	0, 0

all actions are rationalizable, but none are robust to inference.

*Remark 8.* Why is refinement obtained, in light of the results of Weinstein & Yildiz (2007)? The key intuition is that the negative result in Weinstein & Yildiz (2007) relies on construction of tail beliefs that put sufficient probability on payoff functions with dominant actions. But under common inference, it is common certainty that every player puts low probability on “most” payoff functions. So, with high probability, contagion from “far-off” payoff functions with a dominant action cannot begin.

A second explanation for why refinement is obtained is the following. One can show that the perturbations considered in this paper are a subset of perturbations in the uniform-weak topology, which is finer than the topology used in Weinstein & Yildiz (2007). In particular, the sequences of types used to show failure of robustness in Weinstein & Yildiz (2007) do not converge in the uniform-weak topology.

---

<sup>26</sup>The choice of *weak* to describe the latter procedure, and *strong* to describe the former, is explained by Claim 4 (see Appendix B), which says that an action is strongly strict-rationalizable only if it is weakly strict-rationalizable.

<sup>27</sup>The set of such payoffs is nowhere dense in the Euclidean topology on  $U$ .

The broad structure of the proof follows that of Theorem 1, with several new complications that I discuss below. Recall that as the quantity of data tends to infinity, agents have common certainty in a (shrinking) neighborhood of the true payoffs. Thus,  $a_i^*$  is robust to inference if and only if common certainty in a sufficiently small neighborhood of the true payoffs  $u^*$  implies that the action  $a_i^*$  is rationalizable for player  $i$ .

*A necessary condition for robustness to inference.* In analogy with the set  $U_{a_i^*}^{NE}$ , define  $U_{a_i^*}^R$  to be the set of all complete information games in which  $a_i^*$  is rationalizable.<sup>28</sup> Clearly, if  $u^*$  is on the boundary of this set, then common certainty in a neighborhood of  $u^*$  (no matter how small) cannot guarantee rationalizability of  $a_i^*$ . Therefore, a necessary condition for robustness to inference is that  $u^*$  must lie in the interior of  $U_{a_i^*}^R$ . The first lemma says that the interior of  $U_{a_i^*}^R$  is characterized by the set of actions that survive every process of weak strict-rationalizability.

**Lemma 3.**  $u \in \text{Int} \left( U_{a_i^*}^R \right)$  if and only if  $a_i^* \in \mathcal{W}_i^\infty$  in the complete information game with payoffs  $u$ .

Why is  $\text{Int} \left( U_{a_i^*}^R \right)$  characterized by this particular notion of strict rationalizability, and not by others? I provide an example that illustrates why various other natural candidates are not the right notion, and follow this with a brief intuition for the proof of Lemma 3.

Consider the payoff matrices below:

$$\begin{array}{cc}
 & \begin{array}{cc} a_3 & a_4 \end{array} \\
 \begin{array}{c} (u_1) \\ \\ \\ \end{array} & \begin{array}{cc} a_1 & 1, 0 \\ a_2 & 0, 0 \end{array}
 \end{array}
 \qquad
 \begin{array}{cc}
 & \begin{array}{cc} a_3 & a_4 \end{array} \\
 \begin{array}{c} (u_2) \\ \\ \\ \end{array} & \begin{array}{cc} a_1 & 1, 0 \\ a_2 & 0, 0 \end{array}
 \end{array}$$

If *all* strategies that are never a strict best reply are eliminated simultaneously (corresponding to strong strict-rationalizability), then  $a_1$  does not survive in either game.<sup>29</sup> If the criterion is survival of *any* process of iterated elimination of strategies that are never a strict best reply, then  $a_1$  survives in both games.<sup>30</sup> But  $u_1$  is in

<sup>28</sup>Here I abuse notation and write  $U_{a_i^*}^R$  instead of  $U_{i,a_i^*}^R$ .

<sup>29</sup>In the first round, both actions are eliminated for player 2, so  $a_1$  trivially cannot be a best reply for player 1 to any surviving player 2 action.

<sup>30</sup>For example, the order of elimination

$$\begin{array}{cc}
 & \begin{array}{cc} a_3 & a_4 \end{array} \\
 \begin{array}{c} a_1 \\ a_2 \end{array} & \begin{array}{cc} 1, 0 & 1, 0 \\ 0, 0 & 0, 0 \end{array} \longrightarrow \begin{array}{cc} a_3 & a_4 \\ a_1 & 1, 0 \\ a_2 & 0, 0 \end{array} \longrightarrow \begin{array}{cc} a_3 & a_4 \\ a_1 & 1, 0 \\ a_2 & \end{array}
 \end{array}$$

in the first game, and

$$\begin{array}{cc}
 & \begin{array}{cc} a_3 & a_4 \end{array} \\
 \begin{array}{c} a_1 \\ a_2 \end{array} & \begin{array}{cc} 1, 0 & 1, 0 \\ 0, 0 & 1, 0 \end{array} \longrightarrow \begin{array}{cc} a_3 & a_4 \\ a_1 & 1, 0 \\ a_2 & 0, 0 \end{array} \longrightarrow \begin{array}{cc} a_3 & a_4 \\ a_1 & 1, 0 \\ a_2 & \end{array}
 \end{array}$$

the interior of  $U_{a_1}^R$ , while  $u_2$  is not,<sup>31</sup> so neither of these notions provides the desired differentiation.

Now, I provide a brief intuition for the “only-if” direction of Lemma 3. Suppose that action  $a_i^*$  fails to survive some iteration of weak strict-rationalizability. Then, there is some sequence of sets  $(W_i^k)_{k \geq 1}$  satisfying the recursive description in the definition of weak strict-rationalizability, such that  $a_i^* \notin W_i^K$  for  $K < \infty$ . To show that  $a_i^*$  is not robust to inference, I construct a sequence of payoffs  $u^n \rightarrow u$  with the property that  $a_i^*$  fails to be rationalizable in every complete information game  $u^n$ , for  $n$  sufficiently large. The key feature of this construction is translation of weak dominance under the payoffs  $u$  to strict dominance under the payoffs  $u^n$ . This is achieved by iteratively increasing the payoffs to every action that survives to  $W_i^{k+1}$  by  $\epsilon$ , thus breaking ties in accordance with the selection in  $(W_i^k)_{k \geq 1}$ .

So, a necessary condition for robustness to inference is weak strict-rationalizability. Next, I show that a sufficient condition is strong strict-rationalizability, and explain the gap between these two conditions.

*A sufficient condition for robustness to inference.* The reason why weak strict-rationalizability is not sufficient is because, unlike the analogous case for equilibrium, common certainty in the set  $U_{a_i^*}^R$  does not imply rationalizability of  $a_i^*$ .<sup>32</sup> In fact, even if beliefs are concentrated on a (vanishingly) small neighborhood of a payoff function in  $\text{Int}(U_{a_i^*}^R)$ , it may be that  $a_i^*$  fails to be rationalizable. See Appendix D for such an example.

---

in the second.

<sup>31</sup>Action  $a_1$  remains rationalizable in every complete information game with payoffs close to  $u_1$ , so  $u_1 \in \text{Int}(U_{a_1}^R)$ . In contrast, for arbitrary  $\epsilon \geq 0$ , the payoff matrix

$$(u'_2) \quad \begin{array}{cc} & \begin{array}{cc} a_3 & a_4 \end{array} \\ \begin{array}{c} a_1 \\ a_2 \end{array} & \begin{array}{cc} 1, 0 & 1, \epsilon \\ 0, 0 & 1 + \epsilon, \epsilon \end{array} \end{array}$$

is within  $\epsilon$  of  $u_2$  (in the sup-norm), but  $a_1$  is not rationalizable in the complete information game with payoffs  $u'_2$ . So, the payoff  $u_2$  lies on the boundary of  $U_{a_1}^R$ .

<sup>32</sup>A simple example is the following. Consider the following two payoffs:

$$(u_1) \quad \begin{array}{cc} & \begin{array}{cc} a_3 & a_4 \end{array} \\ \begin{array}{c} a_1 \\ a_2 \end{array} & \begin{array}{cc} 1 & 0 \\ \frac{3}{4} & \frac{3}{4} \end{array} \end{array} \quad (u_2) \quad \begin{array}{cc} & \begin{array}{cc} a_3 & a_4 \end{array} \\ \begin{array}{c} a_1 \\ a_2 \end{array} & \begin{array}{cc} 0 & 1 \\ \frac{3}{4} & \frac{3}{4} \end{array} \end{array}$$

Action  $a_1$  is rationalizable for agent 1 in both complete information games, so  $u_1, u_2 \in U_{a_1}^R$ . But action  $a_1$  is strictly dominated by action  $a_2$  if each game is equally likely, since in expectation payoffs are

$$\begin{array}{cc} & \begin{array}{cc} a_3 & a_4 \end{array} \\ \begin{array}{c} a_1 \\ a_2 \end{array} & \begin{array}{cc} \frac{1}{2} & \frac{1}{2} \\ \frac{3}{4} & \frac{3}{4} \end{array} \end{array}$$



*Remark 9.* This example shows moreover that weak strict-rationalizability is not lower hemi-continuous in the uniform-weak topology. Since strong strict-rationalizability is lower-hemicontinuous in the uniform-weak topology (Dekel, Fudenberg & Morris 2006, Chen et al. 2010), this example suggests that subtleties in the definition of strict rationalizability have potentially large implications for robustness.

The reason why common certainty of a shrinking set in  $U_{a_i^*}^R$  need not imply rationalizability of  $a_i^*$  is because the chain of best responses rationalizing action  $a_i^*$  can vary across  $U_{a_i^*}^R$ . In particular, it may be that the true payoffs  $u^*$  lie on the boundary between two open sets of payoff functions, each with different families of rationalizable actions. See Figure 5.3 below for an illustration. These cases are problematic because even though  $a_i^*$  is rationalizable when agents (truly) have common certainty in any payoff functions close to  $u^*$ , it may fail to be rationalizable if agents (mistakenly) believe that payoff functions on different sides of the boundary are common certainty.

On the other hand, if  $a_i^*$  is strongly strict-rationalizable, then it can be justified by a chain of strict best responses that remain constant on some neighborhood of  $u^*$ . It can be shown in this case that common certainty in a vanishing neighborhood of  $u^*$  indeed implies rationalizability of  $a_i^*$ . This provides the sufficient direction.




Figure 4: The set  $U_{a_i^*}^R$  is partitioned such that every agent's set of rationalizable actions is constant across each element of the partition. There are three cases: (1) if  $u^*$  is on the boundary of  $U_{a_i^*}^R$  (e.g.  $u_1$ ), then  $a_i^*$  is not robust to inference; (2) if  $u^*$  is in the interior of  $U_{a_i^*}^R$ , and moreover in the interior of a partition element ( $u_2$ ), then  $a_i^*$  is certainly robust to inference; (3) if  $u^*$  in the interior of  $U_{a_i^*}^R$ , but not in the interior of any partition element ( $u_3$ ), then  $a_i^*$  may not be robust to inference. See Appendix D for an example of the last case.

## 6 How Much Data do Agents Need?

Theorems 1 and 2 characterize the persistence of equilibria and rationalizable actions given sufficiently large quantities of data. But in practice, the quantity of data that agents observe about payoff-relevant parameters is limited. Robustness to

inference is meaningful only if convergence obtains in the ranges of data that we can reasonably expect agents to observe. Therefore, I ask next, how *much* data is needed for reasonable guarantees on persistence?

This section addresses this question by providing lower bounds for  $p_n^R(i, a_i)$  and  $p_n^{NE}(a)$ . These bounds suggest a second, stronger criterion for equilibrium selection, based in the quantity of data needed to reach a desired threshold probability. These bounds also highlight the importance of various features of the solution and the game, including the degree of strictness of the solution, and the complexity of the inference problem.

## 6.1 Bayesian Nash Equilibrium

The following is a measure for the “degree of strictness” of a Nash equilibrium in the complete information game with payoffs  $u^* = g(\theta^*)$ . For any  $\delta \geq 0$ , say that action profile  $a$  is a  $\delta$ -strict Nash equilibrium<sup>33</sup> if

$$u_i^*(a_i, a_{-i}) - \max_{a'_i \neq a_i} u_i^*(a'_i, a_{-i}) > \delta \quad \forall i \in \mathcal{I}.$$

Every strict Nash equilibrium  $a^*$  admits the following cardinal measure of strictness:

$$\delta_a^{NE} = \sup \{ \delta : a \text{ is a } \delta\text{-strict NE} \},$$

which represents the largest  $\delta$  for which  $a$  is a  $\delta$ -strict NE. This parameter describes the amount of slack in the equilibrium  $a$ —action profile  $a$  remains an equilibrium on at least a  $\delta_a^{NE}$ -neighborhood of the payoff function  $u^*$ .

**Proposition 2.** *Suppose  $a^*$  is a  $\delta$ -strict Nash equilibrium for some  $\delta \geq 0$ . Then, for every  $n \geq 1$ ,*

$$p_n^{NE}(a^*) \geq 1 - \frac{2}{\delta_a^{NE}} \mathbb{E}_{P^n} \left( \sup_{\mu \in M} \|h(\mu_{Z_n}) - u^*\|_\infty \right) \quad (5)$$

where  $h(\nu) = \int_{\Theta} g(\theta) d\nu$  for every  $\nu \in \Delta(\Theta)$ .

*Remark 10.* Uniform consistency of  $M$  implies that  $\sup_{\mu \in M} \|h(\mu_{Z_n}) - u^*\|_\infty \rightarrow 0$  a.s., so for any strict Nash equilibrium  $a^*$ , the bound in Proposition 2 converges

<sup>33</sup>Replacing the strict inequality  $>$  with a weak inequality  $\geq$ , this definition reverses the more familiar concept of  $\epsilon$ -equilibrium, which requires that

$$u_i^*(a_i, a_{-i}) - \max_{a'_i \neq a_i} u_i^*(a'_i, a_{-i}) \geq -\epsilon \quad \forall i, \text{ where } \epsilon \geq 0.$$

The concept of  $\epsilon$ -equilibrium was introduced to formalize a notion of approximate Nash equilibria (violating the equilibrium conditions by no more than  $\epsilon$ ). I use  $\delta$ -strict equilibrium to provide a cardinal measure for the *strictness* of a Nash equilibrium (satisfying the conditions with  $\delta$  to spare).

to 1.<sup>34</sup> This implies also that the gap between  $p_n^{NE}(a^*)$  and its lower bound in (5) converges to 0 as the quantity of data  $n$  tends to infinity.

How can we interpret this bound? By assumption, action profile  $a^*$  is an equilibrium in the complete information game with payoffs  $u^*$ . But when  $n < \infty$ , agents may have heterogenous and incorrect beliefs. The probability with which  $a^*$  persists as an equilibrium under these modified beliefs is determined by two components:

$$1 - \underbrace{\frac{2}{\delta_{a^*}^{NE}}}_{(1)} \underbrace{\mathbb{E}_{P^n} \left( \sup_{\mu \in M} \|h(\mu_{Z_n}) - u^*\|_\infty \right)}_{(2)}.$$

First, it depends on the fragility of the solution  $a^*$  to introduction of heterogeneity and error in beliefs. This is reflected in component (1): the bound is increasing in the parameter  $\delta_{a^*}^{NE}$ . Intuitively, equilibria that are “stricter” persist on a larger neighborhood of the true payoffs  $u^*$ . It turns out that common certainty in the  $\delta_{a^*}^{NE}/2$ -neighborhood of  $u^*$  is sufficient to imply that  $a^*$  is an equilibrium (see Lemma 11).

Second, the probability  $p_n^{NE}(a^*)$  depends on the expected error in beliefs. This is reflected in the second component:  $\|\mu_{Z_n} - u^*\|_\infty$  is the (random) error in estimated payoffs using a fixed inference rule  $\mu \in M$ ; so  $\sup_{\mu \in M} \|\mu_{Z_n} - u^*\|_\infty$  is the (random) supremum error across inference rules in  $M$ ; and (2) gives the expected supremum error across inference rules in  $M$ . As  $n$  tends to infinity, this quantity tends to zero,<sup>35</sup> but the speed at which inference rules in  $M$  uniform converge to the truth is determined by the “diversity” of inference rules in  $M$ , and by the statistical complexity of the learning problem.

This first feature, diversity, can be thought of as a property of the relationship between inference rules in  $M$  to each other. Holding fixed the rate at which individual inference rules learn, the lower bound is lower when inference rules in  $M$  *jointly* learn slower. How this occurs, and how much effect this can have on the analyst’s confidence  $p_n^{NE}(a^*)$ , is discussed in detail in Section 6.3.

This second feature, complexity, can be thought of as a property of the relationship between inference rules in  $M$  and the data. For example, in Section 2, the probability  $p(n, r)$  decreases in the dimensionality of the data. More generally, when finite-sample bounds for the uniform rate of convergence of inference rules in  $M$  are available, they can be plugged into the lower bound in Proposition 2. This technique is illustrated below for a new set of inference rules  $M$ .

*Example 6.* Let us consider agents who use ordinary least-squares regression to estimate a relationship between  $p$  covariates and a real-valued outcome variable.

<sup>34</sup>This follows from continuity of the map  $h$  (see Lemma 4).

<sup>35</sup>Since  $M$  is  $\theta^*$ -uniform consistency and  $h$  is continuous.

An observation is a tuple  $(x_i, y_i) \in \mathcal{Z} := \mathbb{R}^p \times \mathbb{R}$ , where

$$y_i = x_i^T \beta + \epsilon_i$$

with  $x_i \sim_{\text{i.i.d.}} \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ ,  $\epsilon_i \sim_{\text{i.i.d.}} \mathcal{N}(0, 1)$ , and  $x_i$  and  $\epsilon_i$  independent. Suppose that the first coordinate of the coefficient vector  $\beta$ , denoted  $\beta_1$ , is payoff-relevant. That is,  $\Theta = \mathbb{R}$ , and the true parameter is  $\theta^* = \beta_1$ .

Recall that the least-squares estimate for the coefficient vector  $\beta$  is

$$\hat{\beta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$$

where  $\mathbf{X}$  is the matrix whose  $i$ -th row is given by  $x_i$ , and  $\mathbf{Y}$  is the matrix whose  $i$ -th row is given by  $y_i$ . Fix a sequence of constants  $\phi_n$  that tends to 0. Let  $M$  consist of the set of inference rules that map the data into a distribution with support in  $B_{\phi_n}(\hat{\beta}_1)$ . That is, every inference rule maps the data into distribution with support in the  $\phi_n$ -neighborhood of the least-squares estimate for  $\beta_1$ .

*Corollary 1.* *Suppose the data-generating process and family of inference rules is as described in the above example. Then, for every complete information game and  $\delta$ -strict Nash equilibrium  $a^*$  (with  $\delta \geq 0$ ),*

$$p_n^{NE}(a^*) \geq 1 - \frac{2K}{\delta_{a^*}^{NE}} (\sigma^2 p (\sqrt{n} + \sqrt{p}) + \phi_n^2)$$

and  $K$  is the Lipschitz constant<sup>36</sup> of the map  $g : \Theta \rightarrow U$ .

So we see that the lower bound is decreasing in the number of covariates  $p$  (i.e. the analyst is less confident in predicting  $a^*$  when the number of covariates is larger). The proof can be found in the appendix.

## 6.2 Rationalizable Actions

We can now repeat the previous exercise for the solution concept of rationalizability. The following is a measure for the “degree” of rationalizability of an action  $a_i$  in the complete information game with payoffs  $u^*$ . For any  $\delta \geq 0$ , say that the family of sets  $(R_j)_{j \in \mathcal{I}}$  is *closed under  $\delta$ -best reply* if for every agent  $j$  and action  $a_j \in R_j$ , there is some distribution  $\alpha_{-j} \in \Delta(R_{-j})$  such that

$$u_j^*(a_j, \alpha_{-j}) > u_j^*(a'_j, \alpha_{-j}) + \delta \quad \forall a'_j \neq a_j. \quad (6)$$

Say that action  $a_i$  is  *$\delta$ -strict rationalizable* for agent  $i$  if there exists some family  $(R_j)_{j \in \mathcal{I}}$ , with  $a_i \in R_i$ , that is closed under  $\delta$ -best reply. Every strictly rationalizable

---

<sup>36</sup>Assuming the sup-norm on  $U$  and the Euclidean norm on  $\Theta$ .

action  $a_i$  admits the following cardinal measure of the degree of strictness:

$$\delta_{a_i}^R = \sup\{\delta : a_i \text{ is } \delta\text{-strict rationalizable}\}.$$
<sup>37</sup>

This parameter describes the amount of slack in the rationalizability of action  $a_i$ —that is, action  $a_i$  remains rationalizable for agent  $i$  on at least a  $\delta_{a_i}^R$ -neighborhood of the payoff function  $u^*$ .

*Remark 11.* This definition is equivalent to requiring that  $a_i$  survive a more general version of strong strict-rationalizability, where the inequality in (4) is replaced by

$$\int_{\Theta \times T_{-i} \times A_{-i}} (u_i(a_i, a_{-i}, \theta) - u_i(a'_i, a_{-i}, \theta)) d\pi > \delta \quad \forall a'_i \neq a_i,$$

so that  $a_i$  yields at least  $\delta$  more than the next best action given the distribution  $\pi$ .<sup>38</sup>

**Proposition 3.** *Suppose action  $a_i^*$  is  $\delta$ -strict rationalizable for some  $\delta > 0$ . Then, for every  $n \geq 1$ ,*

$$p_n^R(i, a_i^*) \geq 1 - \frac{2}{\delta_{a_i^*}^R} \mathbb{E}_{P^n} \left( \sup_{\mu \in M} \|h(\mu_{Z_n}) - u^*\|_\infty \right),$$

where  $h(\nu) = \int g(\theta) d\nu$  for any  $\nu \in \Delta(\Theta)$ .

*Proof.* See appendix. □

Again, we see that the lower bound is increasing in the “strictness” of the solution, as measured through the parameter  $\delta_{a_i^*}^R$ , and in the speed at which expected payoffs using inference rules in  $M$  uniformly converge to the true payoffs  $u^*$ . As before, when finite-sample bounds are available, they can be used to derive closed-form expressions for this bound.

**Corollary 2.** *Suppose the data-generating process and family of inference rules are as described in Example 6. Then, for every complete information game, agent  $i$ , and  $\delta$ -strict rationalizable action  $a_i^*$  (with  $\delta \geq 0$ ),*

$$p_n^R(i, a_i^*) \geq 1 - \frac{2K}{\delta_{a_i^*}^R} (\sigma^2 p (\sqrt{n} + \sqrt{p}) + \phi_n^2)$$

where  $K$  is the Lipschitz constant<sup>39</sup> of the map  $g : \Theta \rightarrow U$ .

<sup>37</sup>I abuse notation here and write  $\delta_{a_i}^R$  instead of  $\delta_{i,a_i}^R$ . Again, this parameter is defined only if  $a_i$  is  $\delta$ -strict rationalizable for some  $\delta \geq 0$ .

<sup>38</sup>A similar procedure is introduced in Dekel, Fudenberg & Morris (2006). The above definition makes the following modifications: first, the inequality to be strict; second,  $\delta$  appears on the right-hand side of the inequality, instead of  $-\delta$ .

<sup>39</sup>Assuming the sup-norm on  $U$  and the Euclidean norm on  $\Theta$ .

### 6.3 Diversity across Inference Rules in $M$

I conclude this section with a brief discussion regarding the dependence of  $p_n^{NE}(a)$  and  $p_n^R(i, a_i)$  on the diversity across inference rules in  $M$ . To isolate this effect from properties of individual inference rules, let us fix the marginal distributions of  $\mu(Z_n)$  for every  $\mu \in M$ , and vary the *joint* distribution of the random variables  $(\mu(Z_n))_{\mu \in M}$ . Proposition 4 below provides upper and lower bounds for  $p_n^{NE}(a)$  and  $p_n^R(i, a_i)$ . These bounds can be understood from the following simple example.

*Example 7.* Recall the game from Section 2 with payoffs

	A	NA
A	$\theta, \theta$	$0, \frac{1}{2}$
NA	$\frac{1}{2}, 0$	$\frac{1}{2}, \frac{1}{2}$

where  $\Theta = \{-1, 1\}$ . Fix a quantity of data  $n < \infty$ , and suppose that  $M$  consists of two inference rules  $\mu_1, \mu_2$  with marginal distributions

$$\begin{aligned}\mu_1(Z_n) &\sim \frac{1}{4}\delta_{-1} + \frac{3}{4}\delta_1 \\ \mu_2(Z_n) &\sim \frac{3}{4}\delta_{-1} + \frac{1}{4}\delta_1\end{aligned}$$

That is, with probability  $\frac{1}{4}$ , data  $\mathbf{z}_n$  is generated such that  $\mu_1(\mathbf{z}_n)$  is degenerate on  $-1$ , and with probability  $\frac{3}{4}$ , data is generated such that  $\mu_1(\mathbf{z}_n)$  is degenerate on 1. (The distribution of  $\mu_2(Z_n)$  is interpreted similarly.) Given these distributions, what are the largest and smallest possible values of  $p_n^{NE}((A, A))$ ?

First observe that action profile (A, A) is an equilibrium if and only if data  $\mathbf{z}_n$  is realized such that  $\mu_1(\mathbf{z}_n) = \mu_2(\mathbf{z}_n) = \delta_1$ . Otherwise, A is strictly dominated for the agent with first-order belief  $\delta_{-1}$ . At one extreme,  $\mu_1(Z_n)$  and  $\mu_2(Z_n)$  may be correlated such that  $\mu_1(\mathbf{z}_n) = \delta_1$  for every dataset  $\mathbf{z}_n$  where  $\mu_2(\mathbf{z}_n) = \delta_1$ . Then,

$$p_n^{NE}(a) = \Pr(\{\mathbf{z}_n : \mu_2(\mathbf{z}_n) = \delta_1\}) = \frac{1}{4}.$$

If instead,  $\mu_1$  and  $\mu_2$  are independent, then

$$p_n^{NE}(a) = \Pr(\{\mathbf{z}_n : \mu_1(\mathbf{z}_n) = \delta_1\}) \Pr(\{\mathbf{z}_n : \mu_2(\mathbf{z}_n) = \delta_1\}) = \left(\frac{1}{4}\right) \left(\frac{3}{4}\right) < \frac{1}{4}.$$

This quantity is further reduced if  $\mu_2(\mathbf{z}_n) = \delta_1$  implies that  $\mu_1(\mathbf{z}_n) = \delta_{-1}$ , in which case  $p_n^{NE}(a) = 0$ .

These observations can be generalized as follows for arbitrary finite  $M$ . For every inference rule  $\mu$  and quantity of data  $n \geq 1$ , define

$$p_{\mu,n}^{NE}(a) := \Pr(h(\mu_{Z_n}) \in U_a^{NE}) \tag{7}$$

This is the probability that action profile  $a$  is a Nash equilibrium if every agent has beliefs degenerate on the prediction of inference rule  $\mu$ . Define  $p_{\mu,n}^R(i, a_i)$  analogously, replacing  $U_a^{NE}$  with  $U_{a_i}^R$  in (7).

**Proposition 4.** *Suppose  $M$  is finite, and the marginal distributions  $(\mu(Z_n))_{\mu \in M}$  are fixed. Then,*

$$1 - \sum_{\mu \in M} p_{\mu,n}^{NE}(a) \leq p_n^{NE}(a) \leq 1 - \min_{\mu \in M} p_{\mu,n}^{NE}(a)$$

and

$$1 - \sum_{\mu \in M} p_{\mu,n}^R(i, a_i) \leq p_n^R(i, a_i) \leq 1 - \min_{\mu \in M} p_{\mu,n}^R(i, a_i).$$

The upper bound corresponds to co-monotonic random variables, and the lower bound, when attainable, corresponds to counter-monotonic random variables. In the co-monotonic case, different inference rules err in inference of payoffs on the same sets of data, whereas in the counter-monotonic case they err on datasets that are as non-overlapping as possible.

## 7 Extensions

The following section provides brief comment on and extension to various inference ruling choices made in the main framework.

### 7.1 Misspecification

Proposition 1 shows that  $\theta^*$ -uniform consistency is both necessary and sufficient for common learning, and I assume in the remainder of the paper that the family of inference rules  $M$  is  $\theta^*$ -uniformly consistent. But continuity in equilibrium sets (and rationalizable sets) does not require common learning. Can we obtain Theorems 1 and 2 under a weakening of this property?

In fact, it is neither necessary that individual inference rules are consistent, nor necessary that inference rules uniformly converge. I introduce a relaxation of uniform consistency below.

**Definition 5** (Almost  $\theta^*$ -uniform consistency). *For any  $\epsilon \geq 0$ , say that the class of inference rules  $M$  is  $(\epsilon, \theta^*)$ -uniformly consistent if*

$$\lim_{n \rightarrow \infty} \sup_{\mu \in M} d_P(\mu(Z_n), \delta_{\theta^*}) \leq \epsilon \text{ a.s.}$$

where  $d_P$  is the Prokhorov metric on  $\Delta(\Theta)$ .

This says that a class of inference rules is almost  $\theta^*$ -uniformly consistent if the set of plausible first order beliefs converges<sup>40</sup> almost surely to a neighborhood of the true parameter. Notice that uniform consistency is nested as the  $\epsilon = 0$  case. The proofs of Theorems 1 and 2 are easily adapted to show the following result. (In reading this, recall that if  $M$  is  $(\epsilon, \theta^*)$ -uniformly consistent, then it is also  $(\epsilon', \theta^*)$ -uniformly consistent for every  $\epsilon' > \epsilon$ .)

**Proposition 5.** *Assume NI and R.*

1. *The rationalizability of action  $a_i^*$  is robust to inference if  $\delta_{a_i^*}^R > 0$  and  $M$  is  $(\delta_{a_i^*}^R, \theta^*)$ -uniformly consistent.*
2. *The equilibrium property of  $a^*$  is robust to inference if  $\delta_{a^*}^{NE} > 0$  and  $M$  is  $(\delta_{a^*}^{NE}, \theta^*)$ -uniformly consistent.*

## 7.2 Private Data

In the main text, I assume that agents observe a common dataset. How do the main results change if agents observe private data? Cripps et al. (2008) have shown that if  $\mathcal{Z}$  is unrestricted, then common learning may not occur even if  $|M| = 1$  (so that  $M$  contains a single inference rule). It is also known that strict Nash equilibria need not be robust to higher-order uncertainty about opponent data (see e.g. Carlsson & van Damme (1993), Kajii & Morris (1997)). Thus, extension to private data requires restrictions on beliefs over opponent data that are beyond the scope of this paper.

In the simplest extension, however, we may suppose that players observe different datasets  $(\mathbf{z}^i)_{i \in \mathcal{I}}$ , independently drawn from the same distribution, but each has an (incorrect) degenerate belief that all opponents have seen the same data that he has. Then, Theorems 1 and 2 hold as stated, and the bounds in Propositions 2 and 3 are revised as follows.

**Proposition 6.** *Suppose  $a^*$  is a  $\delta$ -strict Nash equilibrium for some  $\delta > 0$ . Then, for every  $n \geq 1$ ,*

$$p_n^{NE}(a^*) \geq \left( 1 - \frac{2}{\delta_{a^*}^{NE}} \mathbb{E}_{P^n} \left( \sup_{\mu \in M} \|h(\mu_{Z_n}) - u^*\|_\infty \right) \right)^I$$

where  $I$  is the number of players. Suppose  $a_i^*$  is  $\delta$ -strict rationalizable for some  $\delta$ . Then, for every  $n \geq 1$ ,

$$p_n^R(i, a_i^*) \geq \left( 1 - \frac{2}{\delta_{a_i^*}^R} \mathbb{E}_{P^n} \left( \sup_{\mu \in M} \|h(\mu_{Z_n}) - u^*\|_\infty \right) \right)^I.$$

---

<sup>40</sup>In the Hausdorff distance induced by  $d_P$ .



### 7.3 Limit Uncertainty

In the main text, I assume that agents learn the true parameter as the quantity of data  $n$  tends to infinity, so that the limit game is a complete information game. This approach can be extended such that the limit game has incomplete information. Fix a distribution  $\nu \in \Delta(\Theta)$ —a limit common prior—and rewrite uniform consistency as follows:

**Definition 6** (Limit Common Prior.). *The set of inference rules  $M$  has a limit common prior  $\nu$  if*

$$\sup_{\mu \in M} d_P(\mu(Z_n), \nu) \rightarrow 0 \text{ a.s.}$$

where  $d_P$  is the Prokhorov metric on  $\Delta(\Theta)$ .

Then, taking  $u^* := h(\nu)$  to be the expected payoff under  $\nu$ , all the results in Section 5 follow without modification.

## 8 Related Literature

This paper makes a connection between the literature regarding robustness of equilibrium to specification of agent beliefs, and the literature that studies agents who learn from data. I discuss each of these literatures in turn.

### 8.1 Robustness of equilibrium and equilibrium refinements

The following question has been the focus of an extensive literature: Suppose an analyst does not know the exact game that is being played. Which solutions in his inference rule of the game can be guaranteed to be close to some solution in all nearby games?

Early work on this question considered “nearby” to mean complete information games with close payoffs (Selten 1975, Myerson 1978, Kohlberg & Mertens 1986). Fudenberg, Kreps & Levine (1988) proposed consideration of nearby games in which players themselves have uncertainty about the true game. This approach of embedding a complete information game into games with incomplete information has since been taken in several papers under different assumptions on beliefs. For example: Carlsson & van Damme (1993) consider a class of incomplete information games in which beliefs are generated by (correlated) observations of a noisy signal of payoffs of the game. Kajii & Morris (1997) study incomplete information games in which beliefs are induced by general information structures that place sufficiently high ex-ante probability on the true payoffs.

I ask which solutions of a complete information game persist in nearby incomplete information games, where the definition of nearby that I use differs from the

existing literature in the following ways: First, I place a strong restriction on (interim) higher-order beliefs, which has the consequence that agents commonly learn the true parameter. This contrasts with Carlsson & van Damme (1993) and Kajii & Morris (1997), in which—even as perturbations become vanishingly small—agents consider it possible that other agents have beliefs about the unknown parameter that are very different from their own. In particular, failures of robustness due to standard contagion arguments do not apply in my setting; thus, I obtain rather different robustness results.<sup>41</sup>

Second, while the restriction I place on interim beliefs is stronger in the sense described above, I do not require that these beliefs are consistent with a common prior. This allows for common knowledge disagreement, which is not permitted in either Carlsson & van Damme (1993) or Kajii & Morris (1997).

Finally, the class of perturbations that I consider are motivated by a learning foundation (this aspect shares features with Dekel, Fudenberg & Levine (2004) and Esponda (2013), but agents in this paper learn about payoffs only, and not actions). I interpret the sequence of interim types as corresponding to learning from a fixed number of observations. This motivates a departure from the literature in studying solution sets not just in nearby games (large  $n$ ), but also in far games (small  $n$ ). In particular, I suggest that we can characterize the degree of robustness by looking at the persistence of solutions in small- $n$  games.

## 8.2 Role of higher-order beliefs

A related literature studies the sensitivity of solutions to specification of higher-order beliefs. Early papers in this literature (Mertens & Zamir 1985, Brandenburger & Dekel 1993) considered types to be nearby if their beliefs were close up to order  $k$  for large  $k$  (corresponding to the product topology on types). Several authors have shown that this notion of close leads to surprising and counterintuitive conclusions, in particular that strict equilibria and strictly rationalizable actions are fragile to perturbations in beliefs (Rubinstein 1989, Weinstein & Yildiz 2007).

These findings have motivated new definitions of “nearby” types. Dekel, Fudenberg & Morris (2006) characterize the coarsest metric topology on types under which the desired continuity properties hold. This topology is defined via strategic properties of types, instead of directly on beliefs. Chen et al. (2010) subsequently

---

<sup>41</sup>For example, the construction of beliefs used in Weinstein & Yildiz (2007) to show failure of robustness (Proposition 2) relies on construction of tail beliefs that place positive probability on an opponent having a first-order belief that implies a dominant action. A similar device is employed in Kajii & Morris (1997) to show that robust equilibria need not exist (see the negative example in Section 3.1). These tail beliefs are not permitted under my approach. When the quantity of data is taken to be sufficiently large, it is common certainty (with high probability) that all players have first-order beliefs close to the true distribution, so the process of contagion cannot begin.

developed a (finer) metric topology on types—the uniform-weak topology—which is defined explicitly using properties of beliefs. In this topology, two types are considered close if they have similar first-order beliefs, attach similar probabilities to other players having similar first-order beliefs, and so forth.

The perturbations in beliefs that I allow for are perturbations in the uniform-weak topology. Specifically, the type spaces that I look at—that is, all type profiles with common certainty in the predictions of a set of inference rules  $M$ —converge in this topology to the singleton type space containing the type with common certainty in the true parameter. Thus, robustness to inference can be interpreted as requiring persistence across a subset of perturbations in the uniform-weak topology.<sup>42</sup> A related study is taken in Morris, Takahashi & Tercieux (2012) and Morris & Takahashi (N.d.), where approximate common certainty in the true parameter is considered, instead of common certainty in a neighborhood of the true parameter.

### 8.3 Agents who learn from data

The set of papers including Gilboa & Schmeidler (2003), Billot, Gilboa, Samet & Schmeidler (2005), Gilboa, Lieberman & Schmeidler (2006), Gayer, Gilboa & Lieberman (2007), and Gilboa, Samuelson & Schmeidler (2013) propose an inductive or case-based approach to inference ruleing economic decision-making. The present paper can be interpreted as studying the strategic behaviors of case-based learners when there is uncertainty over the inductive inference rules used by other agents.

There is also a body of work that studies asymptotic disagreement between agents who learn from data. Cripps et al. (2008) study agents who use the same Bayesian inference rule but observe different (private) sequences of data; Al-Najjar (2009) study agents who use different frequentist inference rules to learn from data; and Acemoglu, Chernozhukov & Yildiz (2015) study Bayesian agents who have different priors over the signal-generating distribution. My inference rule of belief formation shares many features with these inference rules, but the main object of study is the convergence of equilibrium sets, instead of the convergence of beliefs.

Finally, Steiner & Stewart (2008) study the limiting equilibria of a sequence of games in which agents use a kernel density estimator to infer payoffs from related games. This paper is conceptually very close, but there are several important differences in the approach. For example, Steiner & Stewart (2008) suppose that agents share a common inference rule and observe endogenous data (generated by past, strategic actors), while I suppose that agents have different inference rules and observe exogenous data. Additionally, the (common) inference rule in Steiner &

---

<sup>42</sup>The characterizations of robustness in this paper are possibly unchanged if agents have common  $p$ -belief in the predictions of inference rules in  $M$ , where  $p \rightarrow 1$  as the quantity of data tends to infinity. I leave verification of this for future work.

Stewart (2008) is not indexed by the quantity of data, so the limit of their learning process is a game with heterogeneous beliefs, whereas the limit of my process is a game with common certainty of the true distribution.

#### 8.4 Model uncertainty

Consideration of inference rule uncertainty in game theory is largely new, but similar ideas have been advanced in several neighboring areas of economics. Eyster & Piccione (2013) study an asset-pricing inference rule in which agents have different incomplete theories of price formation. The set of papers including Hansen & Sargent (2007), Hansen & Sargent (2010), Hansen & Sargent (2012), and Hansen (2014), among others, consider the implications of inference rule uncertainty for various questions in macroeconomics. In their framework, a decision-maker considers a set of inference rules (prior distributions) plausible, and uses a max-min criterion for decision-making.

#### 8.5 Epistemic game theory

I extensively use tools, results, and concepts from various papers in epistemic game theory, including Monderer & Samet (1989), Brandenburger & Dekel (1993), Morris, Rob & Shin (1995), Dekel, Fudenberg & Morris (2007), Chen et al. (2010). The notion of common certainty in a set of first-order beliefs was studied earlier in Battigalli & Siniscalchi (2003).

### 9 Discussion

Directions for future work include the following:

*Endogenous data.* In this paper, data is generated according to an exogenous distribution. An important next step is to consider data generated by actions played by past strategic actors. In this dynamic setting, past actions play a role in coordinating future beliefs via the kind and quantity of data generated.

*Optimal informational complexity design.* Suppose a designer has control over the complexity of information disclosed to agents in a strategic setting. Using the approach developed in this paper, the designer's choice of complexity influences the commonality in beliefs across agents. When will he choose to disclose simpler information, and when will he disclose information that is more complex? If the designer's interests are opposed to those of the agents, should a social planner regulate the kind of information he can provide?

*Confidence in predictions.* An action profile is usually thought of as having the binary quality of either being, or not being, a solution. The approach in this paper may provide a way to qualify such statements with a level of confidence. In this

paper,  $p_n(a)$  describes the analyst's confidence in predicting  $a$  given  $n$  observations. I hope to extend these ideas towards construction of a cardinal measure for the strength of equilibrium predictions across different games.

## Appendix A: Notation and Preliminaries

- If  $(X, d)$  is a metric space with  $A \subseteq X$  and  $x \in X$ , I write

$$d(A, x) = \sup_{x' \in A} d(x', x).$$

- $\text{Int}(A)$  is used for the interior of the set  $A$ .
- Recall that  $u \in U$  is a payoff matrix. For clarity, I will sometimes write  $u_i$  to denote the the payoffs in  $u$  corresponding to agent  $i$ , and  $u(a, \theta)$  to denote  $g(\theta)(a)$ .
- For any  $\mu, \nu \in \Delta(\Theta)$ , the Wasserstein distance is given by

$$W_1(\mu, \nu) = \inf \mathbb{E}(X, Y),$$

where the expectation is taken with respect to a  $\Theta \times \Theta$ -valued random variable and the infimum is taken over all joint distributions of  $X \times Y$  with marginals  $\mu$  and  $\nu$  respectively.

## Appendix B: Preliminary Results

**Lemma 4.** *The function*

$$h(\mu) = \int_{\Theta} g(\theta) d\mu \quad \forall \mu \in \Delta(\Theta)$$

*is continuous.*

*Proof.* By assumption,  $g$  is Lipschitz continuous; let  $K < \infty$  be its Lipschitz constant (assuming the sup-metric on  $U$ ). Suppose  $d_P(\mu, \mu') \leq \epsilon$ ; then,

$$\begin{aligned} \|h(\mu) - h(\mu')\|_{\infty} &= \left\| \int_{\Theta} g(\theta) d(\mu - \mu') \right\|_{\infty} \leq K \sup_{f \in BL_1(\Theta)} \left\| \int_{\Theta} f(\theta) d(\mu - \mu') \right\|_{\infty} \\ &= KW_1(\mu, \mu') \\ &\leq K(\text{diam}(\Theta) + 1)d_P(\mu, \mu') \\ &\leq K(\text{diam}(\Theta) + 1)\epsilon \end{aligned}$$

using the assumption of Lipschitz continuity in the first inequality, and compactness of  $\Theta$  and the Kantorovich-Rubinstein dual representation of  $W_1$  in the following equality. The second inequality follows from Theorem 2 in Gibbs & Su (2002). So  $h$  is continuous.  $\square$

**Lemma 5.** *If  $d_P(F_{Z_n}, \delta_{\theta^*}) \rightarrow 0$  a.s. , then also*

$$d_P(\text{Conv}(F_{Z_n}), \delta_{\theta^*}) \rightarrow 0 \text{ a.s.}$$

*where  $\text{Conv}(F_{Z_n})$  denotes the convex hull of  $F_{Z_n}$ .*

*Proof.* Fix any dataset  $\mathbf{z}_n$ , constant  $\alpha \in [0, 1]$ , and measures  $\mu, \mu' \in F_{\mathbf{z}_n}$ . Again using the dual representation,

$$\begin{aligned} W_1(\alpha\mu + (1 - \alpha)\mu', \delta_{\theta^*}) &= \sup_{f \in BL_1(\Theta)} \left( \int f(\theta) d((\alpha\mu + (1 - \alpha)\mu') - \delta_{\theta^*}) \right) \\ &= \sup_{f \in BL_1(\Theta)} \alpha \left( \int f(\theta) d(\mu - \delta_{\theta^*}) \right) + (1 - \alpha) \left( \int f(\theta) d(\mu' - \delta_{\theta^*}) \right) \\ &\leq \alpha \sup_{f \in BL_1(\Theta)} \left( \int f(\theta) d(\mu - \delta_{\theta^*}) \right) \\ &\quad + (1 - \alpha) \sup_{f \in BL_1(\Theta)} \left( \int f(\theta) d(\mu' - \delta_{\theta^*}) \right) \\ &= \alpha W_1(\mu, \delta_{\theta^*}) + (1 - \alpha) W_1(\mu', \delta_{\theta^*}) \leq \sup_{\mu \in F_{\mathbf{z}_n}} W_1(\mu, \delta_{\theta^*}) \end{aligned}$$

Moreover, using Theorem 2 in Gibbs & Su (2002),

$$d_P(\alpha\mu + (1 - \alpha)\mu', \delta_{\theta^*})^2 \leq W_1(\alpha\mu + (1 - \alpha)\mu', \delta_{\theta^*}),$$

and also

$$\sup_{\mu \in F_{\mathbf{z}_n}} W_1(\mu, \delta_{\theta^*}) \leq (1 + \text{diam}(\Theta)) \sup_{\mu \in F_{\mathbf{z}_n}} d_P(\mu, \delta_{\theta^*}).$$

Thus, for every dataset  $\mathbf{z}_n$ ,

$$d_P(\text{Conv}(F_{\mathbf{z}_n}), \delta_{\theta^*})^2 \leq (1 + \text{diam}(\Theta)) \sup_{\mu \in F_{\mathbf{z}_n}} d_P(\mu, \delta_{\theta^*}),$$

where  $\text{diam}(\Theta)$  is finite by compactness of  $\Theta$ . So  $d_P(F_{\mathbf{z}_n}, \delta_{\theta^*}) \rightarrow 0$  a.s. implies  $d_P(\text{Conv}(F_{\mathbf{z}_n}), \delta_{\theta^*}) \rightarrow 0$  a.s., as desired.  $\square$

**Claim 4.** Fix any agent  $i$ , and let  $t_{\theta^*}$  be the type with common certainty in  $\theta^*$ . If action  $a_i$  is strongly strict-rationalizable for agent  $i$  with type  $t_{\theta^*}$ , then it is also weakly strict-rationalizable for agent  $i$  in the complete information game with payoffs  $u^* = g(\theta^*)$ .

*Proof.* By induction. Trivially  $R_j^1[t_{\theta^*}] = W_j^1 = A_j$  for every agent  $j$ . If  $a_j \notin W_j^2$ , then it is not a strict best response to any distribution over opponent actions, so also  $a_j \notin R_j^2[t_{\theta^*}]$ . Thus,

$$R_j^2[t_{\theta^*}] \subseteq W_j^2 \quad \forall j.$$

Now, suppose  $R_j^k[t_{\theta^*}] \subseteq W_j^k$  for every agent  $j$ , and consider any agent  $i$  and action  $a_i \in R_i^{k+1}[t_{\theta^*}]$ . By construction of the set  $R_i^{k+1}[t_{\theta^*}]$ , there exists some distribution  $\pi$  with  $\text{marg}_{\Theta \times T_{-i}} \pi = \kappa_{t_{\theta^*}}$  and  $\pi(a_{-i} \in R_{-i}^k[t_{-i}]) = 1$  such that

$$\int_{\Theta \times T_{-i} \times A_{-i}} u_i(a_i, a_{-i}, \theta) d\pi > \int_{\Theta \times T_{-i} \times A_{-i}} u_i(a'_i, a_{-i}, \theta) d\pi + \delta \quad \forall a'_i \neq a_i.$$

But since  $R_i^k[t_i] \subseteq W_i^k$ , the distribution  $\pi$  also satisfies  $\pi(a_{-i} \in W_{-i}^k) = 1$ . So  $a_i$  is a  $\delta$ -best response to some distribution  $\pi$  with support in the surviving set of weakly strict-rationalizable actions, implying that  $a_i \in W_i^{k+1}$ , as desired.  $\square$



## Appendix C: Main Results

### 9.1 Proof of Claim 1

I use the following notation. For every dataset  $\mathbf{z}_n = \{(\mathbf{x}_k, \pi(\mathbf{x}_k))\}_{k=1}^n$ , define

$$F(\mathbf{z}_n) = \{\hat{\pi}(\mathbf{0}) : \hat{\pi} \in \Pi \text{ and } \hat{\pi}(\mathbf{x}_k) = \pi(\mathbf{x}_k) \quad \forall k = 1, \dots, n\}$$

and let  $T_{\mathbf{z}_n}$  be the set of hierarchies of belief with common certainty in  $F(\mathbf{z}_n)$ . (See footnote 8 for the definition of  $\Pi$ .) Also, let  $t_{-1}$  be the type with common certainty in  $-1$ , and let  $t_1$  be the type with common certainty in  $1$ . Observe that  $R$  is rationalizable for type  $t_1$  and not for type  $t_{-1}$ .

Suppose  $F(\mathbf{z}_n) = \{-1, 1\}$ . Then  $t_{-1} \in T_{\mathbf{z}_n}$ , so there is a type in  $T_{\mathbf{z}_n}$  for whom  $R$  is not rationalizable. Now suppose  $F(\mathbf{z}_n) = \{1\}$ . Then, the only permitted type is  $t_1$ , so  $R$  is trivially rationalizable for every type in  $T_{\mathbf{z}_n}$ . It follows that  $R$  is rationalizable for every type in  $T_{\mathbf{z}_n}$  if and only if  $F(\mathbf{z}_n) = \{1\}$ ; that is, if and only if every inference rule  $\hat{\pi} \in \Pi$  that exactly fits  $\mathbf{z}_n$  predicts  $\hat{\pi}(\mathbf{0}) = 1$ . For what datasets  $\mathbf{z}_n$  does this hold?

We can reduce this problem by looking at whether the smallest hyper-rectangle that contains every successful observation also contains the origin. This will be the case if and only if for every dimension  $k$ , there exist observations  $(\mathbf{x}_i, 1)$  and  $(\mathbf{x}_j, 1)$  such that  $x_i^k < 0$  and  $x_j^k > 0$  (that is, the  $k$ -th attribute is negative in some observed high yield region, and positive in some observed high yield region). For every  $k$ , this probability is

$$1 - \left[ 2 \left( \frac{2c - c'}{2c} \right)^n - \left( \frac{c - c'}{c} \right)^n \right].$$

Realization of  $k$ -th attributes are independent across dimensions. Thus, the probability that this holds for every dimension is

$$\left( 1 - \left[ 2 \left( \frac{2c - c'}{2c} \right)^n - \left( \frac{c - c'}{c} \right)^n \right] \right)^r$$

as desired.

### 9.2 Proof of Proposition 1

The proof of this proposition follows from two lemmas. The first is a straightforward generalization of Proposition 6 in Chen et al. (2010)<sup>43</sup>, and relates common learning to convergence of types in the uniform-weak topology. The second lemma says that for every dataset  $\mathbf{z}$ , the distance between  $t_{\mathbf{z}}^i$  and  $t_{\theta^*}$  is upper bounded by  $d_P(F_{\mathbf{z}_n}, \delta_{\theta^*})$ .

Throughout, I use  $t_{\theta^*}$  to denote the type with common certainty in  $\theta^*$ .

---

<sup>43</sup>This lemma appears in Chen et al. (2010) for the case in which  $\Theta$  is a finite set and  $d^0$  is the discrete metric, but generalizes to any complete and separable metric space  $(\Theta, d^0)$  when the definition of common learning is replaced by Definition 2.

**Lemma 6.** *Agent  $i$  commonly learns  $\theta^*$  if and only if*

$$d_i^{UW}(t_{Z_n}^i, t_{\theta^*}) \rightarrow 0 \text{ a.s. as } n \rightarrow \infty.$$

Thus, the problem of determining whether an agent  $i$  commonly learns  $\theta$  is equivalent to that of determining whether his random type  $t_{Z_n}^i$  almost surely converges to  $t_{\theta^*}$  in the uniform-weak topology.

**Lemma 7.** *For every dataset  $\mathbf{z}$ .*

$$d_i^{UW}(t_{\mathbf{z}}, t_{\theta^*}) \leq d_P(F_{\mathbf{z}}, \delta_{\theta^*}) \quad (8)$$

*Proof.* Fix any dataset  $\mathbf{z}$ . It is useful to decompose the set of types  $T_{\mathbf{z}}$  into the Cartesian product  $\prod_{k=1}^{\infty} H_{\mathbf{z}}^k$ , where  $H_{\mathbf{z}}^1 = F_{\mathbf{z}}$  and for each  $k > 1$ ,  $H_{\mathbf{z}}^k$  is recursively defined

$$H_{\mathbf{z}}^k = \left\{ t^k \in T^k : (\text{marg}_{T^{k-1}} t^k)(H_{\mathbf{z}}^{k-1}) = 1 \text{ and } \text{marg}_{\Theta} t^k \in H_{\mathbf{z}}^1 \right\}; \quad (9)$$

that is,  $H_{\mathbf{z}}^k$  consists of the  $k$ -th order beliefs of types in  $T_{\mathbf{z}}$ . First, I show that every  $k$ -th order belief in the set  $H_{\mathbf{z}}^k$  is within  $d_P(F_{\mathbf{z}}, \delta_{\theta^*})$  (in the  $d^k$  metric<sup>44</sup>) of the  $k$ -th order belief of  $t_{\theta^*}$ .

**Claim 5.** *Define  $\delta^* = d_P(F_{\mathbf{z}}, \delta_{\theta^*})$ . For every  $k \geq 1$ ,*

$$H_{\mathbf{z}}^k \subseteq \left\{ t_{\theta^*}^k \right\}^{\delta^*} := \left\{ t^k \in T^k : d^k(t, t_{\theta^*}) \leq \delta \right\}.$$

*Proof.* Fix any  $t \in T_{\mathbf{z}}$ . By construction of  $T_{\mathbf{z}}$ , the first-order belief of type  $t$  is in the set  $F_{\mathbf{z}}$ . So it is immediate that

$$d^1(t, t_{\theta^*}) \leq d_P(F_{\mathbf{z}}, \delta_{\theta^*}) = \delta^*. \quad (10)$$

Now suppose  $H_{\mathbf{z}}^k \subseteq \{t_{\theta^*}^k\}^{\delta^*}$ . Then, since  $t^{k+1}(\{t_{\theta^*}^k\}^{\delta^*}) \geq t^{k+1}(H_{\mathbf{z}}^k) = 1$  from (9), and  $t_{\theta^*}^{k+1}(\{t_{\theta^*}^k\})$  by definition of the type  $t_{\theta^*}$ , it follows that

$$t_{\theta^*}^{k+1}(E) \leq t^{k+1}(E^{\delta^*}) + \delta^*$$

for every measurable  $E \subseteq T^k$ . Using this and (10),

$$d^{k+1}(t, t_{\theta^*}) \leq \delta^*. \quad (11)$$

as desired. □

So  $d^k(t, t_{\theta^*}) \leq \delta^*$  for every  $k$ , implying  $d_i^{UW}(t, t_{\theta^*}) = \sup_{k \geq 1} d^k(t, t_{\theta^*}) \leq \delta^*$ . □

---

<sup>44</sup>See Section 3.2.

Thus, the question of convergence of types is reduced to the question of convergence in distributions over  $\Theta$ . The remainder of the argument is now completed:

Fix any map  $t^i : \mathbf{z} \mapsto t_{\mathbf{z}}^i$  such that  $t_{\mathbf{z}}^i \in T_{\mathbf{z}}$  for every  $\mathbf{z}$ . Suppose  $M$  is uniformly consistent; then  $\sup_{\mu \in M} d(\mu_{Z_n}, \delta_{\theta^*}) \rightarrow 0$  a.s.<sup>45</sup>. It follows from Lemma 6 that

$$d_i^{UW}(t_{Z_n}^i, t_{\theta^*}) \rightarrow 0 \text{ a.s.},$$

so that agent  $i$ 's (interim) type  $t_{Z_n}^i$  almost surely converges to  $t_{\theta^*}$ . Using Lemma 7, agent  $i$  commonly learns  $\theta$ .

For the other direction, suppose  $M$  is not uniformly consistent. Then, there exist constants  $\epsilon, \delta > 0$  such that for  $n$  sufficiently large,

$$\sup_{\mu \in M} d(\mu(\mathbf{z}_n), \delta_{\theta^*}) > \epsilon \quad (12)$$

for every  $\mathbf{z}_n$  in a set  $\mathcal{Z}_n^*$  of  $P^n$ -measure  $\delta$ . Define the map  $t^i$  such that for every dataset  $\mathbf{z}_n \in \mathcal{Z}_n^*$ , agent  $i$ 's first-order belief is  $\mu(\mathbf{z}_n)$  for some  $\mu \in M$  satisfying  $d(\mu(\mathbf{z}_n), \delta_{\theta^*}) > \epsilon$  (existence guaranteed by (12)). Then  $d^1(t_{Z_n}^i, t_{\theta^*}) \not\rightarrow 0$ , so also  $d_i^{UW}(t_{Z_n}^i, t_{\theta^*}) \not\rightarrow 0$ , and it follows from Lemma 6 that agent  $i$  does not commonly learn  $\theta$ .

### 9.3 Proof of Claim 3

I prove this claim in two parts. Recall that  $U_a^{NE}$  is the set of all complete information games in which  $a$  is a Nash equilibrium. Thus, the set  $h^{-1}(U_a^{NE})$  is the set of all distributions over  $\Theta$  that induce an expected payoff in  $U_a^{NE}$ . The first claim says that  $\delta_{\theta^*} \in \text{Int}(h^{-1}(U_a^{NE}))$  if and only if  $h(F_{Z_n})$  is almost surely contained in  $U_a^{NE}$  as the quantity of data  $n$  tends to infinity.

**Claim 6.**  $\lim_{n \rightarrow \infty} h(F_{Z_n}) \subseteq U_a^{NE}$  a.s. if and only if  $\delta_{\theta^*} \in \text{Int}(h^{-1}(U_a^{NE}))$ .

*Proof. Sufficiency.* Suppose  $\delta_{\theta^*} \in \text{Int}(h^{-1}(U_a^{NE}))$ . Recall that under uniform consistency,  $W_1(F_{Z_n}, \delta_{\theta^*}) \rightarrow 0$  a.s., so that

$$\lim_{n \rightarrow \infty} F_{Z_n} \subseteq V \text{ a.s.}$$

for every open set  $V$  with  $\delta_{\theta^*} \in V$ . This implies in particular that

$$\lim_{n \rightarrow \infty} F_{Z_n} \subseteq h^{-1}(U_a^{NE}) \text{ a.s.}$$

Using continuity of  $h$  (see Lemma 4), it follows from the continuous mapping theorem that

$$\lim_{n \rightarrow \infty} h(F_{Z_n}) \subseteq U_a^{NE} \text{ a.s.}$$

---

<sup>45</sup>Uniform convergence in  $W_1$  implies uniform convergence in the Prokhorov metric  $d$ . See for example Gibbs & Su (2002).

as desired.

*Necessity.* Suppose  $\delta_{\theta^*} \notin \text{Int}(h^{-1}(U_a^{NE}))$ . Under assumption NI, there exists a constant  $\delta > 0$  independent of  $n$ , and a set  $\mathcal{Z}_n^*$  of measure  $\delta$ , such that

$$\delta_{\theta^*} \in \text{Int}(F_{\mathbf{z}_n}) \quad \forall \mathbf{z}_n \in \mathcal{Z}_n^*.$$

Consider any dataset  $\mathbf{z}_n \in \mathcal{Z}_n^*$ . Since  $\delta_{\theta^*} \notin \text{Int}(h^{-1}(U_a^{NE}))$ , necessarily  $F_{\mathbf{z}_n} \not\subseteq h^{-1}(U_a^{NE})$ . It follows that

$$\lim_{n \rightarrow \infty} P^n(\{\mathbf{z}_n : h(F_{\mathbf{z}_n}) \subseteq U_a^{NE}\}) < 1$$

as desired. □

**Claim 7.**  $\delta_{\theta^*} \in \text{Int}(h^{-1}(U_a^{NE}))$  if and only if  $u^* \in \text{Int}(U_a^{NE})$ .

*Proof.* Suppose  $u^* \in \text{Int}(U_a^{NE})$ . Then, there is an open set  $V$  such that

$$u^* \in V \subseteq U_a^{NE}.$$

Since  $h$  is continuous (see Lemma 4),  $h^{-1}(V)$  is an open set in  $\Delta(\Theta)$ . So

$$\delta_{\theta^*} \in h^{-1}(V) \subseteq h^{-1}(U_a^{NE})$$

implying that  $\delta_{\theta^*} \in \text{Int}(h^{-1}(U_a^{NE}))$ , as desired.

For the other direction, suppose towards contradiction that  $\delta_{\theta^*} \in \text{Int}(h^{-1}(U_a^{NE}))$  but  $u^* \notin \text{Int}(U_a^{NE})$ . Since  $u^*$  is on the boundary of  $U_a^{NE}$ , there exists some agent  $i$  and action  $a'_i \neq a_i$  such that

$$u_i^*(a'_i, a_{-i}) \geq u_i^*(a_i, a_{-i}).$$

Under assumption 4,  $g(\Theta)$  has nonempty intersection with  $S(i, a_i)$ , so there exists some  $\theta \in g^{-1}(S(i, a_i))$ . For every  $\epsilon > 0$ , define

$$\mu_\epsilon = (1 - \epsilon)\delta_{\theta^*} + \epsilon\delta_\theta.$$

The expected payoff under  $\mu_\epsilon$  satisfies

$$\int_U u_i(a'_i, a_{-i}) dg_*(\mu_\epsilon) > \int_U u_i(a_i, a_{-i}) dg_*(\mu_\epsilon)$$

where  $g_*(\nu)$  denotes the push forward measure of  $\nu \in \Delta(\Theta)$  under the map  $g$ . So  $a_i$  is not a best response to  $a_{-i}$  given beliefs  $\mu_\epsilon$  over  $\Theta$ , and therefore  $h(\mu_\epsilon) \notin U_a^{NE}$ . This implies also  $\mu_\epsilon \notin h^{-1}(U_a^{NE})$ . Thus the sequence  $\mu_\epsilon \rightarrow \delta_{\theta^*}$  and has the property that  $\mu_\epsilon \notin h^{-1}(U_a^{NE})$  for every  $\epsilon$ , so  $\delta_{\theta^*} \notin \text{Int}(h^{-1}(U_a^{NE}))$ , as desired. □

## 9.4 Proof of Theorem 2

*Only if:* Define  $U_{a_i^*}^R \subseteq U$  to consist of all payoffs  $u$  such that  $a_i^*$  is rationalizable for player  $i$  in the complete information game with payoffs  $u$ .

**Lemma 8.**  $u \in \text{Int} \left( U_{a_i^*}^R \right)$  if and only if  $a_i^*$  survives every round of weak strict-rationalizability in the complete information game with payoffs  $u$ .

*Proof. Only if:* Suppose  $a_i^*$  fails to survive some iteration of weak strict-rationalizability. Then, there exists a sequence of sets  $\left( W_j^k \right)_{k \geq 1}$  for every agent  $j$  satisfying the recursive description in Section 5.1, such that  $a_i^* \notin W_i^K$  for some  $K < \infty$ . To show that  $u \notin \text{Int} \left( U_{a_i^*}^R \right)$ , I construct a sequence of payoff functions  $u^n$  with  $u^n \rightarrow u$  (in the sup-metric) such that  $a_i^*$  is not rationalizable in any complete information game with payoffs along this sequence, for  $n$  sufficiently large.

For every  $n \geq 1$ , define the payoff function  $u^n$  as follows. For every agent  $j$ , let  $u_j^{n,1}$  satisfy

$$\begin{aligned} u_j^{n,1}(a_j, a_{-j}) &= u_j(a_j, a_{-j}) + \epsilon/n \quad \forall a_j \in W_j^{k-1} \text{ and } \forall a_{-j} \in A_{-j} \\ u_j^{n,1}(a_j, a_{-j}) &= u_j(a_j, a_{-j}) \text{ otherwise.} \end{aligned}$$

Recursively for  $k \geq 1$ , let  $u_j^{n,k}$  satisfy

$$\begin{aligned} u_j^{n,k}(a_j, a_{-j}) &= u_j^{n,k-1}(a_j, a_{-j}) + \epsilon/n \quad \forall a_j \in W_j^{k-1} \text{ and } \forall a_{-j} \in A_{-j} \\ u_j^{n,k}(a_j, a_{-j}) &= u_j^{n,k-1}(a_j, a_{-j}) \text{ otherwise.} \end{aligned}$$

Define  $u^n$  such that  $u_j^n := u_j^{n,K}$  for every player  $j$ .

I claim that  $a_i^*$  is not rationalizable in the complete information game with payoff function  $u^n$ , for any  $n$  sufficiently large. To show this, let us construct for every player  $j$  the sets  $(S_j^{k,n})_{k \geq 1}$  of actions surviving  $k$  rounds of iterated elimination of strictly dominated strategies given payoff function  $u^n$ , and show that for  $n$  sufficiently large,  $S_j^{k,n} = W_j^k$  for all  $k$  and every player  $j$ . I will use the following intermediate results.

**Claim 8.** *There exists  $\gamma > 0$  such that for any  $u'$  satisfying  $\|u' - u\|_\infty < \gamma$ , and for any agent  $j$ , if*

$$u_j(a_j, a_{-j}) > \max_{a'_j \neq a_j} u_j(a_j, a_{-j})$$

*then*

$$u'_j(a_j, a_{-j}) > \max_{a'_j \neq a_j} u'_j(a_j, a_{-j}).$$

*Proof.* Let  $\gamma = \frac{1}{2} \min_{i \in \mathcal{I}} \min_{a_i \in A_i} \left| u_i(a_i, a_{-i}) - \max_{a'_i \neq a_i} u_i(a'_i, a_{-i}) \right|$ , which exists by finiteness of  $\mathcal{I}$  and action sets  $A_i$ . The claim follows immediately.  $\square$

**Corollary 3.** Let  $N = \epsilon K / \gamma$ . Then, for every  $n \geq N$ , if

$$u_j(a_j, a_{-j}) > \max_{a'_j \neq a_j} u_j(a_j, a_{-j})$$

then

$$u_j^{n,k}(a_j, a_{-j}) > \max_{a'_j \neq a_j} u_j^{n,k}(a_j, a_{-j})$$

for every  $k \geq 1$ .

*Proof.* Directly follows from Claim 8, since for every  $j$ ,

$$\|u_j^{n,k} - u_j\|_\infty \leq \|u_j^n - u_j\|_\infty \leq \frac{\epsilon K}{n}$$

by construction. □

The remainder of the proof proceeds by induction. Trivially,  $S_j^{0,n} = W_j^0 = A_j$  for every  $j$  and  $n$ . Now consider any agent  $j$  and action  $a_j \in A_j$ . Suppose there exists some strategy  $\alpha_{-j} \in \Delta(A_{-j})$  such that

$$u_j(a_j, \alpha_{-j}) - \max_{a'_j \neq a_j} u_j(a'_j, \alpha_{-j}) > 0,$$

so that  $a_j$  is a strict best response to  $\alpha_{-j}$  under  $u$ . Then  $a_j \in W_j^1$ , and for  $n \geq N$ , also  $a_j \in S_j^{1,n}$  (using Corollary 3). Suppose  $a_j$  is never a strict best response, but there exists  $\alpha_{-j} \in \Delta(A_{-j})$  such that

$$u_j(a_j, \alpha_{-j}) - \max_{a'_j \neq a_j} u_j(a'_j, \alpha_{-j}) = 0.$$

If  $a_j \in W_j^1$ , then

$$u_j^n(a_j, \alpha_{-j}) - \max_{a'_j \neq a_j} u_j^n(a'_j, \alpha_{-j}) \geq u_j(a_j, \alpha_{-j}) - \max_{a'_j \neq a_j} u_j(a'_j, \alpha_{-j}),$$

so also  $a_i \in S_i^{1,n}$  for  $n \geq N$ . If  $a_j \notin W_j^1$ , then for  $n \geq N$ , there exists an action  $a'_j \neq a_j$  such that  $u_j(a'_j, \alpha_{-j}) = u_j(a_j, \alpha_{-j})$ , but  $u_i^n(a'_j, \alpha_{-j}) > u_i^n(a_j, \alpha_{-j})$ . So  $a_j \notin S_j^{1,n}$ . No other actions survive to either  $W_j^1$  or  $S_j^{1,n}$ . Thus  $S_j^{1,n} = W_j^1$  for all  $n \geq N$ .

This argument can be repeated for arbitrary  $k$ . Suppose  $S_j^{k,n} = W_j^k$  for every  $j$  and  $n \geq N$ , and consider any action  $a_j \in S_j^{k,n}$ . If there exists some strategy  $\alpha_{-j} \in \Delta(S_{-j}^{k,n})$  such that

$$u_j(a_j, \alpha_{-j}) - \max_{a'_j \neq a_j} u_i(a'_j, \alpha_{-j}) > 0,$$

then  $a_j \in W_j^{k+1}$ , and for  $n \geq N$ , also  $a_j \in S_j^{k+1,n}$  (using Corollary 3). Suppose  $a_j$  is not a strict best response to any  $\alpha_{-j} \in \Delta(S_{-j}^{k,n})$ , but there exists  $\alpha_{-j} \in \Delta(S_{-j}^{k,n})$  such that

$$u_j(a_j, \alpha_{-j}) - \max_{a'_j \neq a_j} u_j(a'_j, \alpha_{-j}) = 0.$$

Then, if  $a_j \in W_j^{k+1}$ , action  $a_j$  is a strict best response to  $a_{-j}$  under  $u^n$ , so  $a_j \in S_j^{k+1,n}$ . Otherwise, if  $a_j \notin W_j^{k+1}$ , then there exists some  $a'_j \in W_j^{k+1}$  such that  $u_j^n(a'_j, \alpha_{-j}) > u_j^n(a_j, \alpha_{-j})$ , so also  $a_j \notin S_j^{k+1,n}$ . No other actions survive to either  $W_j^{k+1}$  or  $S_j^{k+1,n}$ , so  $S_j^{k+1,n} = W_j^{k+1}$  for  $n \geq N$ . Therefore  $S_j^{k,n} = W_j^k$  for every  $k$  and  $n \geq N$ , and in particular  $S_j^{K,n} = W_j^K$  for  $n \geq N$ . Since  $a_j \notin W_j^K$ , also  $a_j \notin S_j^{\infty,n}$  for  $n$  sufficiently large, as desired.

Finally, notice that by construction  $\|u^n - u\|_\infty \leq \frac{\epsilon K}{n}$ , which can be rewritten

$$\|u^{n(\epsilon')} - u\|_\infty \leq \epsilon'$$

where  $n(\epsilon') := \frac{\epsilon K}{\epsilon'}$ . Thus, for every  $\epsilon' \geq 0$ , the payoff function  $u_i^{n(\epsilon')} \in B_{\epsilon'}(u)$ , but  $a_i$  is not rationalizable in the complete information game with payoff function  $u_i^{n(\epsilon')}$ . So  $u \notin \text{Int}\left(U_{a_i^*}^R\right)$ , as desired.

*If:* Suppose  $u \notin \text{Int}\left(U_{a_i^*}^R\right)$ . Consider any sequence of payoff functions  $u^n \rightarrow u$ . Since action sets are finite, there is a finite number of possible orders of elimination. This implies existence of a subsequence along which the same order of iterated elimination of strategies removes  $a_i^*$ . Choose any one-at-time iteration of this order of elimination. Then,  $a_i^*$  fails to survive this order of elimination given the limiting payoffs  $u$ , so it is not weakly strict-rationalizable.  $\square$

Next, I show that  $a_i$  is robust to inference only if the true payoff function  $u^*$  is in the interior of  $U_{a_i^*}^R$ .

**Lemma 9.**  $a_i^*$  is robust to inference only if  $u^* \in \text{Int}\left(U_{a_i^*}^R\right)$ .

*Proof.* The following claim will be useful.

**Claim 9.**  $u^* \in \text{Int}\left(U_{a_i^*}^R\right)$  if and only if  $\delta_{\theta^*} \in \text{Int}(h^{-1}(U_a))$ .

*Proof.* See proof of Claim 6.  $\square$

Suppose  $u^* \notin \text{Int}(U_{a_i^*}^R)$ ; then, using Claim 9, also  $\delta_{\theta^*} \notin \text{Int}(h^{-1}(U_{a_i^*}^R))$ . Under assumption NI, there is a constant  $\epsilon > 0$  such that  $\delta_{\theta^*} \in \text{Int}(F_{\mathbf{z}_n})$  for at least an  $\epsilon$ -measure of datasets. Consider any such dataset. Then,  $\delta_{\theta^*} \notin \text{Int}\left(h^{-1}(U_{a_i^*}^R)\right)$ , implies that  $F_{\mathbf{z}_n} \not\subseteq h^{-1}(U_a)$ . Fix any  $u \in F_{\mathbf{z}_n} \setminus h^{-1}(U_{a_i^*}^R)$ . Then  $a_i^*$  is not rationalizable in the complete information game with payoffs  $u$ , so it is also not rationalizable for the type with common certainty in  $u$ .  $\square$

If: If  $a_i^*$  is strongly strict-rationalizable, then there exists a family of sets  $(V_j^k)_{j \in \mathcal{I}}$  is closed under  $\delta$ -strict best reply for some  $\delta \geq 0$ ; that is, for every  $a_j \in V_j^k$ , there exists a distribution  $\alpha_{-j} \in \Delta(V_{-j}^k)$  such that

$$u_j^*(a_j, \alpha_{-j}) > \max_{a'_j \neq a_j} u_j^*(a'_j, \alpha_{-j}) + \delta.$$

Recall the following fixed-point property of the set of rationalizable actions:

**Lemma 10** (Dekel, Fudenberg & Morris (2007)). *Fix any type profile  $(t_j)_{j \in \mathcal{I}}$ . Consider any family of sets  $V_j \subseteq A_j$  such that every action  $a_j \in V_j$  is a best reply to a distribution  $\pi \in \Delta(\Theta \times T_{-j} \times A_{-j})$  that satisfies  $\text{marg}_{\Theta \times T_{-j}} \pi = g(t_j)$  and  $\pi(a_{-j} \in V_{-j}[t_{-j}]) = 1$ . Then,  $V_j \subseteq S_j^\infty[t_j]$  for every agent  $j$ .*

Fix any  $\epsilon > 0$ . Then, for every agent  $j$  and type  $t_j$  with common certainty in  $B_\epsilon(u^*)$ , we have that

$$\begin{aligned} \int u_j(a_j, \alpha_{-j}, \theta) d\kappa_j(t_j) - \max_{a'_j \neq a_j} \int u_j(a'_j, \alpha_{-j}, \theta) d\kappa_j(t_j) \\ \geq \inf_{u \in B_\epsilon(u^*)} \left( u_j(a_j, \alpha_{-j}) - \max_{a'_j \neq a_j} u_j(a'_j, \alpha_{-j}) \right) \\ \geq \delta - 2\epsilon, \end{aligned}$$

which is positive for any  $\epsilon \leq \delta/2$ . So the family of sets  $(V_j^k)_{j \in \mathcal{I}}$  satisfies the conditions in Lemma 10 when  $\epsilon$  is sufficiently small, and it follows that  $a_i^* \in S_i^\infty[t_j]$ , as desired.

## 9.5 Proof of Proposition 2

To simplify notation, set  $\delta := \delta_{a^*}^{NE}$ . By assumption,  $\delta \geq 0$ .

**Lemma 11.**  $B_{\delta/2}(u^*) \subseteq U_{a^*}^{NE}$ .

*Proof.* Consider any payoff function  $u'$  satisfying

$$\|u' - u^*\|_\infty \leq \frac{\delta}{2}. \tag{13}$$

Then for every agent  $i$ ,

$$\begin{aligned} u'_i(a_i^*, a_{-i}^*) - u'_i(a'_i, a_{-i}^*) &= \underbrace{u'_i(a_i^*, a_{-i}^*) - u_i^*(a_i^*, a_{-i}^*)}_{\geq -\delta/2} \\ &\quad + \underbrace{u_i^*(a_i^*, a_{-i}^*) - u_i^*(a'_i, a_{-i}^*)}_{> \delta} + \underbrace{u_i^*(a'_i, a_{-i}^*) - u'_i(a'_i, a_{-i}^*)}_{\geq -\delta/2} \geq 0. \end{aligned}$$



where  $u_i^*(a_i^*, a_{-i}^*) - u_i^*(a_i', a_{-i}^*) > \delta$  follows from the assumption that  $a^*$  is a  $\delta$ -strict NE in the complete information game with payoffs  $u^*$ , and the other two bounds follow from 13. So  $a^*$  is a NE in the complete information game with payoffs  $u'$ , implying that  $u' \in U_{a^*}^{NE}$ .  $\square$

It follows from Lemma 2 that common certainty in  $B_{\delta/2}(u^*)$  is a sufficient condition for  $a^*$  to be a Bayesian Nash equilibrium. Thus,

$$\begin{aligned} p_n^{NE}(a^*) &\geq P^n(\{\mathbf{z}_n : h(F_{\mathbf{z}_n}) \subseteq B_{\delta/2}(u^*)\}) \\ &= P^n\left(\left\{\mathbf{z}_n : \sup_{\mu \in M} \|h(\mu_{\mathbf{z}_n}) - u^*\|_\infty \leq \delta/2\right\}\right) \\ &= 1 - P^n\left(\left\{\mathbf{z}_n : \sup_{\mu \in M} \|h(\mu_{\mathbf{z}_n}) - u^*\|_\infty > \delta/2\right\}\right) \\ &\geq 1 - \frac{2}{\delta} \mathbb{E} P^n\left(\sup_{\mu \in M} \|h(\mu_{\mathbf{z}_n}) - u^*\|_\infty\right) \end{aligned}$$

using Markov's inequality in the final line.

## 9.6 Proof of Proposition 3

To simplify notation, set  $\delta := \delta_{a_i^*}^R$ . By assumption,  $\delta \geq 0$ .

**Lemma 12.**  $B_{\delta/2}(u^*) \subseteq U_{a_i^*}^R$ .

*Proof.* Consider any payoff function  $u'$  satisfying

$$\|u' - u^*\|_\infty \leq \frac{\delta}{2}. \quad (14)$$

By definition of  $\delta_{a_i^*}^R$ , there exists a family of sets  $(R_i)_{i \in \mathcal{I}}$  with the property that for every agent  $j$  and action  $a_j \in R_j$ , there is an action  $\alpha_{-j}[a_j] \in \Delta(R_{-j})$  satisfying

$$u_i^*(a_j, \alpha_{-j}[a_i]) > u_i^*(a_i', \alpha_{-j}[a_j]) + \delta \quad \forall a_i' \neq a_i. \quad (15)$$

I will show that  $(R_j)_{j \in \mathcal{I}}$  satisfies the conditions in Lemma 10 for any type profile  $(t_j)_{j \in \mathcal{I}}$ , where every  $t_j$  has common certainty in  $B_{\delta/2}(u^*)$ . Fix an arbitrary agent  $j$ , and type  $t_j$  with common certainty in  $B_{\delta/2}(u^*)$ . Define the distribution  $\pi \in \Delta(\Theta \times T_{-j} \times A_{-j})$  such that  $\text{marg}_{\Theta \times T_{-j}} \pi = \kappa_j(t_j)$  and  $\text{marg}_{A_{-j}} \pi = \alpha_{-j}[a_j]$ , noting that since  $\alpha_{-j}[a_j] \in \Delta(R_{-j})$ , this implies also that  $\pi(a_{-j} \in R_{-j}) = 1$ .

Since by assumption,  $t_j$  has common certainty in  $B_{\delta/2}(u^*)$ , the support of  $\text{marg}_\Theta \kappa(t_j)$  is contained in  $B_{\delta/2}(u^*)$ . So the expected payoff from playing  $a_j$  exceeds the expected payoff from playing  $a_j' \neq a_j$  by at least

$$\inf_{u \in B_{\delta/2}(u^*)} (u(a_j, \alpha_{-j}) - u(a_j', \alpha_{-j})) \geq -\frac{\delta}{2} \quad (16)$$

It follows that

$$\begin{aligned} \int u_j(a_j, \alpha_{-j}, \theta) d\pi - \int u_j(a'_j, \alpha_{-j}, \theta) d\pi &= \underbrace{\int u_j(a_j, \alpha_{-j}, \theta) d\pi - u_j^*(a_j, \alpha_{-j}, \theta)}_{\geq -\frac{1}{2}\delta} \\ &+ \underbrace{u_j^*(a_j, \alpha_{-j}, \theta) - u_j^*(a'_j, \alpha_{-j}, \theta)}_{>\delta} + \underbrace{\int u_j^*(a'_j, \alpha_{-j}, \theta) d\pi - u_j(a'_j, \alpha_{-j}, \theta)}_{\geq -\frac{1}{2}\delta} \geq 0, \end{aligned}$$

using the inequalities in (15) and (16). It follows that  $a_j$  is a best response to  $\alpha_{-j}$  given distribution  $\pi$ . Repeating this argument for every agent  $j$ , action  $a_j \in R_j$ , and type  $t_j$  with common certainty in  $B_{\delta/2}(u^*)$ , it follows from Lemma 10 that  $R_j \subseteq S_j^\infty[t_j]$  for every agent  $j$ . Since  $a_i^* \in R_i$ , also  $a_i^* \in S_i^\infty[t_i]$ , as desired.  $\square$

It follows from this lemma that  $F_{\mathbf{z}} \subseteq B_{\delta/2}(u^*)$  is a sufficient condition for  $a_i^*$  to be rationalizable in every game in  $\mathcal{G}(\mathbf{z})$ . Thus,

$$\begin{aligned} p_n^R(i, a_i^*) &\geq P^n(\{\mathbf{z}_n : h(F_{\mathbf{z}_n}) \subseteq B_{\delta/2}(u^*)\}) \\ &= P^n\left(\left\{\mathbf{z}_n : \sup_{\mu \in M} \|h(\mu_{\mathbf{z}}) - u^*\|_\infty \leq \delta/2\right\}\right) \\ &= 1 - P^n\left(\left\{\mathbf{z}_n : \sup_{\mu \in M} \|h(\mu_{\mathbf{z}}) - u^*\|_\infty > \delta/2\right\}\right) \\ &\geq 1 - \frac{2}{\delta} \mathbb{E}_{P^n} \left( \sup_{\mu \in M} \|h(\mu_{\mathbf{z}_n}) - u^*\|_\infty \right) \end{aligned}$$

using Markov's inequality in the final line.

## Proof of Corollary 1

From properties of the least-squares estimator,

$$\begin{aligned} \mathbb{E}(|\hat{\beta}_1 - \beta_1|^2) &= \text{Var}(\hat{\beta}_1) \leq \sum_j \text{Var}(\hat{\beta}_j) \\ &= \sigma^2 \sum_k \mathbb{E} \left( (\mathbf{X}^T \mathbf{X})_{kk}^{-1} \right) \\ &= \sigma^2 \mathbb{E} \left( \text{tr}(\mathbf{X}^T \mathbf{X})^{-1} \right) \\ &= \sigma^2 \mathbb{E} \left( \sum_i \lambda_i^{-1} \right) \\ &\leq \sigma^2 p(\sqrt{n} + \sqrt{p}) \end{aligned}$$

where the final line follows from Gordon's theorem for Gaussian matrices (see e.g. Vershynin (2012)). Let  $K$  be the Lipschitz constant of the map  $g : \Theta \rightarrow U$  (assuming the sup-norm on  $U$  and the Euclidean norm on  $\Theta$ ),

$$\begin{aligned} \mathbb{E} \left( \sup_{\mu \in M} \|h(\mu_{Z_n}) - u^*\|_\infty \right) &\leq K \mathbb{E} \left( |\hat{\beta}_1 - \beta_1|^2 + \phi_n^2 \right) \\ &\leq K \left( \sigma^2 p (\sqrt{n} + \sqrt{p}) + \phi_n^2 \right) \end{aligned}$$

and the desired bound follows directly from Proposition 2.

### Proof of Proposition 4

The argument below is for Nash equilibrium; the argument for rationalizability follows analogously. For every inference rule  $\mu \in M$ , define

$$X_\mu^n = \mathbb{1} (h(\mu_{Z_n}) \notin U_a^{NE})$$

to take value 1 if the expected payoff under the (random) distribution  $\mu(Z_n)$  is outside the set  $U_a^{NE}$ . Write  $F_\mu^n$  for the marginal distribution of random variable  $X_\mu^n$ , and  $F_M^n$  for the joint distribution of random variables  $(X_\mu^n)_{\mu \in M}$ . Enumerate the inference rules in  $M$  by  $\mu_1, \dots, \mu_k$ .

By Sklar's theorem, there exists a copula  $C : [0, 1]^k \rightarrow [0, 1]$  such that

$$F_M^n(x_1, \dots, x_k) = C(F_{\mu_1}^n(x_1), \dots, F_{\mu_k}^n(x_k))$$

for every  $x_1, \dots, x_k$ . Using the Frechet-Hoeffding bound,

$$1 - K + \sum_{k=1}^K F_{\mu_k}^n(x_k) \leq C(F_{\mu_1}^n(x_1), \dots, F_{\mu_k}^n(x_k)) \leq \min_{k \in \{1, \dots, K\}} F_{\mu_k}^n(x_k).$$

From Lemma 2,  $p_n^{NE}(a) = F_M^n(0, \dots, 0)$ . It follows that

$$1 - K + \sum_{i=1}^K F_{\mu_i}^n(0) \leq p_n^{NE}(a) \leq \min_{k \in \{1, \dots, K\}} F_{\mu_k}^n(0). \quad (17)$$

Finally, since every  $X_\mu^n \sim \text{Ber}(1 - p_\mu^n)$ , (17) implies

$$1 - \sum_{\mu \in M} p_{\mu, n}^{NE} \leq p_n^{NE}(a) \leq 1 - \min_{\mu \in M} p_{\mu, n}^{NE}$$

as desired.

## Appendix D: An example illustrating the fragility of weak strict-rationalizability

In the following, I present a game in which an action is weakly strict-rationalizable, but fails to be rationalizable along a sequence of perturbed types in the uniform-weak topology.

Consider a game with four players. Each has two actions,  $a$  and  $b$ . Throughout I will use, for example,  $abab$  to denote choice of  $a$  by players 1 and 3, and  $b$  by players 2 and 4. Let payoffs be defined as follows. Player 1's payoffs satisfy

$$u_1(axxx) = \begin{cases} 1 & \text{if } xxx = aaa \text{ or } bbb \\ 0 & \text{otherwise.} \end{cases}$$

$$u_1(bxxx) = \begin{cases} 0 & \text{if } xxx = aaa \text{ or } bbb \\ 1 & \text{otherwise.} \end{cases}$$

That is, player 1 wants to play  $a$  if players 2-4 are all playing  $a$  or all playing  $b$ , and he wants to play  $b$  otherwise. The payoffs to players 2-4 are independent of player 1's action. They are described below (where rows correspond to player 2's actions, columns to player 3, and choice of matrices to player 4), with player 1's payoffs omitted, so that the first coordinate corresponds to player 2's payoff:

$$\begin{array}{cc} & \begin{array}{cc} a & b \end{array} \\ \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 1, 1, 0 & 0, 0, 0 \\ 0, 0, 0 & 0, 0, 0 \end{array} \end{array} \quad \begin{array}{cc} & \begin{array}{cc} a & b \end{array} \\ \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 0, 0, 0 & 0, 0, 0 \\ 0, 0, 0 & 1, 1, 0 \end{array} \end{array} \quad (18)$$

(a) (b)

That is, if player 4 chooses action  $a$ , then players 2 and 3 prefer coordination on  $a$ ; and if player 4 chooses  $b$ , then players 2 and 3 prefer coordination on  $b$ .

Let us first consider the case in which the true payoffs are common certainty, so that this is a game of complete information (denote the payoffs by  $u$ ). Then,  $a$  is rationalizable for player 1. Not only is it rationalizable, but:

- there is a constant  $\epsilon > 0$  such that  $a$  is rationalizable for player 1 in every game  $u'$  with  $\|u' - u\|_\infty \leq \epsilon$ ; that is, rationalizability is preserved on an open set of complete information games.
- $a$  is weakly strict-rationalizable.
- although  $a$  is not strongly strict-rationalizable, it fails to survive this process for the reason that *none* of player 4's actions survive the first round of elimination.<sup>46</sup>

---

<sup>46</sup>In particular,  $a$  is strongly strict-rationalizable in either game in which one of player 4's actions is dropped.

Let  $t_1$  be the type with common certainty in  $u$ . I will now show that there exists a sequence of types  $t_1^n$  such that  $t_1^n \rightarrow t_1$  in the uniform-weak topology, but  $a$  fails to be rationalizable for agent 1 infinitely many times along this sequence. The sequence of types  $t_1^n$  will moreover have the property that every  $t_1^n$  believes that an  $\epsilon_n$ -neighborhood of  $u$  is common certainty, where  $\epsilon_n \rightarrow 0$  as  $n \rightarrow \infty$ .

Define  $t_1^n$  to satisfy two conditions. First, player 1 is certain<sup>47</sup> that: player 2 is certain that the payoffs in (18) are

$$\begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 1, 1, -\epsilon_n & 0, 0, -\epsilon_n \\ 0, 0, -\epsilon_n & 0, 0, -\epsilon_n \end{array}
 \end{array}
 \qquad
 \begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} -\epsilon_n, -\epsilon_n, 0 & 0, -\epsilon_n, 0 \\ 0, 0, 0 & 1, 1, 0 \end{array}
 \end{array}
 \tag{19}$$

(a) (b)

and player 2 is certain, moreover, that player 4 is certain of these payoffs. Second, player 1 is certain that: player 3 is certain that the payoffs in (18) are

$$\begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 1, 1, 0 & 0, 0, 0 \\ -\epsilon_n, -\epsilon_n, 0 & -\epsilon_n, -\epsilon_n, 0 \end{array}
 \end{array}
 \qquad
 \begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 0, 0, -\epsilon_n & 0, 0, -\epsilon_n \\ 0, 0, -\epsilon_n & 1, 1, -\epsilon_n \end{array}
 \end{array}
 \tag{20}$$

(a) (b)

and player 3 is certain, moreover, that player 4 is certain of these payoffs.

Let us now consider the rationalizable actions for players 2 and 3. If player 4 is certain that payoffs are as in (19), then action  $b$  is his uniquely rationalizable action. So player 2, with the beliefs described above, believes with probability 1 that player 4 will play  $b$ . Since he is himself certain of the payoffs in (19), action  $b$  is his uniquely rationalizable action. By a similar argument, if player 4 is certain that payoffs are as in (20), then action  $a$  is uniquely rationalizable. So player 3, with the beliefs described above, believes with probability 1 that player 4 will play  $a$ , and thus considers  $a$  to be his uniquely rationalizable action as well.

So player 1 is certain that player 2 will play  $b$  and that player 3 will play  $a$ . It follows that his uniquely rationalizable action is  $b$ . Since this argument is valid for every  $\epsilon_n > 0$ , action  $a$  is not rationalizable for player 1 of type  $t_i^n$  for any  $n$ . But every  $t_i^n$  believes that  $B_{\epsilon_n}(u)$  is common certainty, so  $t_i^n \rightarrow t_i$  in the uniform-weak topology.

---

<sup>47</sup>Believes with probability 1.

## References

- Acemoglu, Daron, Victor Chernozhukov & Muhamet Yildiz. 2015. “Fragility of Asymptotic Agreement under Bayesian Learning.” *Theoretical Economics* .
- Al-Najjar, Nabil. 2009. “Decisionmakers as Statisticians: Diversity, Ambiguity, and Learning.” *Econometrica* .
- Aumann, Robert J. 1976. “Agreeing to Disagree.” *The Annals of Statistics* .
- Battigalli, Pierpaolo & Marciano Siniscalchi. 2003. “Rationalization and Incomplete Information.” *Advances in Theoretical Economics* .
- Billot, Antoine, Itzhak Gilboa, Dov Samet & David Schmeidler. 2005. “Probabilities as Similarity-Weighted Frequencies.” *Econometrica* .
- Brandenburger, Adam & Eddie Dekel. 1993. “Hierarchies of Belief and Common Knowledge.” *Journal of Economic Theory* .
- Carlin, Bruce I., Francis A. Longstaff & Kyle Matoba. 2014. “Disagreement and asset prices.” *Journal of Financial Economics* .
- Carlin, Bruce I., Shimon Kogan & Richard Lowery. 2013. “Trading Complex Assets.” *The Journal of Finance* .
- Carlsson, Hans & Eric van Damme. 1993. “Global Games and Equilibrium Selection.” *Econometrica* 61(5):989–1018.
- Chen, Yi-Chun, Alfredo di Tillio, Eduardo Faingold & Siyang Xiong. 2010. “Uniform topologies on types.” *Theoretical Economics* .
- Cripps, Martin, Jeffrey Ely, George Mailath & Larry Samuelson. 2008. “Common Learning.” *Econometrica* .
- Dekel, Eddie, Drew Fudenberg & David Levine. 2004. “Learning to Play Bayesian Games.” *Games and Economic Behavior* .
- Dekel, Eddie, Drew Fudenberg & Stephen Morris. 2006. “Topologies on Types.” *Theoretical Economics* .
- Dekel, Eddie, Drew Fudenberg & Stephen Morris. 2007. “Interim Correlated Rationalizability.” *Theoretical Economics* .
- Esponda, Ignacio. 2013. “Rationalizable conjectural equilibrium: A framework for robust predictions.” *Theoretical Economics* .

- Eyster, Erik & Michele Piccione. 2013. “An Approach to Asset Pricing Under Incomplete and Diverse Preferences.” *Econometrica* .
- Fudenberg, Drew, David Kreps & David Levine. 1988. “On the Robustness of Equilibrium Refinements.” *Journal of Economic Theory* .
- Gayer, Gabrielle, Itzhak Gilboa & Offer Lieberman. 2007. “Rule-Based and Case-Based Reasoning in Housing Prices.” *The B.E. Journal of Theoretical Economics* .
- Geanakoplos, John & Heraklis Polemarchakis. 1982. “We Can’t Disagree Forever.” *Journal of Economic Theory* .
- Gibbs, Alison & Francis Su. 2002. “On Choosing and Bounding Probability Metrics.” *International Statistic Review* .
- Gilboa, Itzhak & David Schmeidler. 2003. “Inductive Inference: An Axiomatic Approach.” *Econometrica* .
- Gilboa, Itzhak, Larry Samuelson & David Schmeidler. 2013. “Dynamics of inductive inference in a unified framework.” *Journal of Economic Theory* .
- Gilboa, Itzhak, Offer Lieberman & David Schmeidler. 2006. “Empirical Similarity.” *Review of Economics and Statistics* .
- Hansen, Lars Peter. 2014. “Uncertainty Inside and Outside Economic Models.” *Journal of Political Economy* .
- Hansen, Lars Peter & Thomas J. Sargent. 2007. *Robustness*. Princeton University Press.
- Hansen, Lars Peter & Thomas Sargent. 2010. “Fragile beliefs and the price of uncertainty.” *Quantitative Economics* .
- Hansen, Lars Peter & Thomas Sargent. 2012. “Three Types of Ambiguity.” *Journal of Monetary Economics* .
- Kajii, Atsushi & Stephen Morris. 1997. “The Robustness of Equilibria to Incomplete Information.” *Econometrica* .
- Kandel, Eugene & Neil Pearson. 1995. “Differential Interpretation of Information and Trade in Speculative Markets.” *Journal of Political Economy* .
- Kohlberg, Elon & Jean-Francois Mertens. 1986. “On the Strategic Stability of Equilibria.” *Econometrica* .

- Mankiw, Gregory, Ricardo Reis & Justin Wolfers. 2004. "Disagreement about inflation expectations." *NBER Macroeconomics Annual 2003* .
- Marlon, J.R., A. Leiserowitz & G. Feinberg. 2013. Scientific and Public Perspectives on Climate Change. Technical report Yale Project on Climate Change Communication.
- Mertens, Jean-Francois & Shmuel Zamir. 1985. "Formulation of Bayesian Analysis for Games with Incomplete Information." *International Journal of Game Theory* .
- Monderer, Dov & Dov Samet. 1989. "Approximating Common Knowledge with Common Beliefs." *Games and Economic Behavior* .
- Morris, Stephen, Rafael Rob & Hyun Song Shin. 1995. "p-Dominance and Belief Potential." *Econometrica* .
- Morris, Stephen & Satoru Takahashi. N.d. "Strategic Implications of Almost Common Certainty of Payoffs."
- Morris, Stephen, Satoru Takahashi & Olivier Tercieux. 2012. "Robust Rationalizability Under Almost Common Certainty of Payoffs." *The Japanese Economic Review* .
- Myerson, Roger B. 1978. "Refinements of the Nash equilibrium concept." *International Journal of Game Theory* .
- Rubinstein, Ariel. 1989. "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"." *American Economic Review* 79.
- Selten, Richard. 1975. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *International Journal of Game Theory* .
- Steiner, Jakub & Colin Stewart. 2008. "Contagion through Learning." *Theoretical Economics* .
- Vershynin, Roman. 2012. *Compressed Sensing, Theory and Applications*. Cambridge University Press chapter 5.
- Weinstein, Jonathan & Muhamet Yildiz. 2007. "A Structure Theorem for Rationalizability with Application to Robust Prediction of Refinements." *Econometrica* .
- Wiegel, David. 2009. "How many Southern whites believe Obama was born in America?" *Washington Independent* .



Yu, Jialin. 2011. "Disagreement and return predictability of stock portfolios." *Journal of Financial Economics* .