Games of Incomplete Information Played by Statisticians*

Annie Liang†

August 10, 2016

Abstract

The common prior assumption is a convenient restriction on beliefs in games of incomplete information, but conflicts with evidence that players 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. A key assumption imposes that while players may interpret data in different ways, they have common certainty in the predictions induced by a class of interpretations. The main results characterize which rationalizable actions and Nash equilibria can be predicted when agents observe a finite quantity of data, and how much data is needed to predict different solutions. This quantity, which I refer to as the robustness of the solution, is shown to depend crucially on the degree of strictness of the solution and the "complexity" of inference from data.

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 modeling beliefs gives players a common prior belief over states of the world, and assumes that they use Bayesian updating to form a posterior belief given new information.[1] This approach is known

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

†Microsoft Research and the University of Pennsylvania. Email: liang.annie.h@gmail.com. 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.

[1]The related, stronger, notion of rational expectations assumes moreover that this common prior distribution is in fact the "true" distribution shared by the modeler.

to have strong implications, for example that beliefs that are commonly known must be identical (Aumann 1976), and that repeated communication of beliefs will eventually lead to agreement (Geanakoplos & Polemarchakis 1982). These statements conflict not only with considerable empirical evidence of public and persistent disagreement, but also with the more basic, day-to-day, experience that individuals interpret information in different ways.[2]

This paper generalizes the standard approach by allowing players to form beliefs based on a *set* of learning rules $\mathcal{F}$, where a learning rule is any function f that maps *data* (a sequence of signals) into a belief distribution over the parameter space (a first-order belief). In the main part of the paper, I focus on sets of learning rules that asymptotically recover the true parameter (via a statistical consistency condition defined in Section 3). If all players are Bayesian learners, then the typical set $\mathcal{F}$ is identified with a set of subjective prior beliefs, and the common prior approach is nested as the case in which $\mathcal{F}$ is a singleton. Other interesting choices for $\mathcal{F}$ include sets of different frequentist estimators—for example, players may be case-based learners (Gilboa & Schmeidler 1995, Gilboa, Postlewaite & Schmeidler 2008) who predict unknown outcomes based on past similar situations, but perceive similarity diffferently. The beliefs induced by learning rules in $\mathcal{F}$ describe a *plausible range of uncertainty*. I impose a key assumption to structure the approach: given any realization of the data, players have common certainty in the beliefs induced by learning rules from $\mathcal{F}$.

This approach produces a learning-based refinement on hierarchies of beliefs that is weakly more permissive than the common prior assumption, but not so flexible that "anything goes". Section 3 shows that as the quantity of observed data gets large, beliefs and strategic behavior resemble those in a limit complete information game. Specifically: players commonly learn the true value of the parameter (Proposition 1); the set of (Bayesian Nash) equilibria almost surely converges to the set of strict equilibria in the complete information game (Theorem 1); and the set of (interim correlated) rationalizable actions almost surely converges to a set that is sandwiched between the set of strict rationalizable actions, and a set that I define as "weakly" strict-rationalizable actions (Theorem 2).[3]

Although behavior is constrained asymptotically, the next set of results show that

[2] In financial markets, players 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). players publicly disagree also in matters of politics (Wiegel 2009) and climate change (Marlon, Leiserowitz & Feinberg 2013).

[3] The difference between these two notions of strict rationalizability regards the order of elimination, and is of independent interest in its own right: an extended discussion is relegated to the appendix.

predictions given small finite quantities of data can be substantially different from the limit complete information game. Section 4 studies when limit behavior fails to be a good approximation, and what determines this. Formally, the "robustness" of an equilibrium profile to inference from $n < \infty$ observations is defined as the probability that the profile is an equilibrium when player beliefs are based on a (random) dataset of size n, and the robustness of a rationalizable action is defined likewise. Proposition 2 provides a lower bound on these probabilities for every quantity of data n, with the interpretation that solutions that hold for a larger measure of size-n datasets are more robust. This bound turns out to depend on two important features:

First, the bound depends on the speed at which learning rules in $\mathcal{F}$ jointly learn the true value of the parameter. This speed depends both on the quantity of data required by each *individual* learning rule in order to recover the true parameter—which is determined by the "complexity" of the data relative to f, as measured for example by the dimensionality of the signal space—and also on the degree of correlation across the beliefs induced by learning rules in $\mathcal{F}$. I refer to this latter channel as "opinion diversity," and quantify its effect on robustness in Section 5.1. (b) Second, the bound depends on a cardinal measure of strictness of the solution. Say that an action profile is a *δ-strict NE* if each player's prescribed action is at least δ better than his next best action; and say that an action profile is *δ-strict rationalizable* if it can be rationalized by a chain of best responses, in which each action yields at least δ over the next best alternative. This parameter δ turns out to determine how much estimation error the solution can withstand—the lower the degree of strictness (the smaller the parameter δ), the slower convergence is.

These comparative statics are, in my view, a key advantage to modeling beliefs using the proposed framework. They show that when players form beliefs from data using different learning rules, then new channels—in particular, the amount of common knowledge over how to interpret data, and the "dimensionality" or "complexity" of the learning problem—emerge as determinants of strategic behavior.

These channels are complementary to (and distinct from) the traditional channel of private information, and have new implications for informational design. Section 7 asks within the proposed model: How might a designer be able to manipulate behavior either by choosing the nature of public information, or the way in which individuals interpret it? I present several examples in which provision of extraneous public information prevents agents from coordinating.

The final sections proceed as follows: Section 8 examines modeling choices made in the main text—specifically, the assumption of uniform consistency (Section 8.1), the assumption of common data (Section 8.2), and the assumption of a "true" value of the unknown parameter (Section 8.3)—and discusses the extent to which these assumptions can be relaxed. Section 9 surveys the related literature, placing this paper between the literature that studies the robustness of equilibrium—e.g. Fudenberg, Kreps & Levine (1988), Carlsson & van Damme (1993), Kajii & Morris (1997), and Weinstein & Yildiz (2007)—and the literature that studies the asymptotic properties of learning from data—e.g. Cripps, Ely, Mailath & Samuelson (2008), Al-Najjar (2009), and Acemoglu, Chernozhukov & Yildiz (2015). Section 10 concludes.

2 Preliminaries and Notation

2.1 The game

Fix a finite set of players $\mathcal{I}$ and a finite action set A_i for each player $i \in \mathcal{I}$. As usual, write $A = \prod_{i \in \mathcal{I}} A_i$. Then, the set of possible games can be identified with $U := \mathbb{R}^{|\mathcal{I}| \times |A|}$. Let $\Theta \subset \mathbb{R}^k$ be a compact subset of finite-dimensional Euclidean space, which is related to payoffs under a bounded and Lipschitz continuous embedding[4,5]

$$g : \Theta \to U.$$

For notational convenience, I will describe players as having uncertainty over the parameter space Θ instead of the payoff space U. Throughout this paper, the true value of the parameter is denoted by θ^* and the true payoffs by $u^* = g(\theta^*)$.

Remark 1. In some contexts, the concept of "true" payoffs may not have meaning, and we might prefer to think instead of a true *distribution* over payoffs. In this case, the parameter θ can be interpreted as indexing a family of distributions over payoffs, and the map g as taking parameters into expected payoffs (see Section 8.3).

2.2 Beliefs

A complete description of a player's uncertainty is identified with a hierarchy of beliefs, or more simply, a *type*.

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

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

[4] A map is an embedding if it is a homeomorphism onto its image.

[5] This map g can be interpreted as capturing the known information about the structure of payoffs.

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 Θ (describing the player's uncertainty over Θ, his uncertainty over his opponents' uncertainty over Θ, and so forth), and is referred to as a *type*.

This approach can be generalized for I players, 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 player i, there is a subset of types T_i^* (that satisfy the property of *coherency*[6]) and a function $\kappa_i^* : T_i^* \to \Delta\left(\Theta \times T_{-i}^*\right)$ such that $\kappa_i(t_i)$ preserves the beliefs in t_i; that is, $\mathrm{marg}_{X_{n-1}} \kappa_i(t_i) = t_i^n$ 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 $\kappa_i : T_i \to \Delta(\Theta \times T_{-i})$, represent alternative (smaller) *type spaces*. Since I consider only symmetric type spaces in which there exists a set T such that $T_i = T$ for every i, I will frequently informally refer to T itself as the type space, with the understanding that it is meant to invoke $(T_i, \kappa_i)_{i\in\mathcal{I}}$.

Remark 2. Types are sometimes modeled as encompassing *all* uncertainty in the game. In this paper, I separate strategic uncertainty over opponent actions from structural uncertainty over payoffs.

Common p-belief. Let $T^* = T_1^* \times \cdots \times T_I^*$ denote the set of all type profiles, with typical element $t = (t_1, \dots, t_I)$. Then, $\Omega = \Theta \times T^*$ is the set of all "states of the world." Following Monderer & Samet (1989), for every $E \subseteq \Omega$, let

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

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

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

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, characterized in Battigalli & Sinischalchi (2003). For any $F \subseteq \Delta(\Theta)$, define

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

to be the event in which every player's first-order belief is in F. Then, $\mathcal{C}^1(E_F)$ is the event in which it is common certainty that every player has a first-order belief in F. The set of types t_i given which player i believes that F is common certainty is the projection of $\mathcal{C}^1(E_F)$ onto T_i^*.[7] Since this set is the same for all players, I will refer to the projection of $\mathcal{C}^1(E_F)$ onto T_1^* as "the set of types with common certainty in F."

[6] $\mathrm{marg}_{X_{n-2}} t^n = t^{n-1}$, so that $(t^1, t^2, \dots)$ is a consistent stochastic process.

[7] Notice that when beliefs are allowed to be wrong (as they are in this approach), individual

2.3 Solution concepts

Strict solution concepts will play an important role in this paper. For the complete information game with payoffs u^*, say that action profile a is a *strict Nash equilibrium* if

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

and a *δ-strict Nash equilibrium* if

$$u_i^*(a_i, a_{-i}) > u_i^*(a_i', a_{-i}) + \delta \quad \forall\, i \in \mathcal{I} \text{ and } a_i' \neq a_i.$$[8]

Analogously, say that an action a_i is *strictly rationalizable* for player i if there exists a family of sets $(R_j)_{j\in\mathcal{I}}$ with $a_i \in R_i$, such that for every player j and action $a_j \in R_j$, there is some distribution $\alpha_{-j} \in \Delta(R_{-j})$ satisfying

$$u_j^*(a_j, \alpha_{-j}) > u_j^*(a_j', \alpha_{-j}) \quad \forall\, a_j' \neq a_j. \tag{3}$$

Say that the action is *δ-strictly rationalizable* for player i if there exists a family of sets $(R_j)_{j\in\mathcal{I}}$ with $a_i \in R_i$, such that for every player j and action $a_j \in R_j$, there is some distribution $\alpha_{-j} \in \Delta(R_{-j})$ such that

$$u_j^*(a_j, \alpha_{-j}) > u_j^*(a_j', \alpha_{-j}) + \delta \quad \forall\, a_j' \neq a_j. \tag{4}$$

Additionally, I will use two solution concepts for incomplete information games. The first, *interim Bayesian Nash equilibrium*, is an incomplete information version of Nash equilibrium. Fix any type space $(T_i, \kappa_i)_{i\in\mathcal{I}}$. A strategy for player i is a measurable function $\sigma_i : T_i \to A_i$, and 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 σ with $\sigma_i(t_i) = a_i$ for every $t_i \in T_i$ is a Bayesian Nash equilibrium.

---

perception of common certainty is the relevant object of study. That is, player i can believe that a set of first-order beliefs is common certainty, even if no other player in fact has a first-order belief in this set. Conversely, even if every player indeed has a first-order belief in F, player i may believe that no other player has a first-order belief in this set.

[8]Replacing the strict inequality $>$ with a weak inequality $\geq$, this definition reverses the more familiar concept of ϵ-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 ϵ-equilibrium was introduced to formalize a notion of approximate Nash equilibria (violating the equilibrium conditions by no more than ϵ). I use δ-strict equilibrium to provide a cardinal measure for the *strictness* of a Nash equilibrium (satisfying the conditions with δ to spare).

The second solution concept, *interim correlated rationalizability*, is an incomplete information analogue of rationalizability defined in Dekel, Fudenberg & Morris (2007). For every player 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 $S_{-i}^{k-1}[t_{-i}] = \prod_{j\neq i} S_j^{k-1}[t_{-j}]$. We can interpret π 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 player i of type t_i, or (henceforth) simply *rationalizable*.

3 Approach

Let us enrich the standard description of a game, introduced above, with two new primitives. First, a *data-generating process*—formally, a sequence of n i.i.d. random variables distributed according to a distribution P over signal-space $\mathcal{Z}$. Players commonly observe the realization $\mathbf{z}_n = (z_1, \ldots, z_n)$, which I will refer to as a *dataset*, but may interpret it in different ways. I will write $\mathbf{z}$ when the number of observations is not important, and Z_n to mean the random sequence of n observations.

The second new primitive is a set $\mathcal{F}$ of functions

$$f : \bigcup_{n=1}^{\infty} \mathcal{Z}^n \to \Delta(\Theta)$$

that take every dataset $\mathbf{z}$ into a distribution over Θ, or a first-order belief. I will refer to each function f as a *learning rule*, since it describes a way of extrapolating from data into a belief, and $\mathcal{F}$ a set of learning rules. One interpretation of $\mathcal{F}$ is as a set of different but reasonable ways to interpret data. This interpretation is particularly apt if every learning rule returns the true value of θ^* with sufficient data. In the main part of the paper, I will in fact impose the following, stronger, property of uniform consistency:

Definition 1 (Uniform consistency.)**.** *The family of learning rules $\mathcal{F}$ is θ^*-*uniformly consistent *if*

$$\sup_{f\in\mathcal{F}} d(f(Z_n), \delta_{\theta^*}) \to 0 \ \textit{a.s.}$$

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

This condition says that beliefs induced by all learning rules in $\mathcal{F}$ weakly converge to a point mass on the true value of the parameter; moreover, this convergence is uniform across $\mathcal{F}$, so that the speed of learning is not "too different" across different learning rules. Notice that this is a joint assumption on the data-generating process P and the set of learning rules $\mathcal{F}$.

Remark 3. Every finite family of learning rules $\mathcal{F}$ where $d(f(Z_n), \delta_{\theta^*}) \to 0$ a.s. for every $f \in \mathcal{F}$ is θ^*-uniformly consistent.

For every dataset $\mathbf{z}$, define

$$\Delta_{\mathbf{z}} := \{f(\mathbf{z}) \, : \, f \in \mathcal{F}\} \subseteq \Delta(\Theta)$$

to be the set of beliefs that are induced by learning rules in $\mathcal{F}$. We can think of this as the set of "plausible" first-order beliefs given data $\mathbf{z}$.

The main question of interest is the following: Suppose an external analyst does not know the hierarchies of beliefs that players possess, but believes that given any realization of the data $\mathbf{z}$, players have common certainty in the set $\Delta_{\mathbf{z}}$. (That is, every player has a first-order belief in $\Delta_{\mathbf{z}}$, believes with probability 1 that every other player having a first-order belief in $\Delta_{\mathbf{z}}$, and so forth.) Let the set of types with this property by called $T_{\mathbf{z}}$. What can the analyst predict about players' strategic behaviors, knowing only that they have types in the set $T_{\mathbf{z}}$?[9,10]

3.1 Examples

To give a sense of the flexibility of this approach, let us first consider a few concrete examples of classes of learning rules that satisfy θ^*-uniform consistency.

Bayesian updating with uncommon priors. Players have different subjective beliefs on Θ but agree on the relationship between values of θ and data. Formally, let M be a finite set of prior distributions on Θ that have a grain of truth ($\mu(\theta^*) > 0$ for every $\mu \in M$). Define $\left\{xi^\theta \, : \, \theta \in \Theta\right\}$ to be a set of stochastic processes on $\mathcal{Z}^\infty$, where $\xi_\theta \neq \xi_{\theta'}$ for every $\theta \neq \theta'$ and ξ^{θ^*} corresponds to i.i.d. draws

[9] Notice that no explicit relationship is imposed between types and learning rules. For example, each of the following is a way of generating type profiles in $T_{\mathbf{z}} \times \cdots \times T_{\mathbf{z}}$:

- Every player i is identified with a learning rule $f \in \mathcal{F}$, and the sequence of learning rules $(f_i)_{i \in \mathcal{I}}$ is common knowledge.
- Every player i is identified with an learning rule $f_i \in \mathcal{F}$. Player i knows his own learning rule f_i, but has a nondegenerate belief distribution over the learning rules of other players.
- Every player i is identified with a distribution P_i on $\mathcal{F}$, and draws a learning rule at random from $\mathcal{F}$ from this distribution. The distributions $(P_i)_{i \in \mathcal{I}}$ are common knowledge.

[10] The results in this paper follow without modification if we relax this assumption to common certainty in the convex hull of distributions in $\Delta_{\mathbf{z}}$. See Lemma 4.

from P. The set of learning rules is described by $\mathcal{F} = (f_\mu)_{\mu \in M}$ where each f_μ takes observed data $\mathbf{z}$ into the posterior belief over Θ induced by prior μ and likelihood function $\{P^\theta : \theta \in \Theta\}$.

Confidence intervals. Players use least-squares regression to estimate a relationship between p covariates and a real-valued outcome variable. Data consists of tuples $(x_i, y_i) \in \mathcal{Z} := \mathbb{R}^p \times \mathbb{R}$, where $x_i \sim_{\text{i.i.d.}} \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$, and

$$y_i = x_i^T \beta + \epsilon_i, \qquad \epsilon_i \sim_{\text{i.i.d.}} \mathcal{N}(0, \sigma^2).$$

Assume that x_i and ϵ_i are independent. The first coordinate of the coefficient vector β, denoted β_1, is payoff-relevant, so $\Theta = \mathbb{R}$ and the true value of the parameter is $\theta^* = \beta_1$. The set of learning rules $\mathcal{F}$ consists of all maps from the data into a distribution with support in the 95% confidence interval for the least-squares regression estimate of β_1.

Case-based learning with different similarity functions. Players extrapolate from past instances by taking (different) weighted averages. Let $\mathcal{X} \subseteq \mathbb{R}$ be a set of attributes, which are related to outcomes in Θ under the unknown map $f : \mathcal{X} \to \Theta$. Data is a sequence of observations

$$\mathbf{z}_n = (x_1, f(x_1)), \ldots, (x_n, f(x_n)),$$

where every $x_k \sim_{\text{i.i.d.}} Q$. Suppose that the unknown parameter θ^* is the value of the function f evaluated at a new input x_0.

Fix a kernel function $K : \mathbb{R}^d \to \mathbb{R}$.[11] For every sequence $h = (h_n)_{n \geq 1}$ of constants, define $\hat{f}_{n,h} : \mathcal{X} \to \Theta$ to be the Nadaraya-Watson estimator

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

which produces estimates by taking a weighted average of nearby observations.

Let H be a set of sequences $(h_n)_{n \geq 1}$ that satisfy the assumptions presented in Einmahl & Mason (2005). Each sequence corresponds to a different level of "smoothing" applied to the data. Every learning rule in $\mathcal{F}$ is identified with a sequence from H, and takes the data $\mathbf{z}_n$ into a point belief on the Nadaraya-Watson estimate $\hat{f}_{n,h_n}(x_0)$.

[11] K is measurable and satisfies the conditions

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

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

3.2 Common Learning

Following Cripps et al. (2008), say that players *commonly learn* the true value of the parameter if they have asymptotic common certainty in the true value of the parameter. The formal definition, below, is adapted for the present setting, and says that for every level of confidence p and level of precision ϵ, every type in T_{Z_n} almost surely believes that the ϵ-ball around the true parameter θ^* is common p-belief.[12]

Definition 2 (Common Learning)**.** *Players commonly learn θ^* if*

$$\lim_{n\to\infty} P^n\left(\{\mathbf{z}_n : T_{\mathbf{z}_n} \subseteq \mathcal{C}^p(B_\epsilon(\theta^*))\}\right) = 1.$$

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

The following proposition says that the property of θ^*-uniform consistency is both necessary and sufficient for common learning.

Proposition 1. *Players commonly learn the true parameter θ^* if and only if $\mathcal{F}$ is θ^*-uniformly consistent.*

Common learning is a strong property; for example, it is not satisfied by the sequences of types considered in in Weinstein & Yildiz (2007), Carlsson & van Damme (1993), and Kajii & Morris (1997).[13] Why do we see it here? The key assumption is *common certainty* in the set $\Delta_{\mathbf{z}}$, which translates a restriction on first-order beliefs to a restriction on tail beliefs of arbitrarily high order. As the quantity of data increases, not only does the set of plausible first-order beliefs shrink (as a direct consequence of θ^*-uniform consistency), but in fact the set of every order k of beliefs shrinks *uniformly* in k. That is, it cannot be that players learn θ, believe with high probability that all other players learn θ, but fail to believe that ... other players believe that they learn θ.[14]

An immediate implication is that the $n \to \infty$ limit returns a complete information game, in which agents have common certainty in the true payoffs. Uniform consistency thus imposes an appealing discipline on the flexibility in beliefs generated via $\mathcal{F}$, and is for this reason assumed throughout the rest of the paper. Sections 5 and 6 show that despite this strong convergence in beliefs, strategic behavior given finite quantities of data may differ from their behavior in the limit game.

Remark 4. Section 8.1 describes the maximal relaxation of uniform consistency under which the main results continue to hold.

[12] I take $\epsilon > 0$, so that players believe it is approximate common certainty that the parameter is *close* to θ^*; in Cripps et al. (2008), Θ is finite, so players believe it is approximate common certainty that the parameter is exactly θ^*.

[13] The analogue of $n \to \infty$ is to take the size of the perturbation to 0.

[14] Formally, the set of types T_{Z_n} almost surely converges to the type with common certainty of θ^*, where convergence is in the Hausdorff metric induced by the uniform-weak metric on the universal type space. See Appendix A for the necessary definitions.

4 Asymptotic Behavior

Following, let us consider the predictions that the analyst can make regarding equilibria and rationalizable actions, if players commonly observe a dataset $\mathbf{z}_n$ of size n, and the analyst knows only that players have types in $T_{\mathbf{z}_n}$. A few preliminary definitions are needed.

For every pure-strategy Nash equilibrium a of the complete information game with payoffs u^*, define $p_n^{NE}(a)$ to be the probability (over possible datasets $\mathbf{z}_n$) that the strategy profile

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

is a Bayesian Nash equilibrium.[15] Analogously, define $p_n^R(i, a_i)$ to be the probability (over possible datasets $\mathbf{z}_n$) that action a_i is rationalizable for player i given any type in $T_{\mathbf{z}_n}$; that is,

$$a_i \in S_i^\infty[t_i] \quad \forall\, t_i \in T_{\mathbf{z}_n}.$$

Definition 3. *Say that the rationalizability of action a_i for player i is* robust to inference *if $p_n^R(i, a_i) \to 1$ as $n \to \infty$. Say that the equilibrium property of action profile a is* robust to inference *if $p_n^{NE}(a) \to 1$ as $n \to \infty$.*

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 player i, 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 plausible hierarchy is at least δ. In this way, robustness to inference is a minimal requirement for the rationalizability of a_i to persist when players infer their payoffs from finite data. Analogous statements apply when we replace rationalizability with equilibrium.

Let us first consider two examples in which robustness to inference is trivially met.

Example 1 (Trivial inference.)**.** *Let $\Theta = \mathbb{R}$ and define g to take every $\theta \in \Theta$ into the payoff matrix*

	a_1	a_2
a_1	θ, θ	$0, 0$
a_2	$0, 0$	$\frac{1}{2}, \frac{1}{2}$

[15] 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.

The true value of the parameter θ^* is strictly positive, so (a_1, a_1) is an equilibrium in the limit complete information game. Is it robust to inference?

Suppose $\mathcal{F} = \{f\}$, where the learning rule f is defined such that

$$f(\mathbf{z}) = \delta_{\theta^*} \quad \forall \, \mathbf{z};$$

that is, beliefs are a point mass on θ^* regardless of the realized data. Then, $\Delta_{\mathbf{z}} = \{\delta_{\theta^*}\}$ for every $\mathbf{z}$, and $T_{\mathbf{z}}$ consists only of the type with common certainty in the true value of the parameter. So the incomplete information game with type space $T_{\mathbf{z}}$ reduces to the limit complete information game, and it follows that $p_n^{NE}(a_1, a_1) = 1$ for every $n \geq 1$. The equilibrium property of (a_1, a_1) is trivially robust to inference.

Example 2 (Unnecessary inference.)**.** *Consider the same payoff matrix, but let $\Theta := [0, \infty)$. Then, action profile (a_1, a_1) is a Nash equilibrium in every complete information game in the image of g. It is easy to see that beliefs over Θ are irrelevant—that is,*

$$(\sigma_i)_{i \in \mathcal{I}}, \quad \textit{with } \sigma_i(t_i) = a_i \quad \forall \, i \in \mathcal{I}, \, t_i \in T$$

is a Bayesian Nash equilibrium given any set of types T. So again $p_n^{NE}(a_1, a_1) = 1$ for every $n \geq 1$, and the equilibrium property of (a_1, a_1) is trivially robust to inference.

The first example illustrates a case in which learning is artificial, in the sense that players know the true value of the parameter regardless of what data they see. The second example illustrates a case in which learning is not necessary, because the space of uncertainty has been chosen such that the unknown parameter is irrelevant to the player's strategic incentives. The following two conditions are designed to rule out these cases.

Assumption 1 (Nontrivial Inference.)**.** *There exists a constant $\gamma > 0$ such that for every n sufficiently large,*

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

This property says that for sufficient quantities of data, the probability that δ_{θ^*} is contained in the interior of the set of plausible first-order beliefs $\Delta_{\mathbf{z}_n}$ is bounded away from 0. It rules out Example 1 above, as well as related examples in which every learning rule in $\mathcal{F}$ overestimates, or every learning rule in $\mathcal{F}$ underestimates, the unknown parameter.[16]

[16]This does not rule out classes of *biased* estimators. It may be that in expectation, every learning rule in $\mathcal{F}$ overestimates the true parameter. Assumption 1 requires only that underestimation occurs with probability bounded away from 0.

To rule out the second example, I impose a richness condition on the image of g. For every player 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 player i; that is,

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

Assumption 2 (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.

4.1 Equilibrium

Let us first consider robustness of equilibrium.

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

This theorem suggests a new channel through which strictness produces robustness: inference of payoffs from finite data is subject to estimation error despite arbitrarily large quantities of data. Strict equilibria can withstand small misperceptions of payoffs, and are therefore robust to inference; in contrast, weak equilibria are sensitive to arbitrarily small perturbations in payoffs, and are not robust.

Proof. Define $U^{NE}_{a^*} \subseteq U$ to be the set of payoffs given which a^* is a Nash equilibrium.

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

Proof. Suppose that a^* is not a strict Nash equilibrium in the game with payoffs u. Then, there is some player i and action $a_i \neq a_i^*$ such that $u_i(a_i, a_{-i}^*) \geq u_i(a_i^*, a_{-i}^*)$. Define payoffs u^ϵ to agree with u everywhere, except that $u_i^\epsilon(a_i, a_{-i}^*) = u_i(a_i, a_{-i}^*) + \epsilon$. Then, action a_i is strictly better than a_i^* against a_{-i}^*, so a^* is not an equilibrium in the game with payoffs u^ϵ. Fix any sequence of positive constants $\epsilon_n \rightarrow 0$. Then for every n, $u^{\epsilon_n} \notin U^{NE}_{a^*}$, but $u^{\epsilon_n} \rightarrow u$ as $n \rightarrow \infty$. So $u \notin \text{Int}\left(U^{NE}_{a^*}\right)$, as desired.

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

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

so $u \in B_{\epsilon^*}(u) \subseteq U^{NE}_{a^*}$. Thus, $u \in \text{Int}\left(U^{NE}_{a^*}\right)$, as desired. □

Lemma 2. *The equilibrium property of action profile a^* is robust to inference if and only if $u^* \in \text{Int}\left(U^{NE}_{a^*}\right)$.*

Proof. Suppose $u^* \in \text{Int}\left(U^{NE}_{a^*}\right)$. It will be useful to define the function $h : \Delta(\Theta) \to U$ satisfying

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

which maps first-order beliefs into expected payoff functions. Fix any dataset $\mathbf{z}$ with the property that

$$h(\Delta_{\mathbf{z}}) := \{h(\mu)\ :\ \mu \in \Delta_{\mathbf{z}}\} \subseteq U^{NE}_{a^*}. \tag{5}$$

Since $\max_{i\in\mathcal{I}}(u_i(a_i^*, a_{-i}^*) - \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*)) \geq 0$ for every payoff $u \in U^{NE}_{a^*}$, also

$$\max_{i\in\mathcal{I}} \left(\int_U u_i(a_i^*, a_{-i}^*)\, \mathrm{d}\mu - \max_{a_i \neq a_i^*} \int_U u_i(a_i, a_{-i}^*)\, \mathrm{d}\mu \right) \geq 0$$

for every $\mu \in \Delta(U)$ satisfying $\mu\left(U^{NE}_{a^*}\right) = 1$. Thus, the strategy profile $(\sigma_i)_{i\in\mathcal{I}}$, where $\sigma_i(t_i) = a_i$ for every $i \in \mathcal{I}$ and $t_i \in T_{\mathbf{z}}$, is a Bayesian Nash equilibrium. Repeating this argument for every $\mathbf{z}$ satisfying (5), we see that the measure of such datasets of size n is a lower bound for $p_n^{NE}(a^*)$.

Under uniform consistency of $\mathcal{F}$ and continuity of h (see Lemma 3),

$$d(h(\Delta_{Z_n}), u^*) \to 0 \text{ a.s.}$$

Since u^* is in the interior of $U^{NE}_{a^*}$, this moreover implies that the measure of size-n datasets $\mathbf{z}$ for which $h(\Delta_{\mathbf{z}}) \subseteq U^{NE}_{a^*}$ converges to 1 as $n \to \infty$. Thus, $p_n^{NE}(a^*) \to 1$, and a^* is robust to inference.

In the other direction, suppose that $u^* \notin \text{Int}\left(U^{NE}_{a^*}\right)$. Since u^* is on the boundary of $U^{NE}_{a^*}$, there exists some player i and action $a_i \neq a_i^*$ such that $u_i^*(a_i, a_{-i}^*) \geq u_i^*(a_i^*, a_{-i}^*)$. Under the assumption of richness, there exists a parameter $\theta \in \Theta$ such that a_i is strictly dominant for player i in the game with payoffs $g(\theta)$. Under the assumption of nontrivial inference, for every n sufficiently large, there exists a constant $\epsilon_n > 0$ such that the measure of datasets $\mathbf{z}$ of size n satisfying $B_{\epsilon_n}(\delta_{\theta^*}) \subseteq \Delta_{\mathbf{z}}$ is at least γ. Define the sequence of belief distributions $(\mu_{\epsilon_n})_{n\geq 1}$ where

$$\mu_{\epsilon_n} = (1-\epsilon_n)\delta_{\theta^*} + \epsilon_n \delta_\theta \qquad \text{for every } n.$$

Since every $\mu_{\epsilon_n} \in B_{\epsilon_n}(\delta_{\theta^*})$, if $B_{\epsilon_n}(\delta_{\theta^*}) \subseteq \Delta_{\mathbf{z}}$ then also $\mu_{\epsilon_n} \in \Delta_{\mathbf{z}}$. Thus for n sufficiently large, the measure of size-n datasets $\mathbf{z}$ for which $\mu_{\epsilon_n} \in \Delta_{\mathbf{z}}$ is at least γ. Moreover, for every n,

$$\int_U u_i(a_i, a_{-i}^*)\, \mathrm{d}g_*(\mu_{\epsilon_n}) > \int_U u_i(a_i^*, a_{-i}^*)\, \mathrm{d}g_*(\mu_{\epsilon_n}).$$

So action a_i^* is not a best response to a_{-i}^* given any first-order belief μ_{ϵ_n}, and the strategy profile $(\sigma_i)_{i\in\mathcal{I}}$, where $\sigma_i(t_i) = a_i$ for every $i \in \mathcal{I}$ and $t_i \in T_{\mathbf{z}}$, is not a Bayesian Nash equilibrium. It follows that $\lim_{n\to\infty} p_n^{NE} < 1-\gamma$, so a^* is not robust to inference. □

It immediately follows from the two lemmas that the equilibrium property of action profile a^* is robust to inference if and only if a^* is a strict Nash equilibrium. $\square$

4.2 Rationalizability

Following Theorem 1, a reasonable conjecture is that rationalizability of action a_i^* for player i is robust to inference if and only if a_i^* is strictly rationalizable for player i. A simple example shows this to be false. Suppose that true payoffs are given by

	a_3	a_4
a_1	$1,0$	$1,0$
a_2	$0,0$	$0,0$

Notice that action a_1 is a strictly dominant strategy for player 1 in all close complete information game (where "close" is measured in Euclidean distance between payoffs). Since as the quantity of data increases, player 1 learns the true payoffs, it is easy to show that action a_1 is robust to inference.

But a_1 is not strictly rationalizable for player 1: in the first round of elimination, both of player 2's actions are eliminated, since neither is a strict best reply to any player 1 action. Since there are no surviving player 2 actions, action a_1 trivially fails to be a strict best reply to any surviving player 2 action, and is therefore also eliminated in the second round.

This example illustrates that the procedure of elimination matters. In particular, the definition for strict rationalizability requires that every action that is not a strict best response to any surviving opponent strategy is eliminated at once. This choice has unintuitive consequences in games like the example above, in which a player is indifferent between all of his actions. To resolve this, a new procedure of iterated elimination of strategies that are never a strict best reply is defined below, in which actions are eliminated (at most) one at a time.

First, define $W_i^1 := A_i$ for every player i. Then, for each $k \geq 2$, recursively remove (at most) one action in W_i^k that is not a strict best reply to any opponent strategy α_{-i} with support in W_{-i}^{k-1}. 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 define $\mathcal{W}_i^\infty$ to be the intersection of *all* sets W_i^∞ that can be constructed in this way. Say that an action a_i is *weakly strict-rationalizable* if $a_i \in \mathcal{W}_i^\infty$.[17]

[17] The choice of *weak* to describe the latter procedure is explained by Claim 1 (see Appendix B), which says that an action is strict-rationalizable only if it is weakly strict-rationalizable.

Returning to the example above, we see that there are two patterns of one-at-a-time elimination. One possibility is

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

in which action a_2 is eliminated for player 1 and action a_4 is eliminated for player 2, so that actions a_1 and a_3 remain. Another possibility is

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

in which action a_2 is eliminated for player 1 and action a_3 is eliminated for player 2, so that actions a_1 and a_4 remain. The action a_1 survives both procedures; hence, it is weakly strict-rationalizable.

Theorem 2. *Assume nontrivial inference and richness. Then, the rationalizability of action a_i^* for player i is robust to inference if a_i^* is strictly rationalizable, and only if a_i^* is weakly strict-rationalizable.*

It can be shown that under nontrivial inference and richness, robustness to inference (nearly) coincides with the concept of *robustly rationalizable* proposed in Morris, Takahashi & Tercieux (2012) (see Section 8.2 for a more detailed comparison), and can also be understood as requiring persistence across a subset of types that uniformly converge to the complete information type in the *uniform-weak topology* (proposed and characterized in Chen, di Tillio, Faingold & Xiong (2010)). In light of this, the sufficiency direction of this result has several predecessors—for example, its proof is very similar to the robust rationalizability of strict rationalizable actions in Morris, Takahashi & Tercieux (2012), and for finite $\mathcal{F}$ follows directly from lower hemi-continuity of strict rationalizability in the uniform-weak topology. The necessary direction above (and in particular, the construction of weakly strict rationalizable actions) is new, and its connection with these ideas above is explained following the remarks below.

Remark 5. Why do non-strict rationalizable actions fail to be robust, despite 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 in the proposed approach, 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 intuition for why refinement is obtained is that only perturbations in the uniform-weak topology (Chen et al. 2010) are considered in this paper, while Weinstein & Yildiz (2007) considers perturbations in what is known as the product topology (see Section 8.2 for an extended discussion). In particular, the sequences of types used to show failure of robustness in Weinstein & Yildiz (2007) do not converge in the uniform-weak topology, and hence do not have analogues in the proposed approach.

The example prior to the theorem provided intuition for why strict rationalizability is not necessary for robustness to inference—specifically, that the procedure rules out actions that we might think of as "strictly rationalizable" in an informal sense. Weak strict rationalizability turns out to have a second interpretation, which explains its role as a necessary condition for robustness to inference. Define $U^R_{a_i^*}$ to be the set of payoffs u such that a_i^* is rationalizable for player i in the complete information game with payoffs u. Then, the interior of this set is exactly characterized by the set of payoffs u given which a_i^* is weakly strict-rationalizable. One can show that rationalizability of action a_i^* is robust to inference only if the true payoffs u^* lie in the interior of $U^R_{a_i^*}$; thus, weak strict-rationalizability is necessary for robustness.

But it turns out that weak strict rationalizability is not sufficient, for the surprising reason that an action that is rationalizable given common certainty of the true payoffs may not be rationalizable given common certainty of an arbitrarily small neighborhood of the true payoffs. Such an example is constructed in Appendix D. This example relies on the fact that the chain of best responses rationalizing action a_i^* can vary across $U^R_{a_i^*}$. 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. If players believe that payoff functions on different sides of the boundary are common certainty, then action a_i^* may fail to be rationalizable.

On the other hand, if a_i^* is strictly 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^*, so that robustness is obtained.

Remark 6. A further implication of these comments is that although strict rationalizability is lower-hemicontinuous in the uniform-weak topology (Dekel, Fudenberg & Morris 2006, Chen et al. 2010), the slight relaxation to weak strict-rationalizibility is not.

Taken together, Theorems 1 and 2 demonstrate that strict solutions hold with probability arbitrarily close to 1 as the quantity of data increases to infinity.

5 Finite Data Behavior

Let us turn now to the question of what the analyst can say about behavior if players observe small, finite quantities of data. Do strict solutions continue to hold with high probability, and what determines this? Recall that Theorems 1 and 2 provided conditions under which $p_n^{NE}(a) \to 1$ and $p_n^R(i, a_i) \to 1$. Proposition 2 below complements these results with a lower bound on $p_n^{NE}(a)$ and $p_n^R(i, a_i)$ for every quantity of data n. This bound depends on two key features:

First, it depends on a cardinal measure of strictness. For every strict Nash equilibrium a, define

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

to be the largest value of δ for which a is a δ-strict NE. This parameter describes the amount of slack in the equilibrium property of action profile a; the larger δ_a^{NE} is, the "more strict" we will say that the equilibrium is. Similarly, for every strictly rationalizable action a_i, define

$$\delta_{a_i}^R = \sup\{\delta \, : \, a_i \text{ is } \delta\text{-strictly rationalizable}\}.^{18}$$

to be the largest value of δ such that a_i is δ-strictly rationalizable for player i. Again, this parameter describes the amount of slack in the solution; the larger $\delta_{a_i}^R$ is, the more strict we will say that rationalizability of a_i is.

The second important feature of the bounds below is the speed at which rules in $\mathcal{F}$ jointly learn the true payoffs. To define this, let us first introduce the function

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

that maps first-order beliefs into expected payoff functions. For every realized dataset $\mathbf{z}$, the quantity $\|h(f(\mathbf{z})) - u^*\|_\infty$ is the (sup-norm) distance between the true payoffs and the expected payoffs under belief $f(\mathbf{z})$. We can interpret this as the error in inferred payoffs. Taking the supremum over errors across all $f \in \mathcal{F}$, the quantity

$$\sup_{f \in \mathcal{F}} \|h(f(\mathbf{z})) - u^*\|_\infty$$

is an upper bound on the error from any single learning rule in $\mathcal{F}$. In expectation this quantity converges to 0; that is,

$$\mathbb{E}_{P^n} \left(\sup_{f \in \mathcal{F}} \|h(f(Z_n)) - u^*\|_\infty \right) \to 0.$$

[18] I abuse notation here and write $\delta_{a_i}^R$ instead of δ_{i,a_i}^R. Again, this parameter is defined only if a_i is δ-strictly rationalizable for some $\delta \geq 0$.

This follows from θ^*-uniform consistency of $\mathcal{F}$ and continuity of h (see Lemma 3). But the speed at which this convergence occurs varies substantially over different $(P, \mathcal{F})$ pairs—this is what I will refer to as *the speed of joint learning.*

Proposition 2. *(a) Let a^* be any strict Nash equilibrium in the limit complete information game. Then, for every $n \geq 1$,*

$$p_n^{NE}(a^*) \geq 1 - \frac{2}{\delta_{a^*}^{NE}} \mathbb{E}_{P^n} \left(\sup_{f \in \mathcal{F}} \|h(f(Z_n)) - u^*\|_\infty \right) \tag{6}$$

(b) Let a_i^* be any strictly rationalizable action for player i in the limit complete information game. Then, for every $n \geq 1$,

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

These bounds are increasing in the strictness of the solution (via $\delta_{a^*}^{NE}$ and $\delta_{a_i^*}^R$) and in the speed of joint learning. Intuitively, the speed of joint learning determines how much heterogeneity and error in beliefs we should expect given n observations, and the strictness of the solution determines how much heterogeneity and error in beliefs the solution can withstand. Formally, the key step in the proof is to show that $p_n^{NE}(a^*)$ admits as a lower bound the probability (over size-n datasets) that players have common certainty in the $\delta_{a^*}^{NE}/2$-neighborhood of the true payoffs. (The parallel statement holds for rationalizability, replacing $p_n^{NE}(a^*)$ with $p_n^R(i, a_i^*)$ and $\delta_{a^*}^{NE}$ with $\delta_{a_i^*}^R$.)

The roles played by strictness, and the "complexity" of learning, are in my view the key advantages to modeling beliefs using the proposed framework. In particular, while the degree of strictness resembles classical criteria for equilibrium selection (such as Pareto-dominance and risk-dominance) in that it relies exclusively on the payoff matrix, the speed of joint learning depends on features external to payoffs—specifically, the kind of data that agents see, and the rules they use to learn from the data. Proposition 2 thus provides a way to leverage knowledge about the context of the game towards predicting the likelihood of (limit) equilibria and rationalizable actions. This property will be drawn upon repeatedly in Section 6 to model a third party's ability to influence behavior via manipulation of the complexity of information.

5.1 The Role of Opinion Diversity

It is useful to separate the determinants of the speed of joint learning into two channels: first, the speed at which individual learning rules recover the true payoffs; and second, the "opinion diversity," or the correlation between beliefs induced by learning rules across $\mathcal{F}$. This section isolates the second channel and characterizes its influence on $p_n^{NE}(a)$ and $p_n^R(i, a_i)$.

More precisely, recall that every learning rule f maps data into a first-order belief. Let $f(Z_n)$ be the random belief induced by learning rule f and the measure P^n. Suppose that the distribution of $f(Z_n)$ is fixed for every learning rule in $\mathcal{F}$, but the joint distribution of $(f(Z_n))_{f\in\mathcal{F}}$ is not. Proposition 7 answers: how much can the probabilities $p_n^{NE}(a)$ and $p_n^R(i, a_i)$ vary?

To preview the bounds in the proposition below, let us consider the game

$$\begin{array}{ccc} & a & b \\ a & \theta, \theta & 0, \frac{1}{2} \\ b & \frac{1}{2}, 0 & \frac{1}{2}, \frac{1}{2} \end{array}$$

where $\Theta = \{1, 0\}$, and suppose that the true value of θ is 1. We will be interested in robustness of the equilibrium property of action profile (a, a). Define μ_L to be a point mass on 0 and μ' to be a point mass on 1, and fix two sequences $(Q_1^n)_{n\geq 1}, (Q_2^n)_{n\geq 1}$ of distributions from $\Delta(\Delta(\Theta))$ so that

$$\begin{aligned} Q_1^n(\mu) &= \frac{1}{4n} \text{ and } Q_1^n(\mu') = \left(1 - \frac{1}{4n}\right) \quad \forall\, n \\ Q_2^n(\mu) &= \frac{3}{4n} \text{ and } Q_2^n(\mu') = \left(1 - \frac{3}{4n}\right) \quad \forall\, n \end{aligned}$$

Suppose that any set of learning rules $\mathcal{F} = \{f_1, f_2\}$ is permitted, where $f_1(Z_n) \sim Q_1^n$ and $f_2(Z_n) \sim Q_2^n$. What is the range of possible values for $p_n^{NE}(a, a)$?

Let us consider a few natural choices of $\mathcal{F}$. First, we can choose $\mathcal{F}$ so that $f_1(Z_n)$ and $f_2(Z_n)$ are perfectly correlated for every n. One can show that

$$p_n^{NE}(a, a) = P^n(\{\mathbf{z}_n \mid f_1(\mathbf{z}_n) = f_2(\mathbf{z}_n) = \mu_H\}),$$

so the analysis reduces to determining when both learning rules map to μ'.[19] Since $f_1(Z_n)$ and $f_2(Z_n)$ are perfectly correlated, every dataset that is mapped to μ_H by f_2 is also mapped to μ_H by f_1. Therefore,

$$p_n^{NE}(a, a) = P^n(\{\mathbf{z}_n \mid f_2(\mathbf{z}_n) = \mu_H\}) = 1 - \frac{3}{4n}. \tag{7}$$

[19]Suppose $f_i(\mathbf{z}_n) = \mu_L$ for some $i = 1, 2$. Then, the type with common certainty in μ_L is in $T_{\mathbf{z}_n}$. But a is strictly dominated for every player with this belief; therefore,

$$p_n^{NE}(a, a) \leq P^n(\{\mathbf{z}_n \mid f_1(\mathbf{z}_n) = f_2(\mathbf{z}_n) = \mu_H\})$$

Now suppose $f_1(\mathbf{z}_n) = f_2(\mathbf{z}_n) = \mu_H$. Then, $T_{\mathbf{z}_n}$ is a singleton, consisting only of the type with common certainty in μ_H. Action a is rationalizable for both players of this type. So also

$$p_n^{NE}(a, a) \geq P^n(\{\mathbf{z}_n \mid f_1(\mathbf{z}_n) = f_2(\mathbf{z}_n) = \mu_H\})$$

Therefore, $p_n^{NE}(a, a) = P^n(\{\mathbf{z}_n \mid f_1(\mathbf{z}_n) = f_2(\mathbf{z}_n) = \mu_H\})$ as desired.

We can reduce $p_n^{NE}(a,a)$ by instead choosing $\mathcal{F}$ so that $f_1(Z_n)$ and $f_2(Z_n)$ are independent. Then,

$$\begin{aligned} p_n^{NE}(a,a) &= P^n\left(\{\mathbf{z}_n \mid f_1(\mathbf{z}_n) = \mu_H\}\right) P^n\left(\{\mathbf{z}_n \,:\, f_2(\mathbf{z}_n) = \mu_H\}\right) \\ &= \left(1 - \frac{1}{4n}\right)\left(1 - \frac{3}{4n}\right) < 1 - \frac{3}{4n}. \end{aligned}$$

In fact, we further reduce $p_n^{NE}(a,a)$ by choosing $\mathcal{F}$ so that $f_1(Z_n)$ and $f_2(Z_n)$ are "anti-correlated." Then, the datasets mapped to μ_L by f_1 and f_2 are disjoint. In this case,

$$p_n^{NE}(A,A) = 1 - \frac{1}{4n} - \frac{3}{4n} = 1 - \frac{1}{n}. \tag{8}$$

It turns out that (7) is the largest possible value of $p_n^{NE}(A,A)$ (subject to the constraints on $\mathcal{F}$ described above) and (8) is the smallest possible value.

These observations are generalized as follows for arbitrary finite $\mathcal{F}$. First, recall the previous definitions of U_a^{NE} as the set of payoffs given which a is a Nash equilibrium, and $U_{a_i}^R$ as the set of payoffs given which action a_i is rationalizable for player i. For every distribution $Q \in \Delta(\Delta(\Theta))$ and every quantity of data n, let

$$p_Q^{NE}(a) := Q\left(\{\mu \,:\, h(\mu) \in U_a^{NE}\right), \tag{9}$$

be the probability that a is a Nash equilibrium if every player shares a first-order belief determined by the realization of Q. Define $p_Q^R(i, a_i)$ analogously, replacing U_a^{NE} in (9) with $U_{a_i}^R$.

Proposition 3. *Fix any $K < \infty$ sequences of distributions $(Q_1^n)_{n\geq 1}, \ldots (Q_K^n)_{n\geq 1}$ of distributions from $\Delta(\Delta(\Theta))$. Then, for every $\mathcal{F} = \{f_1, \ldots, f_k\}$ such that every $f_k(Z_n) \sim Q_k^n$, the probabilities $p_n^{NE}(a)$ and $p_n^R(i, a_i)$ satisfy*

$$\begin{aligned} p_n^{NE}(a) &\in \left[1 - \sum_{k=1}^{K} p_{Q_k^n}^{NE}(a), 1 - \min_{k\in\{1,\ldots,n\}} p_{Q_k^n}^{NE}(a)\right] \\ p_n^{R}(i,a_i) &\in \left[1 - \sum_{k=1}^{K} p_{Q_k^n}^{R}(i,a_i), 1 - \min_{k\in\{1,\ldots,n\}} p_{Q_k^n}^{R}(i,a_i).\right] \end{aligned}$$

The upper bounds correspond to the case in which different learning rules perform poorly on sets of data that are as overlapping as possible. That is, if evaluating rule f on data $\mathbf{z}$ produces a belief that is "far" from the degenerate belief on the true parameter θ^*, then it is likely that evaluating other learning rules in $\mathcal{F}$ on the same data would also produce inaccurate beliefs. The lower bound, when attainable, corresponds to the case in which different learning rules perform poorly on sets of data that are "as different" as possible. That is, if evaluating learning rule f on data $\mathbf{z}$ produces a belief that is far from the degenerate belief on the true parameter

θ^*, then it is likely that evaluating other learning rules would produce accurate beliefs. This can be formalized using the idea of co-monotonic and counter-monotonic random variables, and the proof follows from a straightforward application of the Frechet-Hoeffding bound.[20]

6 Application: Data Design

So far, we have taken the data-generating process P and the set of learning rules $\mathcal{F}$ to be exogenously determined. But in practice, both public data and the way in which individuals interpret it are often influenced by external actors—for example, the federal reserve board decides what data to release about various financial and macroeconomic indicators, Consumer Reports determines what data to release about consumer goods, US News decides what data to release about universities, and so forth; moreover, the ways in which people draw inferences from public data are often guided by education and experts. For these reasons, it is important to understand the strategic implications of data design.[21]

The examples below illustrate how an external agent might influence strategic behaviors within the proposed framework, either by controlling the data that players see or the way that players interpret it. These examples focus on an interesting special case of the proposed approach, in which the signal space can be written as $\mathcal{Z} = \mathcal{X} \times \Theta$, where $\mathcal{X}$ is a set of observable features, and Θ is the space of payoff-relevant outcomes. For example,

- the set $\mathcal{X}$ might describe physical characteristics of a laptop (weight, battery life, resolution), while Θ describes quality.
- the set $\mathcal{X}$ might describe various macroeconomic indices (interest rates, the consumer price index), while Θ is inflation next term.
- the set $\mathcal{X}$ describes features of a university (student-faculty ratio, ethnic diversity, graduation rate), while Θ is the value to attending the university.

Players observe a sequence

$$(\mathbf{x}_1, \theta_1), \ldots, (\mathbf{x}_n, \theta_n),$$

where each θ_i is drawn independently and distributed according to an unknown conditional distribution $P(\theta|\mathbf{x} = \mathbf{x}_i)$. The payoff-relevant unknown is the value

[20] Similar techniques are used in Ely (N.d.).

[21] A recent area of research has made significant progress on related questions by formulating the third party's problem as a choice between information structures. The discussion below will depart from this literature in one key way by imposing extrinsic meaning to signals (see Section 8.4 for a brief review and comparison).

of θ at a new out-of-sample feature vector $\mathbf{x}$. In problems like this, a standard approach to inference is to estimate the unknown outcome by inferring a model $\phi : \mathcal{X} \to \Theta$ from the data, and evaluating ϕ at $\mathbf{x}$. There is a vast literature regarding problems like this and approaches for inference of ϕ. The examples below focus on two canonical cases.

Example 1 sets $\Theta = \{1, 0\}$, so that the problem is one of classification: players want to learn which values in $\mathcal{X}$ map to $\theta = 1$ and which values map to $\theta = 0$. I introduce a third party that determines the dimensionality of $\mathcal{X}$, and show that accurate reporting of extraneous observables (an artificial increase in the dimensionality of $\mathcal{X}$) can reduce the probability of coordination. Example 2 considers a related setting in which outcomes are linearly related to a set of covariates. I show that by reporting extraneous covariates (inducing agents to "over-fit" the data), an external analyst can again reduce the probability of coordination. These examples illustrate how standard notions of statistical complexity can be used to model human perception of the "ambiguity" of data, with implications for their strategic behaviors.

6.1 Example 1: Classification

Two plaintiffs are approached by a lawyer to join their cases into a class action suit. Their payoffs are

	Join	Not Join
Join	θ, θ	$0, \frac{1}{2}$
Not Join	$\frac{1}{2}, 0$	$\frac{1}{2}, \frac{1}{2}$

so that not joining yields a certain payoff of $\frac{1}{2}$ and joining alone yields a certain payoff of 0. If both players join, then the suit is taken to court and players receive an unknown payoff of $\theta \in \{1, 0\}$ (interpret $\theta = 1$ to mean success and $\theta = 0$ to mean failure). Is the action 'Join' rationalizable?

For concreteness, let $\mathcal{X} = [-c, c]^{\overline{p}}$, so that every suit is described by $\overline{p}$ characteristics (amount sued for, etc.), each normalized to lie within the interval $[-c, c]$. Every observation $(\mathbf{x}, \theta)$ describes the characteristics and outcome of a past class action suit. Observations are drawn i.i.d. from a distribution P on $\mathcal{X} \times \Theta$ with the properties that: (1) $\text{marg}_{\mathcal{X}} P$ is uniform over $\mathcal{X}$; and (2) for every $\mathbf{x} \in \mathcal{X}$, the conditional distribution $P(\cdot \mid \mathbf{x})$ is a point mass on the value of

$$\phi^*(\mathbf{x}) = \begin{cases} 1 & \text{if } x^k \in [-c', c'] \text{ for every } k = 1, \dots p^* \\ 0 & \text{otherwise} \end{cases}$$

for some $p^* < \overline{p}$. Notice that only p^* of the $\overline{p}$ characteristics matter for the outcome of the suit.

An external agency chooses a transformation of the realized data

$$(\mathbf{x}_1, \theta_1), \ldots (\mathbf{x}_n, \theta_n) \tag{10}$$

that determines what players observe. Specifically the agency chooses the number of characteristics to report, where the agency is obligated to report each of the first p^* characteristics, but can in addition (truthfully) report any of the remaining $\overline{p} - p^*$ characteristics. In the following, I will take these features to be symmetric, so that the agency's choice is simply an integer $p \in \{p^*, p^* + 1, \ldots, \overline{p}\}$. Thus, instead of observing (10), players observe

$$\mathbf{z}_n = ((\tau_p(\mathbf{x}_1), \theta_1), \ldots, (\tau_p(\mathbf{x}_n), \theta_n)), \tag{11}$$

where $\tau_p : (x_1, \ldots, x_{\overline{p}}) \mapsto (x_1, \ldots, x_{p^*})$ is the truncation of vectors in $\mathcal{X}$ to their first p entries.

Players form beliefs about θ from the data in the following way. Take Φ to be the set of all "rectangular classification rules", of which ϕ^* is a member, defined to include every function $\phi : [-c, c]^p \to \Theta$ that can be written as

$$\phi(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in [c_1, c_1'] \times \cdots \times [c_p, c_p'] \\ 0 & \text{otherwise} \end{cases}$$

for some vector $(c_1, c_1', \ldots, c_p, c_p') \in [-c, c]^{2p}$. Given observation of (11),

$$\Phi_{\mathbf{z}} = \{\phi \in \Phi \, : \phi\left(\tau_p(\mathbf{x}_i)\right) = \theta_i \ \forall \ i = 1, \ldots, n\}.$$

is the set of all functions ϕ within the class of rectangular classification rules that exactly fit the observed data. See Figure 1 for an illustration.

Then, the set of all plausible predictions for the class action suit in question—which, let's say, has characteristics $\tilde{\mathbf{x}}$—is given by

$$\{\phi(\tau_p(\tilde{\mathbf{x}})), \phi \in \Phi_{\mathbf{z}}\} \subseteq \Theta,$$

since $\tau_p(\tilde{\mathbf{x}})$ is the reported set of characteristics, and $\phi(\tau_p(\tilde{\mathbf{x}}))$ is the prediction of function ϕ at $\tau_p(\tilde{\mathbf{x}})$. (Observe that since $\Theta = \{1, 0\}$, either this set is fully restrictive, in that every ϕ predicts the same value of θ, or it is completely unrestrictive, returning Θ.) Let the class of learning rules $\mathcal{F}$ be the set of all functions $f : \bigcup_{n=1}^{\infty} \to \Delta(\Theta)$ with the property that every belief $f(\mathbf{z})$ assigns probability 1 to the set of plausible predictions $\{\phi(\tau_p(\tilde{\mathbf{x}})), \phi \in \Phi_{\mathbf{z}}\}$.[22] Then:

[22]The following is a Bayesian interpretation of $\mathcal{F}$. The set of states of the world is $\Omega = \Theta \times \Phi \times \mathcal{Z}^\infty$, so that a state consists of a value of θ, a function ϕ, and an infinite sequence of observations from $\mathcal{Z}$. Let M be the set of all probability distributions over Φ, with the property that the induced

Proposition 4. *Suppose $\tilde{\mathbf{x}}^k \in (-c', c')$ for every $k = 1, \ldots, p^*$, which implies that Join is rationalizable for both players under complete information. Then, for every fixed quantity of data $n \geq 1$, and for $i \in \{1, 2\}$:*

(a) the probability $p_n^R(i, \text{Join})$ is monotonically decreasing in the number of reported characteristics p.

(b) $p_n^R(i, \text{Join}) \to 0$ as $p \to \infty$.

Thus, if the agency wants to minimize the probability that joining is rationalizable given n observations, it should report as many characteristics as possible ($p = \overline{p}$). Moreover, if we allow the agency to report arbitrarily many characteristics p, then $p_n^R(i, \text{Join})$ can be made arbitrarily small for any fixed number of observations. The essential feature of this example is that players do not know *which* or *how many* characteristics ϕ^* depends on. Thus, the more characteristics are reported, the greater the number of models that are "consistent" with the data, and as a result, the greater the ambiguity in how to interpret the data. Since rationalizability of the action Join requires not only that players assign sufficiently high probability to success of the class action suit ($\theta = 1$), but also that they believe with sufficiently high probability that the other player does the same, the dispersion in beliefs introduced by the extraneous variables serves to dissuade joining the class action suit. In practice, uncertainty caused by a lack of understanding or agreement over the determinants of an outcome seem realistic, and provision of "too much" information may indeed be a practical tool for preventing outcomes that require high confidence of similar views.

6.2 Example 2: Regression

Two investors decide whether to participate in a risky investment. Their payoffs are

	Invest	Don't Invest
Invest	θ, θ	$\theta - c, 0$
Don't Invest	$0, \theta - c$	$0, 0$

distribution over parameters in $[-c, c]^{2p}$ is absolutely continuous with respect to the Lebesgue measure.

Conditional on ϕ, a stochastic process ψ_ϕ generates an infinite sequence of i.i.d. draws from P_ϕ, where $\text{marg}_{\mathcal{X}} P_\phi$ is uniform over $\mathcal{X}$, and the conditional distribution $P(\cdot \mid \mathbf{x})$ is a point mass on $\phi(\mathbf{x})$. For every $\mu \in M$, write P_μ for the prior belief over Ω induced by μ and the signal processes $(\psi_\phi)_{\phi \in \Phi}$, with the further restriction that probability 1 is assigned to the set $\{(\theta, \phi, \mathbf{z}) : \theta = \phi(\tilde{\mathbf{x}})\}$.

Let $(\mathcal{H}_n)_{n=1}^\infty$ denote the filtration induced on Ω by datasets $\mathbf{z}_n$ of size n. Then, every $\mathbf{z}_n$ and prior belief $\mu \in M$ generate a posterior belief $P_\mu(\theta \mid \mathcal{H}_n)(\mathbf{z})$ over Θ. Write $\Delta_{\mathbf{z}} \subseteq \Delta(\Theta)$ to be the set of all such posteriors. The set of learning rules $\mathcal{F}$ consists of all maps $f : \bigcup_{n=1}^\infty \mathcal{Z}^n \to \Delta(\Theta)$ such that $f(\mathbf{z}) \in \Delta_{\mathbf{z}}$ for every $\mathbf{z}$. See Appendix C for the argument of equivalence.

where $c > 0$ is known but the return to joint investment $\theta \in \mathbb{R}$ is not. Is the action 'Invest' rationalizable?

Suppose that a central bank collects data on investments and their returns. Each investment is described by p covariates $(x^1, \ldots, x^p) \in \mathbb{R}^p$. Observations are pairs $(x_k^1, \ldots, x_k^p, \theta_k)$, where

$$\theta_k = \phi(x_k) + \epsilon_k = \beta_0 + \beta_1 x_k + \cdots + \beta_{p^*} x_k^{p^*} + \epsilon_k, \quad \epsilon_k \sim \mathcal{N}(0,1)$$

for $k = 1, \ldots, n$. That is, returns are a sum of a linear function of the first p^* covariates and a Gaussian disturbance term.

The central bank reports the first p' covariates describing each observed investment, where $p' \geq p^*$. Following this announcement, players form beliefs about θ by finding the best linear fit to the reported data and projecting the return at the covariates describing the project, which we can denote by $\mathbf{x}^* \in \mathbb{R}^{p'}$. Formally, let $\hat{\beta}^{LS} = \left(\hat{\beta}_0^{LS}, \ldots, \hat{\beta}_p^{LS}\right)$ be the solution to

$$\hat{\beta}^{LS} = \operatorname*{argmin}_{\beta \in \mathbb{R}^p} \sum_{k=1}^{n} |y_k - \beta \cdot (1 \ x_k^1 \ x_k^2 \ \ldots \ x_k^p)^T|^2.$$

The least-squares estimate of ϕ is then

$$\hat{\phi}_{LS}(x) = \hat{\beta}_0^{LS} + \hat{\beta}_1^{LS} x_k^1 + \ldots \hat{\beta}_p^{LS} x_k^p,$$

and the predicted return at $\mathbf{x}^*$ is $\hat{\phi}_{LS}(\mathbf{x}^*)$. Denote the $(1-\alpha)$-th confidence interval for the prediction $\hat{\phi}^{LS}(\mathbf{x}^*)$ by $CI(\mathbf{z})$.[23]

The set of learning rules $\mathcal{F}$ consists of all maps $f : \bigcup_{n=1}^{\infty} \mathcal{Z}^n \to \Delta(\Theta)$ with the property that for every $\mathbf{z}$, the belief $f(\mathbf{z})$ has support in the interval $CI(\mathbf{z})$.

Proposition 5. *Suppose $\phi(\mathbf{x}^*) > 0$. Then, for every fixed quantity of data $n \geq 1$, and for $i \in \{1, 2\}$,*

$$p_n^R(i, \text{Invest}) \geq 1 - \frac{1}{|\theta^*|}\phi(p')$$

for a function ϕ that is monotonically increasing in the number of reported characteristics p'.

Thus, if the bank wants to minimize the probability that 'Invest' is rationalizable given n observations, it should announce as many extraneous covariates as possible ($p = \overline{p}$). As in the previous example, the intuition is that rationalizability of the action 'Invest' requires common q-belief (for sufficiently high q) that the value of θ is positive. The greater the number of extraneous covariates reported, the larger

[23] The $(1-\alpha)$-th confidence interval has the property that, if it were repeatedly calculated on different samples, it would contain the true value $\phi(\mathbf{x}^*)$ for a $(1-\alpha)$ measure of samples

the variance of the least-squares prediction, and the larger the confidence interval around this prediction. This creates a larger range of "plausible" beliefs, given which even individuals who are optimistic about θ may nevertheless choose not to cooperate.

7 Extensions

The following section provides brief comment on various modeling choices made in the main framework.

7.1 Misspecification

The main results hold under a weakening of θ^*-uniform consistency, which I define below:

Definition 4 ((ϵ, θ^*)-uniform consistency.)**.** *For any $\epsilon \geq 0$, say that the class of learning rules $\mathcal{F}$ is* (ϵ, θ^*)-uniformly consistent *if*

$$\lim_{n\to\infty} \sup_{f\in\mathcal{F}} d(f(Z_n), \delta_{\theta^*}) \leq \epsilon \;\; a.s.$$

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

According to this definition, the class of learning rules $\mathcal{F}$ is (ϵ, θ^*)-uniformly consistent if the set of induced first order beliefs converges almost surely (in the Hausdorff distance induced by d) to an ϵ-neighborhood of the true parameter. Notice that θ^*-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 $\mathcal{F}$ is (ϵ, θ^*)-uniformly consistent, then it is also (ϵ', θ^*)-uniformly consistent for every $\epsilon' > \epsilon$.)

Proposition 6. *Assume nontrivial inference and richness.*

1. *Suppose $\mathcal{F}$ is $\left(\delta_{a^*}^{NE}, \theta^*\right)$-uniformly consistent. Then, the equilibrium property of a^* is robust to inference if and only if a^* is a strict equilibrium.*
2. *Suppose $\mathcal{F}$ is $\left(\delta_{a_i^*}^{NE}, \theta^*\right)$-uniformly consistent. Then, the rationalizability of action a_i^* is robust to inference if a_i^* is strictly rationalizable.*

Thus, the main results hold even if players have heterogeneous and incorrect beliefs even in the limit, so long as their limit beliefs are constrained within a $\delta_{a^*}^{NE}$ neighborhood (respectively, $\delta_{a_i^*}^{R}$-neighborhood) of the degenerate belief on θ^*.

7.2 Private Data

This paper studies players who observe a common dataset, but interpret it in different ways. How do the main results change if players instead observe private data? Cripps et al. (2008) have shown that if the set of signals $\mathcal{Z}$ is unrestricted, then common learning may not occur even if $\mathcal{F}$ consists of a single learning rule. So Proposition 1 need not hold. Moreover, Carlsson & van Damme (1993) and Kajii & Morris (1997) (among others) have shown that strict Nash equilibria are not robust to higher-order uncertainty about private opponent information. Thus, Theorems 1 and 2 also will not hold without additional restrictions on beliefs.

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 have (incorrect) degenerate beliefs that all opponents have seen the same data that they have. Then, Theorems 1 and 2 hold without change, and the bounds in Proposition 2 can be revised as follows.

Proposition 7. *Suppose a^* is a strict Nash equilibrium in the limit complete information game. Then, for every $n \geq 1$,*

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

where I is the number of players. Suppose a_i^* is strictly rationalizable in the limit complete information game. Then, for every $n \geq 1$,

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

7.3 Limit Uncertainty

In the main text, we assumed the existence of "true" payoffs, following which the limit as $n \to \infty$ corresponded to a complete information game. This approach can be extended in a simple way so that the limit game is in fact a game of incomplete information. Let $\nu \in \Delta(\Theta)$ be the "true" distribution over uncertainty (a "limit common prior") and rewrite the property of uniform consistency as follows:

Definition 5 (Limit Common Prior.)**.** *The set of learning rules $\mathcal{F}$ has a* limit common prior ν *if*

$$\sup_{f\in\mathcal{F}} d(f(Z_n), \nu) \to 0 \quad a.s.$$

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

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

8 Related Literature

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

8.1 Robustness of Equilibrium

Suppose an analyst does not know the exact game that players are playing, but believes it to be "nearby" to his model of the game. When can he be reasonably certain that the solutions in his model are close to solutions of the true game?

Early work on this question considered the true game to be a complete information game, and "nearby games" to mean other complete information games with payoffs close in the Euclidean norm (Selten 1975, Myerson 1978, Kohlberg & Mertens 1986). Fudenberg, Kreps & Levine (1988) suggested a concept of nearby *incomplete* information games, in which players have vanishing uncertainty about the true payoffs. This approach of embedding a complete information game into games with incomplete information has since been considered in several new ways. For example: Carlsson & van Damme (1993) study a class of incomplete information games in which beliefs are generated by (correlated) observations of a noisy signal of the true payoffs, and Kajii & Morris (1997) study incomplete information games in which beliefs are induced by information structures that place high ex-ante probability on the true payoffs.

The present paper continues in this tradition, studying solutions in incomplete information games that are "nearby" to the true complete information game. The definition of nearby that I use differs from the existing literature in several ways: First, I place a strong restriction on (interim) higher-order beliefs, which has the consequence that players commonly learn the true payoffs. This contrasts with the approaches of Carlsson & van Damme (1993) and Kajii & Morris (1997), in which—even as perturbations become vanishingly small—players consider it possible that other players 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, leading to rather different robustness results.[24]

[24]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.

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 players in the present paper learn about payoffs only, and not actions.

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. Chen et al. (2010) subsequently 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, under the assumptions of nontrivial inference, richness, and θ^*-uniform consistency, all plausible hierarchies converge in the uniform-weak topology to the singleton 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.[25]

Finally, Morris, Takahashi & Tercieux (2012) and Takahashi (N.d.) characterize which rationalizable actions remain rationalizable for all types with approximate common certainty in the true parameter. The property of robustness to inference considered in Section 4.2 can be understood as a continuous analogue of their concept—in the present setting, Θ may not be finite, and the prediction must hold across all types with approximate common certainty in *neighborhoods* of the true

[25] The robustness characterizations in this paper are conjectured to hold also for the relaxation in which players have common p-belief in the predictions of learning rules in $\mathcal{F}$, where $p \to 1$ as the quantity of data tends to infinity.

parameter. These differences do not have implications for the strategic characterization of asymptotic behavior, although they may matter for the rate results in Section 5. Relative to this, the contribution in Theorem 2 is an additional "only if" characterization.

8.3 Players 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 modeling 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 rules used by other players.

There is also a body of work that studies asymptotic disagreement between players who learn from data. Cripps et al. (2008) study players who use the same Bayesian model but observe different (private) sequences of data; Al-Najjar (2009) study players who use different frequentist models to learn from data; and Acemoglu, Chernozhukov & Yildiz (2015) study Bayesian players who have different priors over the signal-generating distribution. My model of belief formation shares many features with these models, 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 players 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 players share a common model and observe endogenous data (generated by past, strategic actors), while I suppose that players have different models and observe exogenous data. Additionally, the learning model in Steiner & 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 payoffs.

8.4 Informational Design

Section 7 is related to the informational design literature (Kamenica & Gentzkow 2011, Aumann & Maschler 1995, Brocas & Carrillo 2007, Rayo & Segal 2009, Alonso & Camara 2016, Ely N.d.), and in particular the recent literature regarding informational design in games (Bergemann & Morris 2016, Mathevet, Perego & Taneva 2016). These papers study the question of how a player can influence the behavior of another (or many others) by controlling their informational environment.

The examples discussed in Section 7, however, cannot be readily mapped into the general frameworks proposed in the papers above. Two key reasons include the following: First, the players in the present paper do not necessarily form beliefs using Bayesian inference (see e.g. Example 2), nor do they necessarily share a common prior (see e.g. Example 1); second, and more importantly, the signals in this paper have *exogenous* meanings, modeled through the set $\mathcal{F}$. This places a constraint on the messages that can be sent, and the posteriors they induce.

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 proposed and characterized in Battigalli & Sinischalchi (2003).

9 Conclusion

This paper proposes and characterizes a learning-based refinement of the universal type space. A set of "plausible" hierarchies of beliefs are defined from a common dataset and a set of rules for extrapolating from the data. The proposed approach is substantially more permissive than the common prior assumption, but restrictive enough still to make predictions. As the quantity of data converges to infinity, beliefs and behavior can be approximated by a limit complete information game. For small quantities of data, the appropriateness of such a reduction depends on the complexity of the problem of learning payoffs and the strictness of limit solutions.

Appendix A: Notation and Definitions

- If (X, d) is a metric space with $A \subseteq X$ and $x \in X$, I use the notation $d(A, x)$ to mean $\sup_{x' \in A} d(x', x)$.
- $\mathrm{Int}(A)$ 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 player i, and $u(a, \theta)$ to denote $g(\theta)(a)$.
- For any measures $\mu, \nu \in \Delta(\Theta)$, 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 μ and ν respectively.

- Let $T_i^k = \Delta(X_{k-1}) = \Delta(\Theta \times T_{-i}^{k-1})$ denote the set of possible k-th order beliefs.[26,27] The *uniform-weak topology* on T_i^*, proposed in Chen et al. (2010), is the metric topology generated by the distance

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

where d^0 is the metric defined on Θ (see Section 2.1)[28] and recursively for $k \geq 1$, d^k is the Prokhorov distance[29] on $\Delta\left(\Theta \times T_{-i}^{k-1}\right)$ induced by the metric $\max\{d^0, d^{k-1}\}$ on $\Theta \times T_{-i}^{k-1}$. Since I consider only symmetric type spaces, I will drop player subscripts throughout, referring to the uniform-weak metric d^{UW} on the set of types T.

[26] 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.

[27] Since type spaces for all agents are identical in this paper, I will consistently drop notation indexing type spaces to players.

[28] In Chen et al. (2010), Θ is finite and d^0 is the discrete metric, but this construction extends to all complete and separable (Θ, d^0).

[29] Recall that the *Levy-Prokhorov* distance ρ between measures on metric space (X, d) is defined

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

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

- Recall that strict rationalizability has the following equivalent definition: Set $\Sigma_i^1 = \Delta(A_i)$ for every player i. For every $k > 1$, recursively define

$$\Sigma_i^k = \{\sigma_i \in \Sigma_i^{k-1} \,|\, \exists\, \sigma_{-i} \in \prod_{k \neq i} \mathrm{Conv}(\Sigma_j^{k-1}) \text{ such that}$$
$$u_i(\sigma_i, \sigma_{-i}) > u_i(\sigma_i', \sigma_{-i}) \,\forall\, \sigma_i' \in \Sigma_i^{k-1}\}$$

Then, all actions in $R_i = \bigcap_{k=1}^{\infty} \Sigma_i^k$ are strictly rationalizable.

Appendix B: Preliminary Results

Lemma 3. *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 $d(\mu, \mu') \leq \epsilon$; then,

$$\begin{aligned} \|h(\mu) - h(\mu')\|_\infty = \left\| \int_\Theta g(\theta) d(\mu - \mu') \right\|_\infty &\leq K \sup_{g \in BL_1(\Theta)} \left\| \int_\Theta g(\theta) d(\mu - \mu') \right\|_\infty \\ &= K W_1(\mu, \mu') \\ &\leq K(\mathrm{diam}(\Theta) + 1) d(\mu, \mu') \\ &\leq K(\mathrm{diam}(\Theta) + 1)\epsilon \end{aligned}$$

using the assumption of Lipschitz continuity in the first inequality, and compactness of Θ and the Kantorovich-Rubinstein dual representation of W_1 in the following equality. The second inequality follows from Theorem 2 in Gibbs & Su (2002). It follows immediately that h is continuous. □

Lemma 4. *If* $d(\Delta_{Z_n}, \delta_{\theta^*}) \to 0$ *a.s. , then also*

$$d\left(\mathrm{Conv}\left(\Delta_{Z_n}\right), \delta_{\theta^*}\right) \to 0 \ \ a.s.$$

where $\mathrm{Conv}(\Delta_{Z_n})$ *denotes the convex hull of* Δ_{Z_n}.

Proof. Fix any dataset $\mathbf{z}_n$ of size n, any constant $\alpha \in [0, 1]$, and any pair of measures $\mu, \mu' \in \Delta_{\mathbf{z}_n}$. Again using the dual representation of the Wasserstein distance W_1,

$$\begin{aligned} W_1(\alpha\mu + (1-\alpha)\mu', \delta_{\theta^*}) &= \sup_{g \in BL_1(\Theta)} \left(\int g(\theta)\, \mathrm{d}((\alpha\mu + (1-\alpha)\mu') - \delta_{\theta^*}) \right) \\ &= \sup_{g \in BL_1(\Theta)} \left[\alpha \left(\int g(\theta)\, \mathrm{d}(\mu - \delta_{\theta^*}) \right) + (1-\alpha) \left(\int g(\theta)\, \mathrm{d}(\mu' - \delta_{\theta^*}) \right) \right] \\ &\leq \alpha \sup_{g \in BL_1(\Theta)} \left(\int g(\theta)\, \mathrm{d}(\mu - \delta_{\theta^*}) \right) \\ &\qquad + \ (1-\alpha) \sup_{g \in BL_1(\Theta)} \left(\int g(\theta)\, \mathrm{d}(\mu' - \delta_{\theta^*}) \right) \\ &= \alpha W_1(\mu, \delta_{\theta^*}) + (1-\alpha) W_1(\mu', \delta_{\theta^*}) \leq W_1(\Delta_{\mathbf{z}_n}, \delta_{\theta^*}) \end{aligned}$$

Using Theorem 2 in Gibbs & Su (2002),

$$d(\alpha\mu + (1-\alpha)\mu', \delta_{\theta^*})^2 \leq W_1(\alpha\mu + (1-\alpha)\mu', \delta_{\theta^*}), \tag{12}$$

and also

$$W_1(\Delta_{\mathbf{z}_n}, \delta_{\theta^*}) \leq K d(\Delta_{\mathbf{z}_n}, \delta_{\theta^*}). \tag{13}$$

where $K := (1 + \text{diam}(\Theta)) < \infty$ by compactness of Θ. Combining (12) and (13) with the previous inequality, it follows that for for every dataset $\mathbf{z}_n$,

$$d(\text{Conv}\left(\Delta_{\mathbf{z}_n}\right), \delta_{\theta^*})^2 \leq K d(\Delta_{\mathbf{z}_n}, \delta_{\theta^*}).$$

So $d(\Delta_{Z_n}, \delta_{\theta^*}) \to 0$ a.s. implies $d(\text{Conv}\left(\Delta_{Z_n}\right), \delta_{\theta^*}) \to 0$ a.s., as desired. □

Claim 1. *If action a_i is strictly rationalizable for player i in the complete information game with payoffs u^*, then it is also weakly strict rationalizable for player i in the complete information game with payoffs u^*.*

Proof. By induction. For arbitrary player j, let $(S_j^k)_{k \geq 1}$ be the sequence of sets of surviving actions corresponding to strict rationalizability (see Section A), and let $(W_j^k)_{k \geq 1}$ be the same for weak strict-rationalizability (see Section 5.2). Trivially $S_j^1 = W_j^1 = A_j$ for every player j. Suppose $a_j \notin W_j^2$; then, it is not a strict best response to any distribution over opponent actions, so also $a_j \notin S_j^2$. Thus,

$$S_j^2 \subseteq W_j^2 \quad \forall\, j.$$

Now, suppose $S_j^k \subseteq W_j^k$ for every player j, and consider any player i and action $a_i \in S_i^{k+1}$. By construction of the set S_i^{k+1}, there exists some $\sigma_{-i} \in \Sigma_{-i}^k$ such that

$$u_i(a_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})\, \forall\, \sigma_i \in \Sigma_i^{k-1}$$

But since $S_i^k \subseteq W_i^k$, the strategy σ_{-i} has support in W_{-k}^k, so a_i is a strict best response to some distribution π with support in the surviving set of weakly strict-rationalizable actions, implying that $a_i \in W_i^{k+1}$ as desired. □

Appendix C: Main Results

.1 Proofs in Section 4

.1.1 *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)[30], 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 $t^i_{\mathbf{z}}$ and t_{θ^*} is upper bounded by $d(\Delta_{\mathbf{z}_n}, \delta_{\theta^*})$.

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

Lemma 5. *Players commonly learn θ^* if and only if*

$$d^{UW}(T_{Z_n}, t_{\theta^*}) \to 0 \;\; a.s. \quad as \; n \to \infty.$$

Thus, the problem of determining whether players commonly learn θ is equivalent to that of determining whether the set of types T_{Z_n} almost surely converges to $\{t_{\theta^*}\}$ in the uniform-weak topology.

Lemma 6. *For every dataset $\mathbf{z}$,*

$$d^{UW}(T_{\mathbf{z}}, t_{\theta^*}) \leq d(\Delta_{\mathbf{z}}, \delta_{\theta^*}) \tag{14}$$

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

$$H^k_{\mathbf{z}} = \left\{ t^k \in T^k \,:\, \left(\mathrm{marg}_{T^{k-1}}\, t^k\right)(H^{k-1}_{\mathbf{z}}) = 1 \text{ and } \mathrm{marg}_{\Theta}\, t^k \in H^1_{\mathbf{z}} \right\}; \tag{15}$$

that is, $H^k_{\mathbf{z}}$ consists of the k-th order beliefs of types in $T_{\mathbf{z}}$. Define $\delta^* = d(\Delta_{\mathbf{z}}, \delta_{\theta^*})$. The following preliminary claims shows that every k-th order belief in the set $H^k_{\mathbf{z}}$ is within δ^* (in the d^k metric[31]) of the k-th order belief of type t_{θ^*}.

Claim 2. *For every $k \geq 1$,*

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

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 $\Delta_{\mathbf{z}}$. So it is immediate that

$$d^1(t, t_{\theta^*}) \leq d(\Delta_{\mathbf{z}}, \delta_{\theta}) = \delta^*. \tag{16}$$

[30]This lemma appears in Chen et al. (2010) for the case in which Θ is a finite set and d^0 is the discrete metric, but generalizes to any complete and separable metric space (Θ, d^0) when the definition of common learning is replaced by Definition 2.

[31]See Appendix A

Now suppose $H^k_{\mathbf{z}} \subseteq \left\{t^k_{\theta^*}\right\}^{\delta^*}$. Consider any measurable set $E \subseteq T^k$. If $t^k_{\theta^*} \in E$, then $t^{k+1}_{\theta^*} \geq \left(\left\{t^k_{\theta^*}\right\}\right) = 1$ by definition of t_{θ^*}. Moreover,

$$t^{k+1}\left(E^{\delta^*}\right) \geq t^{k+1}\left(\left\{t^k_{\theta^*}\right\}^{\delta^*}\right) \geq t^{k+1}\left(H^k_{\mathbf{z}}\right) = 1,$$

using (15) in the final equality and the assumption that $H^k_{\mathbf{z}} \subseteq \{t^k_{\theta^*}\}^{\delta^*}$ in the inequality preceding it. So

$$t^{k+1}_{\theta^*}(E) \leq t^{k+1}\left(E^{\delta^*}\right) + \delta^*. \tag{17}$$

If $t^k_{\theta^*} \notin E$, then $t^{k+1}_{\theta^*}(E) = 0$ (again by definition of t_{θ^*}), so (17) follows trivially. So $t^{k+1}_{\theta^*}(E) \leq t^{k+1}(E^{\delta^*}) + \delta^*$ for every measurable $E \subseteq T^k$. Using this and (16),

$$d^{k+1}(t, t_{\theta^*}) \leq \delta^*. \tag{18}$$

as desired. □

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

Thus, the question of convergence of types is reduced to the question of convergence in distributions over Θ.

Now, suppose that $\mathcal{F}$ is θ^*-uniformly consistent. Then, $d\left(\Delta_{Z_n}, \delta_{\theta^*}\right) \to 0$ a.s.[32] It follows from Lemma 6 that

$$d^{UW}(T_{Z_n}, t_{\theta^*}) \to 0 \text{ a.s.},$$

and from Lemma 5 that players commonly learn θ^*.

For the other direction, suppose $\mathcal{F}$ is not θ^*-uniformly consistent. In the following, I will use the notation $\mathbf{z}_{1:n}$ to mean the first n coordinates of an infinite sequence of observations $\mathbf{z} \in \mathcal{Z}^\infty$, and P^∞ to mean the (unique) measure induced on the product σ-algebra of $\mathcal{Z}^\infty$ by the sequence of measures $(P^n)_{n\geq 1}$ (see Kolmorogov's Extension Theorem). Failure of θ^*-uniform consistency is equivalent to:

$$P\infty\left(\mathbf{z} \in \mathcal{Z}^\infty \mid \lim_{n\to\infty} \sup_{f\in\mathcal{F}} d(f(\mathbf{z}_{1:n}), \delta_{\theta^*}) = 0\right) \neq 1. \tag{19}$$

Define $t^* : \mu \mapsto t^*_\mu$ to map first-order beliefs $\mu \in \Delta(\Theta)$ into the type t^*_μ with common certainty in μ. Then, (19) implies that

$$P\infty\left(\mathbf{z} \in \mathcal{Z}^\infty \mid \lim_{n\to\infty} \sup_{f\in\mathcal{F}} d^{UW}\left(t^*_{f(\mathbf{z}_{1:n})}, t^*_{\delta_{\theta^*}}\right) = 0\right) \neq 1.$$

from which it immediately follows that $d^{UW}(T_{Z_n}, t_{\theta^*})$ does not almost surely converge to 0. Using Lemma 5, players do not commonly learn θ^*, and we are done.

[32]Uniform convergence in W_1 implies uniform convergence in the Prokhorov metric d. See for example Gibbs & Su (2002).

.1.2 *Theorem 2*

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

Lemma 7. *$u \in \text{Int}\left(U^R_{a_i^*}\right)$ if and only if a_i^* is weakly strict-rationalizable in the complete information game with payoffs u.*

Proof. If: Suppose $u \notin \text{Int}\left(U^R_{a_i^*}\right)$. Consider any sequence of payoff functions $u^n \to u$. 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.

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 player j satisfying the recursive description in Section 5.1, such that $a_i^* \notin W_i^K$ for some $K < \infty$. To show that $u \notin \text{Int}\left(U^R_{a_i^*}\right)$, I construct a sequence of payoff functions u^n with $u^n \to u$ (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 player 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 3. *There exists $\gamma > 0$ such that for any u' satisfying $\|u' - u\|_\infty < \gamma$, and for any player 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. □

Corollary 1. *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 3, 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 player j and action $a_j \in A_j$. Suppose there exists some strategy $\alpha_{-j} \in \Delta(A_{-j})$ such that

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

so that a_j is a strict best response to α_{-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 1). Suppose a_j is never a strict best response, but there exists $\alpha_{-j} \in \Delta(A_{-j})$ such that

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

If $a_j \in W_j^1$, then

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

so also $a_i \in S_i^{1,n}$ for $n \geq N$. If $a_j \notin W^j$, then for $n \geq N$, there exists an action $a_j' \neq a_j$ such that $u_j\left(a_j', \alpha_{-j}\right) = u_j\left(a_j, \alpha_{-j}\right)$, but $u_i^n\left(a_j', \alpha_{-j}\right) > u_j^n\left(a_j, \alpha_{-j}\right)$. 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

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

then $a_j \in W_j^{k+1}$, and for $n \geq N$, also $a_j \in S_j^{k+1,n}$ (using Corollary 1). 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

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

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. $\square$

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

Proof. Suppose $u^* \in \text{Int}\left(U_{a_i^*}^R\right)$. Then, there is an open set V such that

$$u^* \in V \subseteq U_{a_i^*}^R.$$

Since h is continuous (see Lemma 3), $h^{-1}(V)$ is an open set in $\Delta(\Theta)$. So

$$\delta_{\theta^*} \in h^{-1}(V) \subseteq h^{-1}\left(U_{a_i^*}^R\right)$$

implying that $\delta_{\theta^*} \in \operatorname{Int}\left(h^{-1}\left(U^R_{a_i^*}\right)\right)$, as desired.

For the other direction, suppose towards contradiction that $\delta_{\theta^*} \in \operatorname{Int}\left(h^{-1}\left(U^R_{a_i^*}\right)\right)$ but $u^* \notin \operatorname{Int}\left(U^R_{a_i^*}\right)$. Since u^* is on the boundary of $U^R_{a_i^*}$, 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 2, $g(\Theta)$ 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 μ_ϵ satisfies

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

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 μ_ϵ over Θ, and therefore $h(\mu_\epsilon) \notin U^R_{a_i^*}$. This implies also $\mu_\epsilon \notin h^{-1}\left(U^R_{a_i^*}\right)$. Thus the sequence $\mu_\epsilon \to \delta_{\theta*}$ and has the property that $\mu_\epsilon \notin h^{-1}\left(U^R_{a_i^*}\right)$ for every ϵ, so $\delta_{\theta^*} \notin \operatorname{Int}(h^{-1}(U^R_{a_i^*}))$, as desired. $\square$

Suppose $u^* \notin \operatorname{Int}(U^R_{a_i^*})$; then, using Claim 4, also $\delta_{\theta^*} \notin \operatorname{Int}(h^{-1}(U^R_{a_i^*}))$. Under assumption NI, there is a constant $\epsilon > 0$ such that $\delta_{\theta^*} \in \operatorname{Int}(\Delta_{\mathbf{z}_n})$ for at least an ϵ-measure of datasets. Consider any such such dataset. Then, $\delta_{\theta^*} \notin \operatorname{Int}\left(h^{-1}(U^R_{a_i^*})\right)$, implies that $\Delta_{\mathbf{z}_n} \nsubseteq h^{-1}(U_a)$. Fix any $u \in \Delta_{\mathbf{z}_n} \backslash h^{-1}(U^R_{a_i^*})$. 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. $\square$

Sufficiency. If a_i^* is strongly strict-rationalizable, then there exists a family of sets $(V_j^k)_{j\in\mathcal{I}}$ is closed under δ-strict best reply for some $\delta \geq 0$; that is, for every $a_j \in V_j^k$, there exists a distribution $\alpha_{-j} \in \Delta(V^k_{-j})$ 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 9 (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 $\operatorname{marg}_{\Theta\times T_{-j}} \pi = g(t_j)$ and $\pi(a_{-j} \in V_{-j}[t_{-j}]) = 1$. Then, $V_j \subseteq S_j^\infty[t_j]$ for every player j.*

Fix any $\epsilon > 0$. Then, for every player 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 9 when ϵ is sufficiently small, and it follows that $a_i^* \in S_i^\infty[t_j]$, as desired.

.2 Proofs in Section 5

.2.1 *Proposition 2*

(a) To simplify notation, set $\delta := \delta_{a^*}^{NE}$. Since a^* is a strict equilibrium, $\delta \geq 0$.

Lemma 10. $B_{\delta/2}(u^*) \subseteq U_{a^*}^{NE}$.

Proof. Consider any payoff function u' satisfying

$$\|u' - u^*\|_\infty \leq \frac{\delta}{2}. \tag{20}$$

Then for every player i,

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

where $u_i^*(a_i^*, a_{-i}^*) - u_i^*(a_i', a_{-i}^*) > \delta$ follows from the assumption that a^* is a δ-strict NE in the complete information game with payoffs u^*, and the other two bounds follow from (20). So a^* is a NE in the complete information game with payoffs u', implying that $u' \in U_{a^*}^{NE}$. □

Recall from the proof of Theorem 1 that $h(\Delta_{\mathbf{z}}) \subseteq U_{a^*}^{NE}$ is a sufficient condition for the strategy profile

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

to be a Bayesian Nash equilibrium. So it follows from Lemma 10 that $h(\Delta_{\mathbf{z}_n}) \subseteq B_{\delta/2}(u^*)$ is also a sufficient condition. Thus,

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

using Markov's inequality in the final line.

(b) To simplify notation, set $\delta := \delta^R_{a_i^*}$. Since a_i^* is strictly rationalizable for player i, $\delta \geq 0$.

Lemma 11. $B_{\delta/2}(u^*) \subseteq U^R_{a_i^*}$.

Proof. Consider any payoff function u' satisfying

$$\|u' - u^*\|_\infty \leq \frac{\delta}{2}. \tag{21}$$

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

$$u_j^*(a_j, \alpha_{-j}[a_i]) > u_j^*(a_j', \alpha_{-j}[a_j]) + \delta \quad \forall\, a_j' \neq a_j. \tag{22}$$

I will show that $(R_j)_{j\in\mathcal{I}}$ satisfies the conditions in Lemma 9 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 player 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 $\mathrm{marg}_{\Theta\times T_{-j}}\, \pi = \kappa_j(t_j)$ and $\mathrm{marg}_{A_{-j}}\, \pi = \alpha_{-j}[a_j]$, 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 $h\left(\mathrm{marg}_\Theta\, \kappa(t_j)\right)$ 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

$$\inf_{u\in B_{\delta/2}(u^*)} \left(u(a_j, \alpha_{-j}) - u(a_j', \alpha_{-j})\right) \geq -\frac{\delta}{2} \tag{23}$$

It follows that

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

using the inequalities in (22) and (23). It follows that a_j is a best response to α_{-j} given distribution π. Repeating this argument for every player j, action $a_j \in R_j$, and type t_j with common certainty in $B_{\delta/2}(u^*)$, it follows from Lemma 9 that $R_j \subseteq S_j^\infty[t_j]$ for every player j. Since $a_i^* \in R_i$, also $a_i^* \in S_i^\infty[t_i]$, as desired. □

Recall from the proof of Theorem 2 that $\Delta_{\mathbf{z}_n} \subseteq U_{a_i^*}^R$ is a sufficient condition for

$$a_i \in S_i^\infty[t_i] \quad \forall\, t_i \in T_{\mathbf{z}_n}.$$

So it follows from Lemma 21 that $\Delta_{\mathbf{z}_n} \subseteq B_{\delta/2}(u^*)$ is also a sufficient condition. Thus,

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

using Markov's inequality in the final line.

.3 Proofs in Section 6

.3.1 *Proposition 4*

The argument is for player 1; the case for player 2 follows identically. For every dataset $\mathbf{z}_n = \{(\mathbf{x}_k, \theta_k\}_{k=1}^n$, define

$$\Delta_{\mathbf{z}_n} = \{\phi \in \Phi \,:\, \phi(\mathbf{x}_k) = \theta_k \quad \forall\, k = 1, \ldots n\}$$

and let $T_{\mathbf{z}_n}$ be the set of hierarchies of belief with common certainty in $\Delta_{\mathbf{z}_n}$. First, I will show that 'Join' is rationalizable for player 1 of any type in $T_{\mathbf{z}_n}$ if and only if $\Delta_{\mathbf{z}_n} = \{1\}$. Suppose that $\Delta_{\mathbf{z}_n} \neq \{1\}$, implying that $0 \in \Delta_{\mathbf{z}_n}$. Then $T_{\mathbf{z}_n}$ includes the type with common certainty in $\theta = 0$. Since Join is not rationalizable for player 1 of this type, it clearly does not hold that Join is rationalizable for player 1 of any type in $T_{\mathbf{z}_n}$. In the other direction, suppose $\Delta_{\mathbf{z}_n} = \{1\}$. Then, $T_{\mathbf{z}_n}$ is a singleton set consisting only of the type with common certainty in $\theta = 1$. Since Invest is rationalizable for player 1 of this type, it trivially follows that Invest is rationalizable for every type in $T_{\mathbf{z}_n}$.

Now, observe that $\Delta_{\mathbf{z}_n} = \{1\}$ if and only if every $\phi \in \Phi$ that satisfies $\phi(\mathbf{x}_k) = \theta_k$ for all $k = 1, \ldots, n$ predicts $\phi(\mathbf{x}^*) = 1$. That is, every rectangular classification rule that exactly fits the data classifies $\mathbf{x}^*$ to '1'. We can reduce this problem by looking at whether the smallest hyper-rectangle that contains every vector $\mathbf{x}$, where $(\mathbf{x}, 1) \in \mathbf{z}_n$, also contains $\mathbf{x}^*$. Say that Condition A is satisfied on dimension k if there exist observations $(\mathbf{x}^i, 1)$ and $(\mathbf{x}^j, 1)$ such that $x_k^i < \mathbf{x}_k^*$ and $x_k^j > \mathbf{x}_k^*$ (that is, a "successful" observation lies on either side of $\mathbf{x}^*$ in dimension k). Then, $p_n^R(i, \text{Join})$ is equal to the probability that Condition A holds on every dimension k.

If $k \in \{1, \ldots, p^*\}$, then Condition A is satisfied on dimension k only if some $\mathbf{x}^i$ satisfying $x_i^k \in [-c', \mathbf{x}_k^*)$ is sampled, and additionally some $\mathbf{x}^j$ satisfying $x_k^j \in (\mathbf{x}_k^*, c']$ is sampled. Since by assumption $x_k^* \in (-c, c')$, the probability that this occurs is

$$1 - \left[\left(\frac{2c - c' - \mathbf{x}_k^*}{2c} \right)^n + \left(\frac{2c - c' + \mathbf{x}_k^*}{2c} \right)^n - \left(\frac{c - c'}{2c} \right)^n \right] := q.$$

If $k \in \{p^* + 1, \ldots, \overline{p}\}$, then Condition A is satisfied on dimension k only if some $\mathbf{x}^i$ satisfying $x_i^k < \mathbf{x}_k^*$ is sampled, and additionally some $\mathbf{x}^j$ satisfying $x_k^j > \mathbf{x}_k^*$ is sampled. The probability that this occurs is

$$1 - \left(\frac{c - \mathbf{x}_k^*}{2c} \right)^n - \left(\frac{x_k^* + c}{2c} \right)^n := r.$$

Now, observe that realizations of characteristics are independent across dimensions. So the probability that Condition A is satisfied on every dimension is

$$p_n^R(i, \text{Join}) = q^{p^*} r^{p - p^*}.$$

Since $r < 1$, $p_n^R(i, \text{Join})$ is strictly and monotonically decreasing in p, as desired.

.3.2 *Footnote 22.*

Write $\mathbf{z} = (\mathbf{x}_k, \phi(\mathbf{x}_k))_{k=1}^n$. I will show that there exists a belief $\nu \in \Delta_{\mathbf{z}}$ such that Join is strictly dominated for player 1 with first-order belief ν if and only if there exists some $\phi \in \Phi$ such that

$$\phi(\mathbf{x}_k) = \theta_k \qquad \text{for every } k = 1, \ldots, n,$$

and moreover, $\phi(\mathbf{x}^*) = 0$.

Suppose there exists some $\phi \in \Phi$ satisfying the conditions above. Consider any prior belief μ with $\mu(\phi) > \frac{1}{2}$. Then, the posterior belief induced by prior μ and the data $\mathbf{z}$ assigns at least probability $\frac{1}{2}$ to ϕ, and hence at least probability $\frac{1}{2}$ to $\theta = 0$. Thus, there exists $\nu \in \Delta_{\mathbf{z}}$ with $\nu(\theta = 0) > \frac{1}{2}$, and Join is strictly dominated for either player with this first order belief.

In the other direction, suppose towards contradiction that there do not exist any functions $\phi \in \Phi$ satisfying the conditions above. Then, for any prior belief μ, the posterior given data $\mathbf{z}$ must put probability 1 on functions ϕ for which $\phi(\mathbf{x}^*) = 1$. Thus, there does not exist a belief $\nu \in \Delta_{\mathbf{z}}$ such that Join is strictly dominated for player 1 with first-order belief ν.

.3.3 *Proposition 5*

Fix an arbitrary $p' > p^*$ and write $\hat{\beta}$ for the OLS estimate of coefficients in

$$y = \beta_0 + \beta_1 x_1 + \cdots + \beta_{p'} x_{p'}$$

and $\tilde{\beta}$ for the OLS estimate of coefficients in

$$y = \beta_0 + \beta_1 x_1 + \cdots + \beta_{p'} x_{p'} + \beta_{p'+1} x_{p'+1}.$$

Claim 5. *For any vector* $\mathbf{u} = (\mathbf{w} \quad \mathbf{z})$ *where* $\mathbf{w} \in \mathbb{R}^{1\times p'}$ *and* $\mathbf{z} \in \mathbb{R}$,

$$Var\left(\mathbf{w}\hat{\beta}\right) \geq Var\left(\mathbf{u}\tilde{\beta}\right).^{33} \tag{24}$$

Proof. Write X for the $n \times p$ matrix stacking row vectors $(\mathbf{x}_1^i, \ldots, \mathbf{x}_{p'}^i)$, where $i \in \{1, \ldots, n\}$, and $\mathbf{x}_{p'+1}$ for the $n \times 1$ column vector of observations $\mathbf{x}_{p'+1}^i$. Write $U = (X \quad \mathbf{x}_{p'+1})$ for the concatenation of these two matrices. Finally, let $\mathbf{y}$ be the $n \times 1$ column vector of outcomes. Then,

$$\hat{\beta} = (X'X)^{-1}X'\mathbf{y} \quad \text{and} \quad \tilde{\beta} = (U'U)^{-1}U'\mathbf{y}.$$

Observe that

$$\begin{aligned}
\mathrm{Var}(\mathbf{w}\hat{\beta}) &= \mathrm{Var}(\mathbf{w}(X'X)^{-1}X'\mathbf{y}) \\
&= \mathrm{Var}(\mathbf{w}(X'X)^{-1}X'(X\beta + \epsilon)) \\
&= \mathrm{Var}(\mathbf{w}(X'X)^{-1}X'\epsilon)) \\
&= \sigma^2(\mathbf{w}(X'X)^{-1}X')(\mathbf{w}(X'X)^{-1}X')' \\
&= \sigma^2\mathbf{w}(X'X)^{-1}\mathbf{w}'
\end{aligned}$$

and similar manipulations yield that

$$\mathrm{Var}(\mathbf{u}\tilde{\beta}) = \sigma^2\mathbf{u}(U'U)^{-1}\mathbf{u}'.$$

Further define $R = (U'U)^{-1}$ and

$$Q = \begin{bmatrix} (X'X)^{-1} & O_{K_1\times K_2} \\ O_{K_2\times K_1} & O_{K_2\times K_2} \end{bmatrix}.$$

where each $O_{k\times k'}$ is a zero matrix of size $k \times k'$. The, the inequality in (24) holds if and only if the matrix

$$\Delta := R - Q = \begin{bmatrix} X'X & X'Z \\ Z'X & Z'Z \end{bmatrix}^{-1} - \begin{bmatrix} (X'X)^{-1} & O_{K_1\times K_2} \\ O_{K_2\times K_1} & O_{K_2\times K_2} \end{bmatrix}$$

[33] Proof of this claim benefitted from discussions with Ben Golub and Iosif Pinelis.

is positive semidefinite. To show this, write

$$\begin{pmatrix} U & V \\ V' & T \end{pmatrix} = \begin{pmatrix} X'X & X'Z \\ Z'X & Z'Z \end{pmatrix}^{-1}.$$

From properties of block matrix inversion,

$$\begin{aligned} V &= -A^{-1}BT \\ U &= A^{-1} + A^{-1}BTB'A^{-1} \\ T &= (D - B'A^{-1}B)^{-1} \end{aligned}$$

where $A := X'X$, $B := X'Z$, $D = Z'Z$.

Now consider any row vector $(\mathbf{c} \quad \mathbf{d})$. Algebraic manipulations yield

$$\begin{aligned} (\mathbf{c} \quad \mathbf{d}) \; \Delta \begin{pmatrix} \mathbf{c} \\ \mathbf{d} \end{pmatrix} &= (\mathbf{c} \quad \mathbf{d}) \begin{pmatrix} U - A & V \\ V' & T \end{pmatrix} \begin{pmatrix} \mathbf{c} \\ \mathbf{d} \end{pmatrix} \\ &= (B'A^{-1}\mathbf{x} - \mathbf{d})'T(B'A^{-1}\mathbf{x} - \mathbf{d}) \geq 0 \end{aligned}$$

using in the last inequality that T is positive definite (as a diagonal block in a positive definite matrix). Since this holds for arbitrary $(\mathbf{c} \quad \mathbf{d})$, we have that Δ is positive semidefinite as desired. □

.4 Remaining Proofs

Proposition 3

The argument below is for Nash equilibrium; the argument for rationalizability follows analogously. Enumerate the learning rules in $\mathcal{F}$ by $f_1, \ldots, f_K$. For every $k = 1, \ldots, K$, define

$$X_k^n = \mathbb{1}\left(h\left(f_k(Z_n)\right) \notin U_a^{NE}\right).$$

Write G_k^n for the distribution of random variable X_k^n, and G^n for the joint distribution of random variables $(X_k^n)_{k=1}^K$. By Sklar's theorem, there exists a copula $C : [0,1]^K \to [0,1]$ such that

$$G^n(x_1, \ldots, x_K) = C\left(G_1^n(x_1), \ldots, G_K^n(x_K)\right)$$

for every $(x_1, \ldots, x_K) \in \mathbb{R}^K$. Using the Frechet-Hoeffding bound,

$$1 - K + \sum_{k=1}^{K} G_k^n(x_k) \leq C\left(G_1^n(x_1), \ldots, G_K^n(x_K)\right) \leq \min_{k \in \{1,\ldots,K\}} G_k^n(x_k).$$

Since $p_n^{NE}(a) = G^n(0, \ldots, 0)$, it follows that

$$1 - K + \sum_{i=1}^{K} G_k^n(0) \leq p_n^{NE}(a) \leq \min_{k \in \{1,\ldots,K\}} G_k^n(0). \tag{25}$$

Finally, since every $X_k^n \sim \mathrm{Ber}\left(1 - p_{Q_k^n}^{NE}\right)$ by definition of $p_{Q_k^n}^{NE}$, (25) implies

$$1 - \sum_{k=1}^{K} p_{Q_k^n}^{NE} \leq p_n^{NE}(a) \leq 1 - \min_{k \in \{1,\ldots,K\}} p_{Q_k^n}^{NE}$$

as desired.

Appendix D

The following is an example of a game with the property that some action a is weakly strict-rationalizable given common certainty of the true payoffs, but not rationalizable along a sequence of perturbed types in the uniform-weak topology.

This game has four players, each of whom 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.) Payoffs u are defined as follows. Player 1's payoffs are:

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

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

Notice that if players 2-4 are playing the same action, then player 1's best response is a. Otherwise, his best response is b. The payoffs to players 2-4 are the same given either player 1 action, and are described below (where rows correspond to player 2's actions, columns to player 3, and choice of matrices to player 4). Player 1's payoffs are omitted, so that the first coordinate corresponds to player 2's payoff:

$$\begin{array}{ccc} & a & b \\ a & 1,1,0 & 0,0,0 \\ b & 0,0,0 & 0,0,0 \\ & (a) & \end{array} \qquad \begin{array}{ccc} & a & b \\ a & 0,0,0 & 0,0,0 \\ b & 0,0,0 & 1,1,0 \\ & (b) & \end{array} \tag{26}$$

Notice that 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 payoffs u are common certainty, so that this is a game of complete information. Action a is rationalizable for player 1 in this game; moreover,

- there is a constant $\epsilon > 0$ such that a is rationalizable for player 1 in every game u' with $\|u' - u\|_\infty \leq \epsilon$, so that rationalizability is preserved on a neighborhood of the complete information game with payoffs u.
- a is weakly strict-rationalizable.
- although a is not strictly rationalizable, it fails to survive this process for the reason that *none* of player 4's actions survive the first round of elimination.

Let t_1 be the type with common certainty in the payoffs u. I will now show that there exists a sequence of types t_1^n such that $t_1^n \to t_1$ in the uniform-weak topology, but action a fails to be rationalizable for agent 1 everywhere along this sequence.

Fix a sequence of positive constants ϵ_n that tend to 0 as $n \to \infty$. Let type t_1^n be any hierarchy of beliefs that satisfies the following three properties: First, the first-order beliefs of t_1^n are a point mass on u. Second, t_1^n assigns probability 1 to the event that player 2 believes it is common certainty that payoffs are:

$$\begin{array}{ccc} & a & b \\ a & 1,1,-\epsilon_n & 0,0,-\epsilon_n \\ b & 0,0,-\epsilon_n & 0,0,-\epsilon_n \\ & \multicolumn{2}{c}{(a)} \end{array} \qquad \begin{array}{ccc} & a & b \\ a & -\epsilon_n,0,0 & -\epsilon_n,0,0 \\ b & 0,0,0 & 1,1,0 \\ & \multicolumn{2}{c}{(b)} \end{array} \tag{27}$$

Third, t_1^n assigns probability 1 to the event that player 3 believes it is common certainty that payoffs are:

$$\begin{array}{ccc} & a & b \\ a & 1,1,0 & 0,-\epsilon,0 \\ b & 0,0,0 & 0,-\epsilon_n,0 \\ & \multicolumn{2}{c}{(a)} \end{array} \qquad \begin{array}{ccc} & a & b \\ a & 0,0,-\epsilon_n & 0,0,-\epsilon_n \\ b & 0,0,-\epsilon_n & 1,1,-\epsilon_n \\ & \multicolumn{2}{c}{(b)} \end{array} \tag{28}$$

Let us now consider the rationalizable actions for players 2 and 3, from the perspective of type t_1^n. If player 4 assigns probability 1 to the payoffs in (27), then b is his uniquely rationalizable action. So (t_1^n believes with probability 1 that) player 2 believes with probability 1 that player 4 will play b. Since (t_1^n believes with probability 1 that) player 2 also assigns probability 1 to the payoffs in (27), (t_1^n believes with probability 1 that) action b is player 2's uniquely rationalizable action. By a similar argument, if player 4 assigns probability 1 to payoffs in (28), then a is his uniquely rationalizable action. So player 3 believes with probability 1 that player 4 will play a, and thus considers a to be his own uniquely rationalizable action.

Thus, type t_1^n believes that player 2's uniquely rationalizable action is b, and player 3's uniquely rationalizable action is a, given which player 1's unique best response is b. Repeating this argument for every n, it follows that action a is not rationalizable for player 1 all along the sequence of types $(t_i^n)_{n\geq 1}$. But every t_i^n believes that $B_{\epsilon_n}(u)$ is common certainty, so $t_i^n \to t_i$ in the uniform-weak topology.

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* .

Alonso, Ricardo & Odilon Camara. 2016. "Bayesian Persuasion with Heterogeneous Priors." *Journal of Economic Theory* .

Aumann, Robert J. 1976. "Agreeing to Disagree." *The Annals of Statistics* .

Aumann, Robert & Michael Maschler. 1995. *Repeated Games with Incomplete Information*. MIT Press.

Battigalli, Pierpaolo & Marciano Sinischalchi. 2003. "Rationalization and Incomplete Information." *Advances in Theoretical Economics* .

Bergemann, Dirk & Stephen Morris. 2016. "Bayes correlated equilibrium and the comparison of information structures in games." *Theoretical Economicsi* .

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* .

Brocas, Isabelle & Juan D. Carrillo. 2007. "Influence through Ignorance." *The RAND Journal of Economics* .

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* .

Einmahl, Uwe & David M. Mason. 2005. "Uniform in Bandwith Consistency of Kernel-Type Function Estimators." *The Annals of Statistics* .

Ely, Jeffrey. N.d. "Beeps." Working Paper.

Esponda, Ignacio. 2013. "Rationalizable conjectural equilibrium: A framework for robust predictions." *Theoretical Economics* .

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, Andrew Postlewaite & David Schmeidler. 2008. "Probability and Uncertainty in Economic Modeling." *Journal of Economic Perspectives* .

Gilboa, Itzhak & David Schmeidler. 1995. "Case-Based Decision Theory." *The Quarterly Journal of Economics* .

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* .

Kajii, Atsushi & Stephen Morris. 1997. "The Robustness of Equilibria to Incomplete Information." *Econometrica* .

Kamenica, Emir & Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* .

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.

Mathevet, Laurent, Jacopo Perego & Ina Taneva. 2016. "Information Design: the Belief Hierarchy Approach.".

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 & 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* .

Rayo, Luis & Ilya Segal. 2009. "Optimal Information Disclosure." *Journal of Political Economy* .

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* .

Takahashi, Satoru. N.d. “Strategic Implications of Almost Common Certainty of Payoffs.”.

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* .