

Games of Incomplete Information Played by Statisticians

Annie Liang*

January 1, 2018

Abstract

This paper proposes a foundation for heterogeneous beliefs in games, in which disagreement arises not because players observe different information, but because they learn from common information in different ways. Players may be misspecified, and may moreover be misspecified about how others learn. The key assumption is that players nevertheless have some common understanding of how to interpret the data; formally, players have common certainty in the predictions of a *class* of learning rules. The common prior assumption is nested as the special case in which this class is a singleton. 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 number of observations depends on the degree of strictness of the solution and the “complexity” of inference from data.

1 Introduction

How are beliefs formed? And how do individuals come to form beliefs over the beliefs of others? Predictions of play in incomplete information games depend crucially on our answers to these questions. The classic approach gives players a common prior over states of the world, and assumes that they use Bayesian updating to form a posterior belief given new information.¹ But this approach imposes strong restrictions on the extent to which players can disagree: for example, it implies that beliefs that

*University of Pennsylvania. Email: anliang@upenn.edu. I am deeply grateful to Drew Fudenberg, Jerry Green, David Laibson, and Sendhil Mullainathan for their continual advice and support. Drew Fudenberg’s extensive guidance was essential in the development of this paper. Comments from Jetlir Duraj, Siddharth George, Ben Golub, Scott Kominers, Jonathan Libgober, Harry Pei, Andrei Shleifer, Dov Samet, Tomasz Strzalecki, Satoru Takahashi, Anton Tsoy, and Muhamet Yildiz improved this paper and are gratefully acknowledged.

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

are commonly known are identical (Aumann, 1976), and that repeated communication of beliefs will eventually lead to agreement (Geanakoplos and Polemarchakis, 1982). These implications are hard to believe when taken literally: they conflict with considerable empirical evidence that individuals publicly disagree, and also with our basic, day-to-day, experience that individuals interpret information in different ways.² Perhaps more importantly, there are conflicts between the predictions of the common prior theory and certain observed economic behaviors—for example, the well known result that the common prior assumption precludes speculative trade (Milgrom and Stokey, 1982).

The main contribution of this paper is to propose a simple framework that relaxes the common prior assumption in a structured way. The proposed approach takes a statistical view of belief formation: economic actors observe a sequence of data (generated via a stochastic process) and extrapolate from that data to predict payoff-relevant unknowns. For example, investors forecast future stock returns based on past returns. Theoretically, given infinite data about a stationary environment, agents can learn a complete theory of the relationship between observables and outcomes. But when data is “partial” (for example, outcomes are observed for settings that are similar but not identical to the current environment) and limited in quantity, there can be many plausible and competing viewpoints on what the data implies.

The proposed framework thus permits “ambiguity” in how to interpret the data, formalized as a *class* of rules for learning from data, which structure the potential disagreement. I define a *learning rule* to be any map from data into beliefs over payoffs. For example, the common prior assumption can be described by a learning rule that Bayesian updates from a common view of the relationship between data and payoff-relevant unknowns. In general, I allow for a set of learning rules (potentially misspecified). The premise of this approach is that even in settings where there is not a single accepted way to interpret data, there is often domain knowledge about a class of reasonable approaches. For example, in the forecasting example above, there may be basic competing theories for how the returns are generated—corresponding to different parametrized families of return processes; or, players may take different averages of the historical data, based on different assumptions about which time periods are most relevant.³

Thus, the standard approach is enriched with two new primitives: a *class of*

²In financial markets, players publicly disagree in their interpretations of earnings announcements (Kandel and Pearson, 1995), valuations of financial assets (Carlin et al., 2013), forecasts for inflation (Mankiw et al., 2004), forecasts for stock movements (Yu, 2011), and forecasts for mortgage loan prepayment speeds (Carlin et al., 2014). players publicly disagree also in matters of politics (Wiegel, 2009) and climate change (Marlon et al., 2013).

³See for example case-based learning (Gilboa and Schmeidler, 1995; Gilboa et al., 2008).

learning rules and a *data-generating process*. Given any data set, the class of learning rules determines a plausible range of disagreement about payoffs. A final assumption, *Common Inference*, structures the approach: for every realization of the data, players have common certainty in the beliefs induced by the class of learning rules.⁴ Thus, although players can disagree, and even have common knowledge of disagreement, the extent of their disagreement is constrained.

Section 5 characterizes the conditions on the class of learning rules that guarantee common learning (Cripps et al., 2008).⁵ It is necessary and sufficient that beliefs induced by different learning rules weakly converge to a limiting belief, and that this convergence occurs *uniformly* over the set of learning rules. Loosely, this requires that the set of learning rules is not too large. The case in which players commonly learn receives special emphasis in this paper, and I refer to the limiting belief as a *limiting common prior*.⁶

Sections 6 and 7 focus on the restrictions on strategic behavior that are imposed by a class of learning rules, a data-generating process, and Common Inference. The key object of study is the probability (with respect to the data generating process) that a strategic prediction will hold under Common Inference. Specifically, I consider the probability that the prediction holds uniformly over all types consistent with Common Inference (call these predictions *robust*), and also the probability that the prediction holds for some type consistent with Common Inference (call these predictions *plausible*). Informally, the larger the two probabilities are, the more confidence we should have in the corresponding prediction. In this way, the probabilities serve as different *continuous* metrics for the strength of a prediction. Both measures are indexed to the number of observations, so that the strength of a prediction varies depending on how many observations players have seen.

The case in which players commonly observe a large quantity of data is considered in Section 6. The main takeaway is that when there is a limiting common prior, then strategic predictions that hold strictly given infinite data can also be made (with high probability) when players observe a large finite quantity of data. Formally, under the assumption of a limiting common prior (and an additional technical condition), the probability that an action profile is a plausible equilibrium converges to 1 (as the number of observations grows large) if and only if it is an equilibrium under the limiting common prior; the probability that an action profile is a robust equilibrium converges to 1 if and only if it is a *strict* equilibrium under

⁴That is, all players assign probability 1 to this set of beliefs, believe with probability 1 that all other players assign probability 1 to this set of beliefs, and so forth.

⁵That is, players' own beliefs converge to this distribution, they believe with high probability that all other players' beliefs converge to this distribution, and so forth.

⁶One can interpret this limit as the point at which learning has removed all differences in beliefs that are not due to differences of information. From this point forward, players indeed share a common perception of the world.

the limiting common prior. Similar statements hold for rationalizability, although there are additional subtleties in the appropriate notion of “strictness”: I introduce a new definition for “weakly” strict-rationalizable actions, which is a necessary condition for rationalizable predictions to be robust with high probability when the number of observations is large (the standard concept of strict rationalizability is a sufficient condition).⁷ Actions are plausibly rationalizable with high probability when the number of observations is large if and only if they are rationalizable under the limiting common prior. Thus, the simplifying assumption of a common prior is largely without loss if we believe that players observe a large amount of data, and there is a limiting common prior.

But when the number of observations is limited (the more practically relevant setting), then play can be quite different from the limit game. In Section 7, I provide a lower bound for the probability that a strategic prediction is robust, when the number of observations is some arbitrary (and potentially small) n . This bound depends on two key properties: (a) First, it is increasing in the speed at which different learning rules jointly learn the true value of the parameter. Thus, the more “complex” the learning problem, the lower the probability. (b) Second, the bound is increasing in a cardinal measure of strictness of the solution that I define. 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 action⁸; 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 bounds 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.

Section 8 explores the case in which the data-generating process is not exogenously determined, but can be influenced by an external actor—for example, the federal reserve board decides what data to release about various financial and macroeconomic indicators. How might a designer be able to manipulate behavior either by choosing the nature of public information? I present two examples that build on the proposed framework, and demonstrate that provision of extraneous public information can deter coordination.

Section 9 examines modeling choices made in the main text and discusses the

⁷The difference between these two notions of strict rationalizability regards the order of elimination, and is of independent interest.

⁸This reverses the more familiar concept of ϵ -equilibrium, which requires that each player’s prescribed action is no more than ϵ worse than the best action.

extent to which these can be relaxed—specifically, by allowing for an “approximate” limiting common prior, and allowing for i.i.d. private data.

Section 10 surveys the related literature. This paper primarily builds on a literature that studies the robustness of equilibrium predictions to the specification of player beliefs (Rubinstein, 1989; Monderer and Samet, 1989; Carlsson and van Damme, 1993; Kajii and Morris, 1997; Weinstein and Yildiz, 2007), but considers questions of robustness motivated by a learning foundation. This learning-based approach shares features with Dekel et al. (2004) and Esponda (2013), with a key difference that players in the present paper learn about payoffs only and not actions. Another precedent is Steiner and Stewart (2008), which characterizes the limiting equilibria of a sequence of games in which players infer payoffs from related games. I consider a related setting, where the main question is what happens when players infer payoffs using *heterogeneous* rules. Additionally, the focus on large data limits in Section 6 is especially related to Cripps et al. (2008), which characterizes the beliefs of Bayesian learners as the quantity of data grows large. The primary object of study in this paper is strategic behavior instead of beliefs.

Finally, the depiction of agents as “statisticians” or “machine learners” relates to recent work including Gilboa and Schmeidler (2003), Gayer et al. (2007), Al-Najjar (2009), Al-Najjar and Pai (2014), Acemoglu et al. (2015), Spiegel (2016), and Olea et al. (2017). I examine the strategic implications when multiple agents form beliefs given common certainty of a diverse model class.

2 Example

Two investors decide whether to invest in a new product. There is an underlying demand process $D : \mathbb{R}_+ \rightarrow \mathbb{R}$ that is sampled at n different times $t_1 < \dots < t_n$. Players *commonly* observe the sequence

$$(D(t_1), \dots, D(t_n))$$

Payoffs are given by

	Invest	Don't Invest
Invest	y, y	$y - 1, 0$
Don't Invest	$0, y - 1$	$0, 0$

where $y = D(T)$ is the unknown demand at some large time $T > t_n$. In this game, rationalizability of investment depends on our model of players' beliefs about y , and also their beliefs about others' beliefs; how should we model these?

The standard approach gives players a common prior over the possible values of y , and assumes that conditional on the value of y , players know the stochastic process

that generates demand. This has the implication that, since all information is public, investors share the same beliefs given any number of observations. Moreover, given sufficiently many observations, players will commonly learn y , so we can replace analysis of the game above with analysis of a limiting complete information game (where the value of y is common knowledge). In particular, if the true value of y is strictly positive, then investment will be rationalizable given enough data.

The analyst in the narrative above is indifferent to the historical pattern of demand realizations, and also to the structure of the demand process itself. It is irrelevant, for example, whether the historical data resembles the left or right panel of the figure below, although intuitively we may expect lower commonality of beliefs given the latter sequence of observations.




Figure 1: Two example sequences of demand realizations.

Let's distinguish between these cases by adding structure to the setting in the following way. Suppose that the true process samples uniformly over times in the interval $[0, \underline{T}]$, where $\underline{T} < T$, and the observed demand at time t is a d -th degree polynomial plus Gaussian error. That is,

$$D(t) = \beta_0 + \beta_1 t + \dots + \beta_d t^d + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2)$$

where both the vector of coefficients, and also the order of the polynomial d are unknown. Players form forecasts for $D(T)$ by fitting a function \hat{D} to the observed data and using the estimated function to predict demand at time T . Specifically, for each $d \leq \bar{d}$, let \hat{D}_d be the best d -th order polynomial fit to the observed data.⁹ The key assumption is that players commonly know that the order of the polynomial is not more than some \bar{d} , so that they have common certainty in the set of beliefs with support in $\{\hat{D}_d(T)\}_{d=1}^{\bar{d}}$.

In such a setting, the actual model order d , the size of the model class (as determined by \bar{d}), and the number of common observations n are all crucial.

⁹Assume that players minimize least-squares error, although this choice is not critical.

In fact, if $\bar{d} = \infty$, then there is no number of observations given which we can predict investment based on these assumptions alone. This because for every sequence $D(t_0), \dots, D(t_n)$, there is a polynomial of order d' that perfectly interpolates between the observed sequence and some $D(T) < 0$. In this sense, the model class is too rich, and common learning fails (see Section 5).

Suppose $d < \bar{d} < \infty$, so that some players may use overspecified regression models.¹⁰ Then, players will indeed commonly learn the true value of x . As in the common prior case, given sufficiently many observations, analysis of this game can be reduced to analysis of the limiting complete information game. So long as $y > 0$ (so that investment is *strictly* rationalizable in the limiting complete information game), then investment will be rationalizable given a large number of observations (see Section 6). However, the number of observations necessary is increasing in both d and \bar{d} . This is because the rate at which players commonly learn is made slower both by the true underlying complexity of the model (the true order d), and also by the number of extraneous variables used to fit the data (the size $\bar{d} - d$) (Hastie et al., 2009).

Finally suppose that $\bar{d} < d < \infty$, so that all players have underspecified regression models.¹¹ In this case, players will generally not commonly learn the true value of y . Prediction of investment as rationalizable may nevertheless make sense: the key question in this case is whether there exists an N such that for every number of observations $n \geq N$, each under-specified model of order $d < \bar{d}$ predicts that $y > 0$. If so, then even without common learning, we have that investment will be rationalizable given a sufficiently large number of observations (again see Section 6).

Thus, there are many reasons not to predict investment even if $y > 0$ and all agents use models that will eventually recover y . In particular, the set of such models can be too large, or the rate at which they recover y can be too slow. These assessments depend on the underlying model class and data-generating process, as well as on the number of public observations. For example, in the right panel of the figure, the number of public datapoints is small relative to the complexity of the underlying model. In the extended framework, the analyst should have low confidence in predicting that agents commonly perceive investment to be rationalizable.

The following section develops a general framework which includes the setting above as a special case: players commonly observe data, and have beliefs that are consistent with a common model class.

¹⁰There are plausible regression models that contain at least one extraneous predictor.

¹¹All regression models miss one or more predictors.

3 Setup and Notation

3.1 Basic Game

Consider a game with a finite set of players \mathcal{I} and a finite set of actions A_i for each player i , where $A = \times_i A_i$ is the set of action profiles. The state space is a compact subset of the set of possible payoff matrices $\Theta \subset \mathbb{R}^{|A| \cdot |\mathcal{I}|}$, endowed with the Euclidean distance metric. I will refer to θ alternatively as the *state*, the *payoffs*, or the (*complete information*) *game*. Notice that somewhat unusually, the state θ is the payoff matrix, and not a parameter that the payoff function depends on. To evoke more familiar notation, I often write $u_i(a, \theta)$ for $\theta_i(a)$; that is, the payoffs that player i receives from action profile a when the payoff matrix is θ . Finally, for every belief $\mu \in \Delta(\Theta)$, the expected payoff matrix is $\mathbb{E}_\mu[\theta]$.

3.2 Description of Beliefs

Because the state space $(\Theta, \|\cdot\|_2)$ is complete and separable, we construct a full description of player uncertainty over Θ as follows (Brandenburger and Dekel, 1993). Suppose for a moment that there are two players, and recursively define

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

so that each X_k is the set of possible k -th order beliefs. Let $T_0 = \prod_{n=0}^{\infty} \Delta(X_n)$. An element $(t^1, t^2, \dots) \in T_0$ is a *hierarchy of beliefs* over Θ (describing the player's uncertainty over Θ , his uncertainty over his opponents' uncertainty over Θ , and so forth), and will be referred to simply as a *type*.

The above 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 and Zamir (1985) have shown that for every player i , there is a subset of types T_i^* (that satisfy the property of *coherency*¹²) and a function $\kappa_i^* : T_i^* \rightarrow \Delta(\Theta \times T_{-i}^*)$ such that $\kappa_i(t_i)$ preserves the beliefs in t_i ; that is, $\text{marg}_{X_{n-1}} \kappa_i(t_i) = t_i^n$ for every n . Notice that T_{-i}^* is used here to denote the set of profiles of opponent types.

The tuple $(T_i^*, \kappa_i^*)_{i \in \mathcal{I}}$ is known as the *universal type space*. Other tuples $(T_i, \kappa_i)_{i \in \mathcal{I}}$ with $T_i \subseteq T_i^*$ for every i , and $\kappa_i : T_i \rightarrow \Delta(\Theta \times T_{-i})$, represent alternative (smaller) *type spaces*. Since I consider only symmetric type spaces in which there exists a set

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

T such that $T_i = T$ for every i , the set T itself will be informally referred to as the type space, with the understanding that it is meant to suggest $(T_i, \kappa_i)_{i \in \mathcal{I}}$.¹³

Remark 1. 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.

3.3 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 and Samet \(1989\)](#), for every $E \subseteq \Omega$, let

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

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

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

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

I use in particular the concept of common certainty in a set of first-order beliefs, characterized in [Battigalli and Siniscalchi \(2003\)](#). For any $F \subseteq \Delta(\Theta)$, define

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

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^* .¹⁴ 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 .”

3.4 Solution Concepts

Fix a complete information game with payoffs θ . The action profile a is a *Nash equilibrium* in this game if for every player i ,

$$u_i(a_i, a_{-i}, \theta) \geq u_i(a'_i, a_{-i}, \theta) \quad \forall a'_i \in A_i,$$

¹³For more than two players, the statement that $T_i = T \forall i$ should be understood as saying that each T_i is equivalent to T under an appropriate permutation of player indices.

¹⁴Notice that when beliefs are allowed to be wrong (as they are here), individual *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.

and it is a *strict Nash equilibrium* if the inequality above is strict for every $a'_i \neq a_i$.

The family of sets of actions $(R_j)_{j \in \mathcal{I}}$, where every $R_j \subseteq A_j$, is *closed under best reply* if for every player j and action $a_j \in R_j$, there is some distribution $\alpha_{-j} \in \Delta(R_{-j})$ such that

$$u_j(a_j, \alpha_{-j}, \theta) \geq u_j(a'_j, \alpha_{-j}, \theta) \quad \forall a'_j \in A_j.$$

The family $(R_j)_{j \in \mathcal{I}}$ is closed under *strict best reply* if the inequality above holds strictly for every $a'_j \neq a_j$. An action a_i is *rationalizable* for player i if $a_i \in R_i$ for a family $(R_j)_{j \in \mathcal{I}}$ that is closed under best reply, and a_i is *strictly rationalizable* if $a_i \in R_i$ for some family $(R_j)_{j \in \mathcal{I}}$ that is closed under strict best reply.

Now, fix an incomplete information game with type space $(T_i, \kappa_i)_{i \in \mathcal{I}}$, so that a strategy for player i is a measurable function $\sigma_i : T_i \rightarrow A_i$. The strategy profile $(\sigma_1, \dots, \sigma_I)$ is a *Bayesian Nash equilibrium* if

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

so that every action a_i is a best reply to the strategy σ and the belief κ_i over player types.

I will use the following incomplete information notion of rationalizability: 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 $a_i \in BR_i \left(\operatorname{marg}_{\Theta \times A_{-i}} \pi \right)$ for some $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ satisfying (1) $\operatorname{marg}_{\Theta \times T_{-i}} \pi = \kappa_i(t_i)$ 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 , the actions in

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

are *interim correlated rationalizable* for player i of type t_i , or (henceforth) simply *rationalizable* (Dekel et al., 2007).

4 Approach

4.1 Restriction on Beliefs

The proposed approach derives a set of beliefs from two new primitives: a *data-generating process*, and a *set of rules* for how to extrapolate beliefs from realized data.

The data-generating process is formally a sequence of random variables $Z = (Z_t)_{t \geq 1}$. Players commonly observe the realizations of the first n random variables Z_1, \dots, Z_n , which I refer to as a *data set*, but may interpret it in different ways. I use the convention that $Z^n = (Z_1, \dots, Z_n)$ is the random sequence of the first n realizations, \mathbf{z}_n is a typical realization of Z^n , and \mathcal{Z}_n is the set of possible realizations of (Z_1, \dots, Z_n) . Subscripts indicating the number of observations are dropped when they are not important.

Under the common prior assumption, players learn from data by Bayesian updating a common prior over the states and observations. I generalize this idea to *learning rules*, where a learning rule is any map from data sets into first-order beliefs:

$$\mu : \bigcup_{n=1}^{\infty} \mathcal{Z}_n \rightarrow \Delta(\Theta).$$

Given data set \mathbf{z}_n , every learning rule μ produces a first-order belief $\mu(\mathbf{z}_n)$; naturally, these beliefs need not be the same. Throughout, I restrict consideration to learning rules that further satisfy the following two regularity conditions:

Assumption 1 (Convergence). $\mu(Z^n)$ almost surely converges (in the weak topology) to some limiting belief μ^∞ .

Assumption 2 (Richness). Let μ^∞ be the limiting distribution defined above. There exists a sequence of positive numbers ξ_n such that

$$\xi_n \cdot (\mu(Z^n) - \mu^\infty) \rightarrow^d \nu$$

where ν assigns strictly positive measure to every open set (in the weak topology) in a neighborhood of the zero measure.

Assumption 1 says that beliefs induced by learning rule μ almost surely converge to a limiting belief as the quantity of observations increases. Assumption 2 requires that this convergence occurs with positive probability “from any direction”. This is a weak technical condition, which loosely guarantees that the path of convergence is not too asymmetric around the limiting beliefs; in particular, players cannot over- or under-estimate any payoffs with probability 1.^{15,16} Both assumptions are satisfied by

¹⁵A more detailed intuition for Assumption 2 is as follows. Note that for each realization of data \mathbf{z}_n , the expression $\xi_n \cdot (\mu(\mathbf{z}_n) - \mu^\infty)$ is a (scaled, signed) measure over Θ , describing the differences between belief $\mu(\mathbf{z}_n)$ and the limiting belief μ^∞ . This difference is the zero measure when the two beliefs are the same. Assumption 2 requires that for n arbitrarily large, there is a neighborhood around the zero measure in which every open set receives strictly positive probability. This guarantees that not only does $\mu(\mathbf{z}^n)$ converge to μ^∞ a.s., but the path of its convergence cannot exclude any particular direction of convergence with probability 1.

¹⁶Notice that biased estimators are permitted as learning rules—for example, every learning rule may over-estimate payoffs in θ in expectation. Assumption 2 requires that this overestimation does not occur with probability 1.

essentially all learning rules and data-generating processes that come up in practice.

Typically it is assumed that players interpret data using the same learning rule μ , so that given identical information, players hold identical beliefs. The main departure in this paper is to allow for a set of learning rules \mathcal{M} . For example:

Bayesian Updating with Different Models

Players observe signal realizations from a set \mathcal{Z} . Different learning rules correspond to different models for the signal-generating distribution. Let $\mathcal{M} = \{\mu_i\}$, where each learning rule μ_i is identified with a prior distribution π_i over $\Theta \times \mathcal{Z}^\infty$. Given data \mathbf{z} , the belief induced by learning rule μ_i is the marginal of the posterior belief over Θ (updating from π_i and \mathbf{z}).

Case-Based Learning with Different Similarity Functions

Suppose that $\mathcal{X} \subseteq \mathbb{R}$ is a set of attributes relevant to payoffs (e.g. physical covariates of a patient seeking health insurance). Players commonly observe a sequence of attribute vectors and the associated payoffs:

$$\mathbf{z}_n = (\mathbf{x}_1, \theta_1), \dots, (\mathbf{x}_n, \theta_n).$$

Attribute vectors are drawn i.i.d. from a known distribution, and payoffs are determined according to an unknown mapping f ; that is, each $\theta_k = f(\mathbf{x}_k)$. The attribute vector \mathbf{x}^* describing the present game is known, but the payoffs are not.

Define \mathcal{M} to be a set of learning rules, where each learning rule corresponds to an approach for weighting past observations. Let $g_i : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$ be a similarity function on attributes, so that $g_i(\mathbf{x}, \mathbf{x}')$ describes the distance between attributes \mathbf{x} and \mathbf{x}' . The predicted payoff matrix at \mathbf{x}^* is the weighted average $\frac{1}{n} \sum_{k=1}^n \theta_k \left(e^{-\lambda g_i(\mathbf{x}_k, \mathbf{x}^*)} \right) / \left(\sum_{k'} e^{-\lambda g_i(\mathbf{x}_{k'}, \mathbf{x}^*)} \right)$. Each μ_i maps \mathbf{z}_n into a degenerate belief on the corresponding weighted average.

Fix an arbitrary set \mathcal{M} of learning rules satisfying Assumptions 1 and 2. Even if players observe a *common* dataset \mathbf{z} , learning rules $\mu \in \mathcal{M}$ may determine different, plausible, first-order beliefs $\mu(\mathbf{z})$. The set of expected payoffs induced by learning rules from \mathcal{M} is given by

$$\Theta(\mathbf{z}) := \{ \mathbb{E}_{\mu(\mathbf{z})}[\theta] : \mu \in \mathcal{M} \}.$$

Throughout, I will consider these the possible payoffs given \mathbf{z} . The main restriction imposed on beliefs is the following.

Assumption 3 (Common Inference). *Let \mathbf{z} be the realization of data. Then, all agents have common certainty in the set of payoff matrices $\Theta(\mathbf{z})$.*

That is, all players assign probability 1 to payoffs in $\Theta(\mathbf{z})$, believe with probability 1 that all other players assign probability 1 to $\Theta(\mathbf{z})$, and so forth.¹⁷ Notice that when \mathcal{M} consists of a single Bayesian learning rule, then Common Inference reduces to the common prior assumption.

I emphasize that Common Inference is a restriction on final beliefs over Θ , and does not directly build in any structural assumptions about use of learning rules. For example, we can construct beliefs consistent with Common Inference by supposing that each player i uses a different learning rule μ_i to form his first-order beliefs, and these learning rules $(\mu_i)_{i \in \mathcal{I}}$ are common knowledge. Alternatively, players may hold beliefs that are induced by a convex combination of learning rules in \mathcal{M} , and they may also have uncertainty over which learning rules are used by other players to form beliefs.

A key feature of Common Inference is that it permits types with common knowledge disagreement: players may know that (all know that...) they hold different first-order beliefs.¹⁸ Such types are not permitted under the common prior assumption, even when we allow for private and different information (Aumann, 1976). In this way, Common Inference represents a relaxation of the common prior assumption, where the permitted extent of disagreement is governed by a fixed set of belief updating rules.

The key question of this paper is the following: What are the restrictions on strategic behavior that are imposed by a set of learning rules \mathcal{M} , a data-generating process Z , and Common Inference?

4.2 Robust and Plausible Strategic Predictions

Below I propose concepts for the predictions that the analyst can make regarding equilibria and rationalizable actions, when players commonly observe a (random) dataset of size n and have beliefs obeying Common Inference.

It will be useful to write $T(\mathbf{z})$ for the set of types satisfying Common Inference when the realized data is \mathbf{z} . Say that a prediction is “robust” if it can be made *for all* player types from $T(\mathbf{z})$, and “plausible” if it can be made *for some* player types from $T(\mathbf{z})$.

¹⁷This is a Δ -rationalizable set in the sense of Battigalli and Siniscalchi (2003), where $\Delta = \{\mu(\mathbf{z}) : \mu \in \mathcal{M}\}$.

¹⁸For a simple example, suppose it is common knowledge that player 1 forms first-order beliefs using μ_1 , and player 2 forms first-order beliefs using a different learning rule μ_2 .

Robust Predictions. For every action profile \mathbf{a} , define $\underline{p}_n^{NE}(\mathbf{a})$ to be the probability (over possible datasets \mathbf{z}_n) that play of \mathbf{a} constitutes a Bayesian Nash equilibrium whenever players have types in $T(\mathbf{z}_n)$;¹⁹ that is,

$$\underline{p}_n^{NE}(\mathbf{a}) = \Pr(\{\mathbf{z}_n : \sigma \text{ with } \sigma_i(t_i) = a_i \forall i, t_i \in T(\mathbf{z}_n) \text{ is a BNE.}\}) \quad (3)$$

Then, $\underline{p}_n^{NE}(\mathbf{a})$ is a continuous measure of how likely it is that \mathbf{a} is an equilibrium when players observe n realizations. Informally, the higher $\underline{p}_n^{NE}(\mathbf{a})$ is, the more confidence an analyst should have in predicting that \mathbf{a} is an equilibrium.

Similarly, define $\underline{p}_n^R(i, a_i)$ 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,

$$\underline{p}_n^R(i, a_i) = \Pr(\{\mathbf{z}_n : a_i \in S_i^\infty[t_i] \forall t_i \in T(\mathbf{z}_n)\}). \quad (4)$$

Again, the higher $\underline{p}_n^R(\mathbf{a})$ is, the the more confidence an analyst should have in predicting that a_i is rationalizable.

Plausible Predictions. Now, in contrast to the approach taken above, consider a maximally permissive perspective. Define

$$\begin{aligned} \bar{p}_n^{NE}(a) = \Pr(\{\mathbf{z}_n : \exists \text{ belief-closed type space } (T_i, \kappa_i)_{i \in \mathcal{I}} \\ \text{s.t. each } T_i \subseteq T(\mathbf{z}_n), \text{ and } \sigma \text{ with } \sigma_i(t_i) = a_i \forall i, t_i \in T_i \text{ is a BNE.}\}) \end{aligned} \quad (5)$$

This is the probability (over possible datasets \mathbf{z}_n) that constant play of \mathbf{a} constitutes a Bayesian Nash equilibrium for *some* (belief-closed) type space, where all player types are from $T(\mathbf{z}_n)$. Similarly, define

$$\bar{p}_n^R(i, a_i) = \Pr(\{\mathbf{z}_n : a_i \in S_i^\infty[t_i] \text{ for some } t_i \in T(\mathbf{z}_n)\}). \quad (6)$$

This is the probability (over possible datasets \mathbf{z}_n) that action a_i is rationalizable for player i for some type in $T(\mathbf{z}_n)$.

In the special case in which agents have a common prior, these definitions have the following simple interpretation:

Example 1. (Common Prior.) Suppose that players share a common and correct prior over $\Theta \times \mathcal{Z}^\infty$. Write μ for the learning rule that maps any sequence of realizations $\mathbf{z}_n = (z_1, \dots, z_n)$ into the induced posterior belief under the common prior.

¹⁹The implicit type space is $(T_i, \kappa_i)_{i \in \mathcal{I}}$ where each $T_i = T(\mathbf{z}_n)$. Notice that the definition of $T(\mathbf{z}_n)$ implies that this type space is belief-closed.

Each realization \mathbf{z}_n determines an interim game, where players all have common certainty in the induced posterior. (Notice that the set of plausible types $T(\mathbf{z}_n)$ is a singleton for every \mathbf{z}_n .) The common prior determines a distribution over data sets \mathbf{z}_n , and hence a distribution over possible interim games. The probabilities $\bar{p}_n^{NE}(\mathbf{a}) = \underline{p}_n^{NE}(\mathbf{a})$ are equal, and they are the measure of size- n datasets \mathbf{z}_n (under the common prior) with the property that action profile \mathbf{a} is an equilibrium in the corresponding interim game. Similarly, $\bar{p}_n^R(i, a_i) = \underline{p}_n^R(i, a_i)$ is the measure of size- n datasets \mathbf{z}_n (under the common prior) with the property that action a_i is rationalizable for player i in the corresponding interim game. Note that the properties that $\bar{p}_n^{NE}(\mathbf{a}) = \underline{p}_n^{NE}(\mathbf{a})$ and $\bar{p}_n^R(i, a_i) = \underline{p}_n^R(i, a_i)$ are a consequence of \mathcal{M} being a singleton. In general, these probabilities will not be the same.

4.3 Interpretation of Model

I discuss below different interpretations of the framework described above. One perspective is that the analyst knows the set of frameworks used by agents to learn about a particular economic unknown, but there is randomness in the outcome of the data. He wants to predict which actions agents might take after observing the realized data.

Under a second interpretation, the data-generating process and rules for learning from data produce a “natural” measure over interim types. Recall that in the common prior setting (Example 1), the ex-ante distribution over interim types is determined by the common prior. When players have different priors, or do not use Bayesian updating to form beliefs, then this approach is not viable. Nevertheless, the approach outlined in Section 4.2 produces a way to quantify which interim types are more likely. The probabilities defined in Section 4.2 then evaluate the strength of a prediction by the probability that players have types given which the prediction holds.

There is, additionally, a natural measure-theoretic notion of genericity in this framework, where a set of interim types are “generic” if this set has probability 1 under the data-generating process. This notion of genericity contrasts with the topological notion of genericity taken in many recent works (Dekel et al., 2006; Weinstein and Yildiz, 2007; Ely and Peski, 2011), where typicality means denseness in a particular topology. One attractive feature of the proposed perspective is that the probabilities defined in (3-6) provide different *continuous* measures of typicality.

5 Common Learning

Following Cripps et al. (2008), say that players *commonly learn* a distribution P if they have asymptotic common certainty in P . Formally, for every probability $p \in$

$[0, 1)$ and level of precision $\epsilon > 0$, players commonly learn P if every ϵ -neighborhood of P (in the weak topology) is eventually common p -belief for all types in $T(Z^n)$. This definition is generalized from the one used in [Cripps et al. \(2008\)](#), where Θ was a finite set (equipped with the discrete measure) and P was a degenerate measure.

Definition 1 (Common Learning). *Players commonly learn the distribution P if*

$$\lim_{n \rightarrow \infty} \mathbb{P}(T(Z^n) \subseteq \mathcal{C}^p(\{P\}^\epsilon)) = 1 \quad \forall p \in [0, 1], \epsilon > 0$$

where $\{P\}^\epsilon$ is the ϵ -neighborhood of P (in the weak topology).

Clearly, for common learning to occur, each individual learning rule $\mu \in \mathcal{M}$ must eventually deliver a belief arbitrarily close to P ; that is, $\lim_{n \rightarrow \infty} d(\mu(Z^n), P) = 1$ for every $\mu \in \mathcal{M}$, where d is the Prokhorov metric on $\Delta(\Theta)$.²⁰ The stronger condition that

$$\sup_{\mu \in \mathcal{M}} d(\mu(Z^n), P) \rightarrow 0 \text{ a.s.} \quad (7)$$

requires not only that beliefs induced by learning rules in \mathcal{M} weakly converge to P , but that this convergence is uniform across \mathcal{M} . Notice that this is a joint assumption on the data-generating process and the set of learning rules.

The proposition below says that (7) is the only requirement for common learning: players commonly learn P if and only if beliefs induced by learning rules in \mathcal{M} uniformly weakly converge to P .

Proposition 1. *Players commonly learn the true distribution P if and only if (7) holds.*

When the condition in (7) is satisfied by some P , I refer to P as a limiting common prior.

Definition 2 (Limiting Common Prior.). *Say that $P \in \Delta(\Theta)$ is a limiting common prior if it satisfies the condition in (7).*

Common learning is a strong property; for example, it is not satisfied by the sequences of types considered in in [Weinstein and Yildiz \(2007\)](#), [Carlsson and van Damme \(1993\)](#), and [Kajji and Morris \(1997\)](#).²¹ The reason we see it here is because players have *common certainty* in the set $\Theta(\mathbf{z})$ at any \mathbf{z} , which translates a restriction on first-order beliefs into a restriction on tail beliefs. As the quantity of data increases, not only does the set of plausible first-order beliefs shrink (as a direct consequence of (7)), but the set of plausible beliefs of every order shrinks *uniformly* across orders. Formally, the set of types $T(Z^n)$ almost surely converges to the type with common certainty of P , where convergence is in the Hausdorff metric induced by the uniform-weak metric ([Chen et al., 2010](#)) on the universal type space.

²⁰Otherwise, the type that has common certainty in the belief induced by μ does not eventually have common p -belief in small neighborhoods of P .

²¹The analogue of $n \rightarrow \infty$ is to take the size of the perturbation to 0.

6 Strategic Behavior: Asymptotics

Let us begin by considering the limiting strategic predictions that can be made when players have commonly observed a large number of realizations. Say that predictions are “robust to inference” if they eventually hold in all plausible interim games with probability arbitrarily close to 1.

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

Thus, strategic predictions are robust to inference if the analyst believes that the prediction holds with high probability *given sufficient data*. Conversely, if the prediction is not robust to inference, then there exists a constant $\delta > 0$ such that for *any* finite quantity of data, the probability that the prediction fails for some types consistent with Common Inference is at least δ . In this way, robustness to inference is a minimal requirement for a prediction to not require assumption that players have beliefs coordinated by an infinite quantity of data.

We can additionally ask when it is the case that a prediction holds for *some* player types consistent with Common Inference, so long as the number of observations is sufficiently large. I define below the analogous concept of “plausibility under inference” to capture this.

Definition 4. *Say that the equilibrium property of action profile \mathbf{a} is plausible under inference if $\overline{p}_n^{NE}(\mathbf{a}) \rightarrow 1$ as $n \rightarrow \infty$. Say that the rationalizability of action a_i for player i is plausible under inference if $\overline{p}_n^R(i, a_i) \rightarrow 1$ as $n \rightarrow \infty$.*

Section 6.1 characterizes robustness and plausibility under inference for the solution concept of Nash equilibrium, and Section 6.1 provides characterizations for rationalizability. In both cases, I begin by considering general sets of learning rules, and then turn to the special case of a limiting common prior (Assumption 2).

6.1 Equilibrium

For a given prediction of equilibrium—specifically, that some action profile \mathbf{a} is a Bayesian Nash equilibrium—robustness to inference is completely characterized by whether the set of plausible payoffs is eventually contained in the following set:

Definition 5. *Let $\Theta_{\mathbf{a}}^{NE}$ be the set of all payoffs θ with the property that when θ is common knowledge, then action profile \mathbf{a} is a Nash equilibrium.*

The interior of $\Theta_{\mathbf{a}}^{NE}$ consists of all payoffs given which a is a *strict* Nash equilibrium, and its boundary consists of all payoffs given which a is a *weak* (and not strict) Nash equilibrium.

Lemma 1. *The equilibrium property of action profile \mathbf{a} is robust to inference if and only if*

$$\mathbb{P}(\Theta(Z^n) \subseteq \Theta_{\mathbf{a}}^{NE}) \rightarrow 1 \text{ as } n \rightarrow \infty \quad (8)$$

Notice that failure of asymptotic learning (beliefs converge to an “incorrect” distribution over payoffs), and also failure of asymptotic agreement (players disagree even given infinite data), can both be consistent with robustness to inference. What is necessary and sufficient is that players are eventually certain that the expected payoffs are in $\Theta_{\mathbf{a}}^{NE}$, know that they are all certain of this, and so forth.

To better understand the condition in (8), let us consider the special case of a limiting common prior (Assumption 2).

Proposition 2. *Suppose there is a limiting common prior P . Then, the equilibrium property of action profile \mathbf{a} is robust to inference if and only if \mathbf{a} is a strict Nash equilibrium in the incomplete information game with common prior P .*

Thus, prediction of \mathbf{a} is robust to inference if and only if \mathbf{a} is a strict NE in the limiting common prior game. This result recalls [Monderer and Samet \(1989\)](#), which showed that strict equilibria in a complete information game are robust to approximate common certainty of payoffs; here, I consider the related exercise of weakening common certainty in a belief P to common certainty in a (shrinking) neighborhood of a belief P . (Note that in contrast to [Monderer and Samet \(1989\)](#), players are permitted to assign probability 0 to the actual payoffs θ arbitrarily close to the limit.)

Intuitively, under the assumption of a limiting common prior, it is eventually approximate common certainty that the true game is nearby to the limiting common prior game. If action profile \mathbf{a} is a *strict* equilibrium in that limiting game, then that action profile continues to be an equilibrium given common certainty in the set of nearby payoffs. Conversely, if \mathbf{a} is only a weak equilibrium (or not an equilibrium at all), then no matter the number of observations, there is positive probability (bounded below) that \mathbf{a} is not an equilibrium. The technical condition introduced in Assumption 2 is necessary for this latter “only if” direction; for example, if players were permitted to have beliefs that consistently overweight the payoffs to \mathbf{a} , then \mathbf{a} would be robust to inference even if it were a weak equilibrium in the limiting game.

The next results characterize the weaker condition of plausibility under inference. First define:

Definition 6. *For each player i and action profile \mathbf{a} , let $B_i(a_i, a_{-i}) \subseteq \Theta$ be the set of all payoff matrices given which a_i is a best response to a_{-i} .*

The lemma below says that prediction of \mathbf{a} is plausible under inference so long as some payoff in each $B(a_i, a_{-i})$ is eventually plausible.

Lemma 2. *The equilibrium property of profile \mathbf{a} is plausible under inference if and only if*

$$\mathbb{P}(\Theta(Z^n) \cap B_i(a_i, a_{-i}) \neq \emptyset \quad \forall i) \rightarrow 1 \text{ as } n \rightarrow \infty \quad (9)$$

This condition is considerably weaker than the requirement that eventually $\Theta(Z^n) \cap \Theta_{\mathbf{a}}^{NE} \neq \emptyset$. That is, action profile \mathbf{a} can be a plausible equilibrium even if players agree that there it would not be played given complete information of the true payoffs. This is demonstrated in the simple example below:

Example 2. Consider the 2-player game

$$\begin{array}{cc} & l & r \\ u & x, -x & 0, 0 \\ d & 0, 0 & 1, 1 \end{array}$$

where $x \in \{-1, 1\}$. Players share a common prior that each value of x is equally likely, and observe signals from $S = \{s_L, s_H\}$, which are generated according to the following information structure:

$$\begin{array}{cc} & s_L & s_H \\ x = -1 & p & 1 - p \\ x = +1 & 1 - p & p \end{array}$$

Different learning rules μ_p correspond to Bayesian updating from different values of p in the information structure above. Set $\mathcal{M} = \{\mu_p\}_{p \in P}$ where $P = [\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]$ for some $\epsilon > 0$.

For any sequence of realizations from $\{s_L, s_H\}$, there is a learning rule in \mathcal{M} that assigns higher posterior belief to $x = -1$, and a learning rule that assigns higher posterior belief to $x = 1$. Thus, it is always consistent with Common Inference that player 1's expected value of x is positive, but player 2's expected value of x is negative, and these beliefs are common knowledge. So, action profile (u, l) is a plausible equilibrium for all sequences of signals. But in either possible state of the world ($x = -1$ and $x = 1$), the action profile (u, l) is not a Nash equilibrium in the complete information game.

Notice however that in Example 2, the different learning rules yield increasingly divergent beliefs as the quantity of data gets large. Assumption of a limiting common prior rules this out; in this case, it turns out that the only surviving predictions are Nash equilibria in the limiting game.

Proposition 3. *Suppose there is a limiting common prior P . Then, the equilibrium property of action profile \mathbf{a} is plausible under inference if and only if \mathbf{a} is a Nash equilibrium in the incomplete information game with common prior P .*

Thus, under Assumption 2, the difference between the less permissive notion of *robustness to inference* and the more permissive notion of *plausibility under inference* grows small.

6.2 Rationalizability

Again, I begin by considering the more stringent condition of robustness to inference. Following the previous section, define $\Theta_{a_i}^R$ to be the set of all complete information games in which some action a_i is rationalizable.

Definition 7. For each player i and action $a_i \in A_i$, let $\Theta_{a_i}^R$ be the set of all payoffs given which action a_i is rationalizable for player i .

The earlier results for equilibrium (Lemma 1 and Proposition 2) suggest the following parallel statements for rationalizability: (1) robustness to inference is characterized by whether the set of plausible expected payoffs is eventually contained in $\Theta_{a_i}^R$, and (2) if there is a limiting common prior, then robustness to inference is characterized by whether a_i is strictly rationalizable in the limiting game. Neither statement turns out to hold.

Let us begin by discussing (1). Notice that although a_i is rationalizable at every payoff in $\Theta_{a_i}^R$, the chain of best responses rationalizing action a_i can vary across the set. The simple example below shows that even if a_i is rationalizable in each of two different games, it can fail to be rationalizable when the player has uncertainty over which of the two payoffs will be realized. This is closely related to the well-known fact that the set of rationalizable actions is not convex.

Example 3. Consider a two-player game. Player 1 assigns equal probability to the following two games:

	a_3	a_4		a_3	a_4
a_1	1, 1	-1, 0		a_1	-1, 0 1, 1
a_2	0, 1	1, 0		a_2	1, 0 0, 1

In the first game, action a_3 is uniquely rationalizable action for player 2, and in the second, action a_4 is uniquely rationalizable. So player 1 believes with probability 1/2 that player 2 will choose a_3 , and with probability 1/2 that player 2 will choose a_4 . Thus player 1's expected payoffs are given by

	a_3	a_4
a_1	0, 1/2	0, 1/2
a_2	1/2, 1/2	1/2, 1/2

and a_2 is his uniquely rationalizable action. This follows even though a_1 is a rationalizable action in both games.

Thus, common certainty in $\Theta_{a_i}^R$ does not imply rationalizability of a_i . In fact, even common certainty in *an arbitrarily small open set* within $\Theta_{a_i}^R$ does not guarantee rationalizability of a_i (see an example of this surprising fact in Appendix B.5).²² The key reason is that the set of payoffs can span two sets of payoff functions with different families of rationalizable actions. Actions that are rationalizable when players perceive different, arbitrarily close, payoffs, need not be rationalizable given any common perception of payoffs.

One way to address this is to require eventual common certainty in an open set of payoffs, across which the chain of best responses rationalizing action a_i remains constant. This condition turns out to be unnecessarily restrictive—an action a_i may be robust to inference even if every chain of best responses rationalizing action a_i fails to survive an arbitrarily small perturbation to payoffs. This is demonstrated in the example below.

Example 4. Consider the following complete information game

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

and notice that a_1 is rationalizable. Moreover, since a_1 is *strictly dominant* at all nearby payoffs (in Euclidean distance), a_1 remains rationalizable given common certainty of a small enough neighborhood of these payoffs. But consider the following perturbations:

	a_3	a_4		a_3	a_4	
a_1	1, $-\epsilon$	1, 0		a_1	1, 0	1, $-\epsilon$
a_2	0, $-\epsilon$	0, 0		a_2	0, 0	0, $-\epsilon$

Action a_1 remains rationalizable in both games, but a_1 is not a best reply to a_3 in the game on the left, and a_1 is not a best reply to a_4 in the game on the right. Thus, there is no chain of best responses rationalizing a_1 that holds on any (arbitrarily small) neighborhood of the original payoffs. Relatedly, action a_1 turns out not to be *strictly rationalizable* in the original game.

Lemma 3 organizes the above observations. It says that rationalizability of an action is robust to inference if players eventually have common certainty in some set of payoffs, across which a_i can be rationalized using the same best response chain, and only if players eventually have common certainty in $\Theta_{a_i}^R$.

Lemma 3. *The rationalizability of action a_i for player i is robust to inference if*

$$\mathbb{P}(\Theta(Z^n) \subseteq V) \rightarrow 1 \text{ as } n \rightarrow \infty \tag{10}$$

²²A very nice two-player example in the concurrent work of [Chen and Takahashi \(2017\)](#) shows this as well.

for some set V , where a_i can be rationalized using the same chain of best-responses for all payoffs in V , and only if

$$\mathbb{P}(\Theta(Z^n) \subseteq \Theta_{a_i}^R) \rightarrow 1 \text{ as } n \rightarrow \infty. \quad (11)$$

Again, to better understand these conditions, let us consider the case in which there is a limiting common prior (Assumption 2). In this case, the sufficient condition above reduces to strict rationalizability of a_i in the limiting game, and the necessary condition reduces to a property that I now introduce:

Recall that strict rationalizability can be defined as the limit of a process of iterative elimination of actions that are never a strict best reply. It is well known that this procedure is sensitive to the manner of elimination. Consider specifically all the orders of elimination in which *only one* action is eliminated at a time. Formally, 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 $\alpha_{-i} \in \Delta(W_{-i}^{k-1})$. 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^∞ to be the intersection of *all* sets W_i^∞ that can be constructed in this way.

Definition 8. Say that an action a_i is weakly strict-rationalizable if $a_i \in \mathcal{W}_i^\infty$.

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

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

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

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

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.

Weak-strict rationalizability turns out to be the right condition because it characterizes the interior of U_{i,a_i}^R ; that is, action a_i is weakly-strict rationalizable in every game in the interior of U_{i,a_i}^R (see Lemma 5 in the appendix). From this, and Lemma 3, it follows that:

Proposition 4. *Suppose there is a limiting common prior P . Then, rationalizability of action a_i for player i is robust to inference if a_i is strictly rationalizable in the incomplete information game with common prior P , and only if a_i is weakly-strict rationalizable in the incomplete information game with common prior P .*

Observe that refinement is obtained despite the negative results of [Weinstein and Yildiz \(2007\)](#). One way to understand this is to observe that (with probability 1) any sequence of types t^n from $T(Z^n)$ converges in the uniform-weak topology ([Chen et al., 2010](#)), while the negative result in [Weinstein and Yildiz \(2007\)](#) relies on types that converge only in the (coarser) product topology; see [Section 10](#) for an extended discussion. Loosely speaking, the tail beliefs of types in [Weinstein and Yildiz \(2007\)](#) are permitted to put high probability on payoff functions that receive low probability at all lower orders. In the approach considered in this paper, players have common certainty in a (shrinking) set of payoffs, so higher order beliefs are required to have similar supports to the lower order beliefs. In fact, the sufficiency direction of this result follows almost directly from lower hemi-continuity of strict rationalizability in the uniform-weak topology ([Chen et al., 2010](#)), although the necessity direction requires new arguments.

Another precedent for the sufficient result appears in [Morris et al. \(2012\)](#) and [Takahashi \(2017\)](#), where the concept of *robustly rationalizable* is proposed. This concept turns out to be (nearly) equivalent to robustness to inference.²³ Finally, the concurrent paper [Chen and Takahashi \(2017\)](#) provides a necessary condition related to weak-strict rationalizability.

Turning now to plausibility of inference, for every distribution $\nu \in \Delta(\Theta \times A_{-i})$, define

$$BR_i(\nu) := \operatorname{argmax}_{a_i \in A_i} \mathbb{E}_\nu u(a_i, a_{-i}, \theta)$$

to be player i 's set of best replies given belief ν .

Lemma 4. *Rationalizability of action a_i for player i is plausible under inference if and only if there exists such that*

$$\mathbb{P}(\exists \text{ family } (R_i)_{i \in \mathcal{I}} \text{ s.t. } a_i \in BR_i(\nu_i) \text{ for some } \nu_i \in \Delta(\Theta(Z^n) \times R_{-i}) \forall i) \rightarrow 1 \text{ as } n \rightarrow \infty \quad (12)$$

Again, let us restrict to a limiting common prior ([Assumption 2](#)).

²³The present setting differs from this prior work in that Θ may not be finite, and the notion of robustness is across all types with approximate common certainty in *neighborhoods* of the true parameter. These differences are not crucial for the characterization of asymptotic strategic behavior, but they do matter for the rate results in the subsequent section.

Proposition 5. *Suppose there is a limiting common prior. Then, rationalizability of action a_i for player i is robust to inference if and only if a_i is rationalizable in the incomplete information game with common prior P .*

Thus with large quantities of data, we again see that the difference between the less permissive notion of *robustness to inference* and the more permissive notion of *plausibility under inference* is small. These concepts diverge substantially in the next section, where we consider small quantities of data.

7 Strategic Behavior: Finite-Sample

Propositions 2 and 4 tell us that when there is a limiting common prior, and players observe a large quantity of data, then it is largely without loss to make the simplifying assumption that players share a common prior (as if having observed infinite data). Specifically, predictions eventually hold for *all* plausible types only if they hold “strictly” in the limiting common prior game, and they eventually for *some* plausible types only if they hold in the limiting common prior game.

These conclusions no longer hold when the number of common observations is small. The equilibrium (and rationalizable) sets for interim games given small n observations need not be the same as the equilibrium (and rationalizable) sets in the limiting game. I seek below to quantify the probability with which they differ. Here I again emphasize the utility of $\bar{p}_{\mathbf{a}}^{NE}$ (and the related measures defined in Section 4) as *continuous* metrics; for example, we can ask how large $\bar{p}_{\mathbf{a}}^{NE}$ is for a fixed action profile \mathbf{a} and number of observations n . Informally, we should have more confidence in predictions that hold with high probability for small n , than predictions that only hold when players commonly observe a large number of realizations.

Throughout this section, I impose the additional assumption that observations are i.i.d. from a finite set.

Assumption 4. $Z_1, \dots, Z_n \sim_{i.i.d.} Q$ and take values in a finite set \mathcal{Z} .

Robust Predictions. For every action profile \mathbf{a} and distribution $P \in \Delta(\Theta)$, define

$$\bar{\delta}_{\mathbf{a}} := \max \{ \delta : \mathbb{E}_Q(\theta) \in \Theta_{\mathbf{a}}^{NE} \quad \forall Q \in \{P\}^{\epsilon} \},$$

where $\{P\}^{\epsilon}$ denotes the δ -neighborhood of distribution P , measured in the Prokhorov metric. The parameter $\bar{\delta}_{\mathbf{a}}$ is the radius of the maximal neighborhood of P such that \mathbf{a} is a Nash equilibrium given any common prior in this neighborhood. The larger the parameter $\bar{\delta}_{\mathbf{a}}$ is, the larger the set of (common) beliefs on which the equilibrium \mathbf{a} holds.

The analogous parameter for rationalizability is this: Define $\bar{\delta}_{a_i}$ to be the maximal value of δ such that there exists a family $(R_i)_{i \in \mathcal{I}}$ where for all $Q \in \{P\}^\epsilon$, $(R_i)_{i \in \mathcal{I}}$ is closed under best response in the complete information game with payoffs $\mathbb{E}_Q(\theta)$. This is the maximal neighborhood of P such that a_i is rationalizable (using the same chain of best responses) given any common prior in this neighborhood.

Proposition 11 uses these parameters to lower bound $\underline{p}_n^{NE}(a)$ and $\underline{p}_n^R(i, a_i)$ for all quantities of data n .

Proposition 6. *Suppose there is a limiting common prior P . If \mathbf{a} is a strict Nash equilibrium in the game with common prior P , and a_i is rationalizable for player i , then*

$$\underline{p}_n^{NE}(\mathbf{a}) \geq 1 - \frac{1}{\bar{\delta}_{\mathbf{a}}} \mathbb{E} \left(\sup_{\mu \in \mathcal{M}} d(\mu(Z^n), P) \right) \quad \forall n \geq 1 \quad (13)$$

and

$$\underline{p}_n^R(i, a_i) \geq 1 - \frac{1}{\bar{\delta}_{a_i}} \mathbb{E} \left(\sup_{\mu \in \mathcal{M}} d(\mu(Z^n), P) \right) \quad \forall n \geq 1 \quad (14)$$

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

Thus, if \mathbf{a} is a Nash equilibrium in the limiting common prior game, then given that players have observed n realizations, the probability that it is also a Bayesian Nash equilibrium is at least the probability in (17) above.

Two observations are in order. First, the bounds in Proposition 11 are increasing in the strictness of the parameters $\bar{\delta}_{\mathbf{a}}$ and $\bar{\delta}_{a_i}$: thus, the “stricter” the prediction, the fewer observations are necessary for the predictions to hold. They are also decreasing in

$$\mathbb{E} \left(\sup_{\mu \in \mathcal{M}} d(\mu(Z^n), P) \right),$$

which is the expected distance from P to the farthest plausible belief. So the quicker players *commonly* learn, the fewer the observations are necessary.

Plausible Predictions. Turning now to plausibility given n observations, define for each action profile \mathbf{a} the set of all sequences of data given which \mathbf{a} is a plausible equilibrium:

$$\mathbb{Z}_{\mathbf{a}} := \{\mathbf{z} : \Theta(\mathbf{z}) \cap \Theta(i, a_i) \neq \emptyset \quad \forall i\}$$

Additionally, write $\hat{Q}(\mathbf{z}_n)$ for the empirical measure over the signal set \mathcal{Z} that is associated with data sequence \mathbf{z}_n . Notice that the probability $\bar{p}_n^{NE}(a)$ is exactly the probability that the realized data (of size n) is in the set \mathbb{Z}_a . Finally, let

$$Q_{\mathbf{a}}^* = \operatorname{argmin}_{\hat{Q} \in \{\hat{Q}(\mathbf{z}) : \mathbf{z} \in \mathbb{Z}_{\mathbf{a}}\}} D(\hat{Q} \| Q)$$

be the empirical measure (associated with a data set in \mathbb{Z}_a^n) that is closest in Kullback-Leibler distance²⁴ to the actual signal-generating distribution Q (see Assumption 4).

Similarly, for player i and action a_i , define

$$\mathbb{Z}_{a_i}^R := \{\mathbf{z} : \exists \text{ family } (R_i)_{i \in \mathcal{I}} \text{ s.t. } \forall i \in \mathcal{I}, \\ a_i \in BR(\nu_i) \text{ for some } \nu_i \in \Delta(\Theta(\mathbf{z}) \times R_{-i})\}$$

to be all datasets given which action a_i is a plausible rationalizable action for player i . Let

$$Q_{a_i}^* = \underset{\hat{Q} \in \{\hat{Q}(\mathbf{z}) : \mathbf{z} \in \mathbb{Z}_{a_i}^R\}}{\operatorname{argmin}} D(\hat{Q} \| Q)$$

be the empirical measure (associated with a data set in $\mathbb{Z}_{a_i}^R$) that is closest in Kullback-Leibler distance to the actual signal-generating distribution Q .

Then, application of Sanov's theorem directly yields:

Proposition 7. *Suppose there is a limiting common prior P . If \mathbf{a} is not a strict Nash equilibrium in the game with common prior P , and a_i is not rationalizable for player i , then for every $n \in \mathbb{Z}_+$,*

$$\bar{p}_n^{NE}(a) \leq (n+1)^{|\mathcal{Z}|} 2^{-nD(Q_a^* \| Q)}$$

and

$$\bar{p}_n^R(i, a_i) \leq (n+1)^{|\mathcal{Z}|} 2^{-nD(Q_{a_i}^* \| Q)}$$

8 Application: Data Design

So far, we have taken the data-generating process and the set of learning rules to be exogenously determined. 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.

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 $\mathcal{Z} = \mathcal{X} \times Y$, where \mathcal{X} is a set of observable features, and Y is the space of payoff-relevant outcomes. For example,

²⁴The Kullback-Leibler distance from distribution q to Q is

$$D(\hat{Q} \| Q) = \sum_{-\infty}^{\infty} \hat{Q}(x) \log \frac{Q(x)}{\hat{Q}(x)}$$

- the set \mathcal{X} might describe physical characteristics of a laptop (weight, battery life, resolution), while Y 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

$$(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n),$$

where each y_i is drawn independently and distributed according to an unknown conditional distribution $P(y|\mathbf{x} = \mathbf{x}_i)$. The payoff-relevant unknown is the value 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 $\phi : \mathcal{X} \rightarrow Y$ from the data, and evaluating ϕ at \mathbf{x} . A large literature examines problems like this, and approaches for inference of ϕ . The examples below focus on two canonical cases.

Example 1 sets $Y = \{0, 1\}$, so that the problem is one of classification: players want to learn which values in \mathcal{X} map to $y = 0$ and which values map to $y = 1$. 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, 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.

8.1 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	y, y	$0, \frac{1}{2}$
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 $y \in \{0, 1\}$ (interpret $y = 0$ to mean loss and $y = 1$ to mean success). Is joining the suit a rationalizable action?

For concreteness, let $\mathcal{X} = [-c, c]^{\bar{p}}$, so that every suit is described by \bar{p} characteristics, each normalized to lie within the interval $[-c, c]$. Every observation (\mathbf{x}, y) describes the characteristics and outcome of a past class action suit. Observations are drawn i.i.d. from a distribution Q on $\mathcal{X} \times \Theta$ with the properties that: (1) $\text{marg}_{\mathcal{X}} Q$ is uniform over \mathcal{X} ; and (2) there is some $p^* < \bar{p}$ such that the conditional distribution $Q(\cdot | \mathbf{x})$ is a point mass on 1 if $x^k \in [-c^*, c^*]$ for each $k \leq p^*$, and it is a point mass on 0 otherwise. Notice that only p^* of the \bar{p} characteristics matter for the outcome of the suit.

An external agency chooses a transformation of the realized data

$$(\mathbf{x}^1, \theta^1), \dots, (\mathbf{x}^n, \theta^n) \quad (15)$$

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 $\bar{p} - p^*$ characteristics. In the following, I will take these features to be symmetric, so that the agency's choice is simply an integer $p \in \{p^*, p^* + 1, \dots, \bar{p}\}$. Thus, instead of observing (15), players observe

$$\mathbf{z}_n = ((\mathbf{x}_{1:p}^1, \theta^1), \dots, (\mathbf{x}_{1:p}^n, \theta^n)), \quad (16)$$

where $\mathbf{x}_{1:p}$ is the truncation of vector \mathbf{x} to its 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 $\phi : [-c, c]^p \rightarrow \Theta$ that can be written as

$$\phi(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in [\underline{c}_1, \bar{c}_1] \times \dots \times [\underline{c}_p, \bar{c}_p] \\ 0 & \text{otherwise} \end{cases}$$

where $\underline{c}_1, \bar{c}_1, \dots, \underline{c}_p, \bar{c}_p \in [-c, c]$. Given data set \mathbf{z} , let $\Phi_{\mathbf{z}}$ be the set of all rectangular classification rules that exactly fit the observed data. Examples are shown in the figure below.

The data provided about the current suit is the truncated vector $\mathbf{x}_{1:p}^*$. Define \mathcal{M} be the set of all functions μ with the property that each $\mu(\mathbf{z})$ has support in the predictions made by classification rules consistent with the data: $\{\phi(\mathbf{x}^*), \phi \in \Phi_{\mathbf{z}}\}$.²⁵ Suppose joining is rationalizable for both players under complete information; that is, indeed $\mathbf{x}_p^* \in [\underline{c}^*, \bar{c}^*]$ for every $p \leq p^*$. Then:

²⁵The following is a Bayesian interpretation of \mathcal{M} . The set of states of the world is $\Omega = Y \times \Phi \times \mathcal{Z}^\infty$, so that a state consists of a value of y , a function ϕ , and an infinite sequence of observations from \mathcal{Z} . Let M be the set of all probability distributions over $Y \times \Phi$, with the property that the induced 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 ,

Proposition 8. For every quantity of data $n \geq 1$, and for both players $i \in \{1, 2\}$:

- (a) the probability $p_n^R(i, \text{Join})$ is monotonically decreasing in the number of reported characteristics p^* .
- (b) $p_n^R(i, \text{Join}) \rightarrow 0$ as the number of characteristics $p^* \rightarrow \infty$.

Thus, if the agency wants to minimize the probability that joining is rationalizable given n observations, it should report as many characteristics as possible ($p = \bar{p}$). 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 ($y = 1$), 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 deters participation in 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.

8.2 Example 2: Regression

Consider the game described in Section 2, but suppose that each investment is described by p covariates $(x^1, \dots, x^p) \in \mathbb{R}^p$. Observations are pairs $(x_k^1, \dots, x_k^p, y_k)$, where

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

for $k = 1, \dots, n$. 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

where $\text{marg}_{\mathcal{X}} P_\phi$ is uniform over \mathcal{X} , and the conditional distribution $P(\cdot | \mathbf{x})$ 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 $\{(\theta, \phi, \mathbf{z}) : \theta = \phi(\mathbf{x}^*)\}$.

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 $P_\mu(\theta | \mathcal{H}_n)(\mathbf{z})$ over Θ . Let $\Delta_{\mathbf{z}} \subseteq \Delta(Y)$ be the set of all such posteriors. The set of learning rules \mathcal{M} consists of all maps $\mu : \bigcup_{n=1}^\infty \mathcal{Z}^n \rightarrow \Delta(Y)$ such that $\mu(\mathbf{z}) \in \Delta_{\mathbf{z}}$ for every \mathbf{z} . See Appendix D.2 for the argument of equivalence.

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 $\hat{\beta}^{LS} = (\hat{\beta}_0^{LS}, \dots, \hat{\beta}_p^{LS})$ be the least-squares solution

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

and let $\hat{\phi}_{LS}(x) = \hat{\beta}^{LS} \cdot x$ be the associated function. 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 $CI(\mathbf{z})$.

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

Proposition 9. *Suppose $\phi(\mathbf{x}^*) > 0$. Then, for every fixed quantity of data $n \geq 1$, and for $i \in \{1, 2\}$,*

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

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 ($p = \bar{p}$). 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 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.

9 Extensions

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

9.1 Misspecification

The main results hold under a weakening of Assumption 2, which I define below:

Definition 9 ((Approximate Limiting Common Prior.). *For any $\epsilon \geq 0$, say that the class of learning rules \mathcal{M} has a $(1 - \epsilon)$ -limiting common prior P if*

$$\lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{M}} d(\mu(Z^n), P) \leq \epsilon \text{ a.s.}$$

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

According to this definition, the class of learning rules \mathcal{M} has a $(1 - \epsilon)$ -limiting common prior P if the set of induced first order beliefs converges almost surely (in the Hausdorff distance induced by d) to an ϵ -neighborhood of P . Notice that Assumption 2 is nested as the $\epsilon = 0$ case. Propositions 2 and ?? can be weakened to show the following result. (In reading this, recall that if \mathcal{M} has a $(1 - \epsilon)$ -limiting common prior P , then it also has a $(1 - \epsilon')$ -limiting common prior P for every $\epsilon' > \epsilon$.)

Proposition 10. (a) *Suppose \mathcal{M} has a $(1 - \delta_{\mathbf{a}})$ -limiting common prior P . Then, the equilibrium property of \mathbf{a} is robust to inference if and only if \mathbf{a} is a strict equilibrium in the game with common prior P .*

(b) *Suppose \mathcal{M} has a $(1 - \delta_{a_i}^R)$ -limiting common prior P . Then, the rationalizability of action a_i is robust to inference if a_i is strictly rationalizable in the game with common prior P .*

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 $\delta_{\mathbf{a}^*}^{NE}$ neighborhood (respectively, $\delta_{a_i^*}^R$ -neighborhood) of the degenerate belief on θ^* .

9.2 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{M} consists of a single learning rule. So Proposition 1 need not hold. Moreover, Carlsson and van Damme (1993) and Kajii and Morris (1997) (among others) have shown that strict Nash equilibria are not robust to higher-order uncertainty about private opponent information. Thus, Propositions 2 and 4 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, Propositions 2 and 4 continue to hold, and the bounds in Proposition 11 can be revised as follows (where I is the number of players).

Proposition 11. *Suppose there is a limiting common prior P . If \mathbf{a} is a strict Nash equilibrium in the game with common prior P , and a_i is rationalizable for player i , then*

$$\underline{p}_n^{NE}(\mathbf{a}) \geq \left(1 - \frac{1}{\delta_{\mathbf{a}}} \mathbb{E} \left(\sup_{\mu \in \mathcal{M}} d(\mu(Z^n), P) \right) \right)^I \quad \forall n \geq 1$$

and

$$\underline{p}_n^R(i, a_i) \geq \left(1 - \frac{1}{\delta_{a_i}} \mathbb{E} \left(\sup_{\mu \in \mathcal{M}} d(\mu(Z^n), P) \right) \right)^I \quad \forall n \geq 1$$

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

10 Related Literature

Suppose an analyst does not know the exact beliefs that players hold. Can he be reasonably certain that the solutions in his model are close to the solutions given the actual beliefs? Early answers to this question considered the strategic properties of types whose beliefs were close up to order k for large k (Rubinstein, 1989; Mertens and Zamir, 1985; Brandenburger and Dekel, 1993). Several authors have demonstrated that this notion of nearby (which corresponds to the product topology on types) leads to surprising and counterintuitive conclusions, in particular, even strict equilibria and strictly rationalizable actions are fragile to perturbations (Rubinstein, 1989; Weinstein and Yildiz, 2007).

Dekel et al. (2006), Chen et al. (2010), and Chen et al. (2017) subsequently developed and characterized finer metric topologies on types under which the desired continuity properties hold. In particular, the *uniform-weak topology* proposed in Chen et al. (2010) considers two types to be close if they have similar first-order beliefs, attach similar probabilities to other players having similar first-order beliefs, and so forth. Under the assumption of a limiting common prior, all types consistent with Common Inference converge in the uniform-weak topology. Thus, the property of robustness to inference, considered in Section 6, can be interpreted as requiring persistence across a subset of perturbations in the uniform-weak topology.

This approach contrasts with Carlsson and van Damme (1993) and Kajii and 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.²⁶ In addition, I do not require that these beliefs are consistent with a common prior, thus allowing for common knowledge disagreement. Collectively, these differences lead to rather different robustness results.

²⁶For example, the construction of beliefs used in Weinstein and 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 and 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.

Finally, the definition of robustness to inference for rationalizability (nearly) coincides with the concept of *robustly rationalizable* proposed in [Morris et al. \(2012\)](#), and subsequently characterized in the concurrent work of [Takahashi \(2017\)](#) and [Chen and Takahashi \(2017\)](#). A more detailed discussion of this relationship appears in [Section 6.1](#) following [Proposition 4](#).

11 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.

A Supplementary Material to Section 5

A.1 Proof of Proposition 1

Preliminaries. Let $T_i^k = \Delta(X_{k-1}) = \Delta(\Theta \times T_{-i}^{k-1})$ denote the set of k -th order beliefs for player i .²⁷ Let d_i^0 be the Euclidean norm on Θ (see Section 2.1), and recursively for $k \geq 1$, define d_i^k to be the Prokhorov distance²⁸ on $\Delta(\Theta \times T_{-i}^{k-1})$ induced by the metric $\max\{d_i^0, d_i^{k-1}\}$ on $\Theta \times T_{-i}^{k-1}$. As in the main text, since only symmetric type spaces are considered, player subscripts are dropped throughout.

Additionally, I use E^δ to mean the δ -neighborhood of E (where the metric should be clear from context).

Uniform convergence implies common learning (only if): Fix any dataset \mathbf{z} . We can decompose the set of types $T(\mathbf{z})$ into the Cartesian product $\prod_{k=1}^\infty H_{\mathbf{z}}^k$, where $H_{\mathbf{z}}^1 = \Delta_{\mathbf{z}}$ and for each $k > 1$, $H_{\mathbf{z}}^k$ is recursively defined

$$H_{\mathbf{z}}^k = \left\{ t^k \in T^k : (\text{marg}_{T^{k-1}} t^k)(H_{\mathbf{z}}^{k-1}) = 1 \text{ and } \text{marg}_{\Theta} t^k \in H_{\mathbf{z}}^1 \right\}. \quad (17)$$

This is the set of k -th order beliefs of types in $T(\mathbf{z})$. Define

$$\Delta_{\mathbf{z}} := \{ \mu(\mathbf{z}) : \mu \in \mathcal{M} \}$$

to be the set of first-order beliefs induced from learning rules in \mathcal{M} , and

$$\delta^* := d_1(\Delta_{\mathbf{z}}, P)$$

to be the largest distance between P and any belief in $\Delta_{\mathbf{z}}$. I show below that for every dataset \mathbf{z} , the distance between any type $t \in T(\mathbf{z})$ and the type t_P with common certainty of P is upper bounded by δ^* .

Claim 1. *For every $k \geq 1$ and $t \in T(\mathbf{z})$,*

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

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

$$d^1(t, t_P) \leq d(\Delta_{\mathbf{z}}, P) = \delta^*. \quad (18)$$

²⁷Working 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.

²⁸Recall that the *Levy-Prokhorov* distance ρ between measures on metric space (X, d) is defined

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

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

Now suppose $H_{\mathbf{z}}^k \subseteq \{t_P^k\}^{\delta^*}$. Consider any measurable set $E \subseteq T^k$. If $t_P^k \in E$, then $t_P^{k+1}(E) = 1$ by definition of t_P . Also

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

using (19) in the final equality and the assumption that $H_{\mathbf{z}}^k \subseteq \{t_P^k\}^{\delta^*}$ in the inequality preceding it. So

$$t_P^{k+1}(E) \leq t^{k+1}(E^{\delta^*}) + \delta^*. \quad (19)$$

If $t_P^k \notin E$, then $t_P^{k+1}(E) = 0$ (again by definition of t_P), so (21) follows trivially. Thus $t_P^{k+1}(E) \leq t^{k+1}(E^{\delta^*}) + \delta^*$ for every measurable $E \subseteq T^k$. Using this and (20),

$$d^{k+1}(t, t_P) \leq \delta^* \quad (20)$$

as desired. \square

Now consider any $\epsilon \geq 0$ and $p \in [0, 1)$ and choose $\delta^* \leq \min(\epsilon, 1 - p)$. I will show that $\{P\}^\epsilon$ is common p -belief for all types in $\{t_P\}^{\delta^*}$. Fix an arbitrary $t \in \{t_P\}^{\delta^*}$. Trivially,

$$\text{marg}_{\Theta} t \in \{P\}^{\delta^*} \subseteq \{P\}^\epsilon$$

Moreover,

$$t^1(\{P\}^\epsilon) \geq t^1(\{P\}^{\delta^*}) \geq 1 - \delta^* \geq p$$

where the second inequality follows from $d_1(t, t_P) \leq \delta^*$.²⁹ Thus, every type in $\{t_P\}^{\delta^*}$ assigns at least probability p to the event that every player has a first-order belief in $\{P\}^\epsilon$, or $\{t_P\}^{\delta^*} \subseteq \mathcal{B}^p(\{P\}^\epsilon)$. Now suppose that $\{t_P\}^{\delta^*} \subseteq [\mathcal{B}^p]^k(\{P\}^\epsilon)$. By a similar argument,

$$t^{k+1}([\mathcal{B}^p]^k(\{P\}^\epsilon)) \geq t^{k+1}(\{t_P\}^{\delta^*}) \geq t^{k+1}(\{t_P^k\}^{\delta^*}) \geq 1 - \delta^* \geq p,$$

So also $\{t_P\}^{\delta^*} \subseteq [\mathcal{B}^p]^{k+1}(\{P\}^\epsilon)$. Thus, $\{t_P\}^{\delta^*} \subseteq [\mathcal{B}^p]^k(\{P\}^\epsilon)$ for every order k , and it follows that $\{t_P\}^{\delta^*} \subseteq C^p(\{P\}^\epsilon)$ as desired.

Under condition (7), the set

$$\left\{ \mathbf{z} : \lim_{n \rightarrow \infty} d(\Delta_{\mathbf{z}_n}, P) = 0 \right\}$$

has measure 1. Thus, for every $\epsilon > 0$ and $p \in [0, 1]$, there is a measure-1 set of sequences \mathbf{z} such that $\{P\}^\epsilon$ is common p -belief for all types in $T(\mathbf{z}_{1:n})$ for n sufficiently large.

²⁹By definition, $t_P^1(\{P\}) = 1$, and since $d_1(t, t_P) \leq \delta^*$, necessarily $t_P^1(\{P\}) \leq d_1(\{P\}^{\delta^*}) + \delta^*$, which implies that $d_1(\{P\}^{\delta^*}) \geq 1 - \delta^*$.

Common learning implies uniform convergence (if): Suppose condition (7) is not satisfied; then, there exists some $\underline{\epsilon} > 0$ such that

$$Z^* := \left\{ \mathbf{z} \mid \lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{M}} d(\mu(\mathbf{z}_{1:n}), P) \leq \underline{\epsilon} \right\}$$

has positive measure. Choose any $\epsilon < \underline{\epsilon}$. Then, for all $\mathbf{z} \in Z^*$, the set $\{P\}^\epsilon$ fails to be eventually common p -belief for all $p \in [0, 1]$ and for all types in $T(\mathbf{z})$. So common learning fails for some types in $T(Z^n)$ with positive probability, for n arbitrarily large.

B Supplementary Material to Section 6

Below, I use $\sigma_{\mathbf{a}}$ to denote the strategy profile in which every player i 's strategy is constant on the action a_i . This notation assumes implicitly that players have type sets $(T_i)_{i \in \mathcal{I}}$, and formally I mean that $\sigma_i(t_i) = a_i$ for every $t_i \in T_i$ and player i . The dependence on the type space is dropped throughout for notational simplicity. When data \mathbf{z}_n is realized, assume always that $T_i = T(\mathbf{z}_n)$ for each player i .

B.1 Proof of Lemma 1

Consider the event in which the set of plausible payoffs satisfies $\Theta(\mathbf{z}_n) \subseteq \Theta_{\mathbf{a}}^{NE}$. Since every player's first-order belief has support in $\Theta_{\mathbf{a}}^{NE}$, every player i perceives a_i to be a best reply to a_{-i} . Thus, the strategy profile $\sigma_{\mathbf{a}}$ is a Bayesian Nash equilibrium. This establishes that $\Theta(\mathbf{z}_n) \subseteq \Theta_{\mathbf{a}}^{NE}$ is sufficient for $\sigma_{\mathbf{a}}$ to be a BNE, so the assumption in (8) directly implies that $\underline{p}_n^{NE}(a) \rightarrow 1$.

For the other direction, suppose that for all n sufficiently large, the measure of datasets

$$\{\mathbf{z}_n : \Theta(\mathbf{z}_n) \not\subseteq \Theta_{\mathbf{a}}^{NE}\}$$

is at least some constant $\delta > 0$ (independent of n). Conditioning on the event that $\Theta(\mathbf{z}_n) \not\subseteq \Theta_{\mathbf{a}}^{NE}$, there is some plausible payoff matrix $\theta \notin \Theta_{\mathbf{a}}^{NE}$. In this game θ , there exists some player i and action $a'_i \in A_i$ satisfying $u_i(a_i, a_{-i}, \theta) < u_i(a'_i, a_{-i}, \theta)$. So, action a_i is not a best reply against a_{-i} when player i has common certainty in θ . But since $\theta \in \Theta(\mathbf{z}_n)$, the type with common certainty in θ is in the set $T(\mathbf{z}_n)$. Thus, if $\Theta(\mathbf{z}_n) \not\subseteq \Theta_{\mathbf{a}}^{NE}$, then the strategy profile $\sigma_{\mathbf{a}}$ fails to be a BNE for some plausible type. Since the probability of the event $\Theta(\mathbf{z}_n) \not\subseteq \Theta_{\mathbf{a}}^{NE}$ is at least δ , we have that $\underline{p}_n^{NE}(a) \not\rightarrow 1$, concluding the proof.

B.2 Proof of Proposition 2

Suppose a is a strict Nash equilibrium in game with common prior P . Then, the expected payoff matrix $\mathbb{E}_P(\theta)$ is in the interior of $\Theta_{\mathbf{a}}^{NE}$. Moreover, by assumption

$\sup_{\mu \in \mathcal{M}} d(\mu(Z^n), P) \rightarrow 0$ a.s., so also

$$\sup_{\mu \in \mathcal{M}} \|\mathbb{E}_{\mu(\mathbf{z}_n)}[\theta] - \mathbb{E}_P[\theta]\|_\infty \rightarrow 0 \text{ a.s.}$$

Thus, the set $\Theta(\mathbf{z}_n) = \{\mathbb{E}_{\mu(\mathbf{z}_n)}(\theta) : \mu \in \mathcal{M}\}$ converges a.s. to $\{\mathbb{E}_P(\theta)\}$ in the Hausdorff metric induced by $\|\cdot\|_\infty$. This means that $\Theta(\mathbf{z}_n)$ is almost surely contained in Θ_a^{NE} and the first part of the corollary directly follows from Proposition 1.

Now suppose that a is not a strict NE in the game with common prior P , so that $\mathbb{E}_P(\theta) \notin \text{Int}(\Theta_a^{NE})$. By assumption, for every n sufficiently large, there is a positive measure of realizations \mathbf{z}_n given which $\mathbb{E}_P(\theta)$ is not an extremal point in $\Theta(\mathbf{z}_n)$. Combining this with the arguments above, the probability that $\Theta(\mathbf{z}_n)$ contains some payoff matrix $\theta \in \Theta_a^{NE}$ is bounded away from 0. Again apply Proposition 1 and we are done.

B.3 Proof of Lemma 2

Consider the event in which the set of plausible payoffs $\Theta(\mathbf{z}_n)$ has nonempty intersection with each $B_i(a_i, a_{-i})$. Let each player i have common certainty in a payoff matrix $\theta \in B_i(a_i, a_{-i})$. Then a_i is a best response to a_{-i} for each player i given his type. So the probability that $\Theta(\mathbf{z}_n) \cap B_i(a_i, a_{-i})$ for each i is a lower bound on $\underline{p}_{\mathbf{a}}^{NE}$, so that the condition in (9) directly implies $\underline{p}_{\mathbf{a}}^{NE} \rightarrow 1$.

Conversely, suppose that $\Theta(\mathbf{z}_n) \cap B_i(a_i, a_{-i}) = \emptyset$ for some player i . Then, there is no plausible payoff matrix $\theta \in B_i(a_i, a_{-i})$, so player i cannot perceive a_i to be a best reply to a_{-i} given his beliefs. Thus if (9) does not hold, then $\underline{p}_{\mathbf{a}}^{NE} \rightarrow 0$.

B.4 Proof of Proposition 2

Suppose there is a limiting common prior P . Then, $\Theta(\mathbf{z}_n) \rightarrow \{\mathbb{E}_P(\theta)\}$ in the Hausdorff metric. Under condition (9), necessarily $\mathbb{E}_P(\theta) \in B_i(a_i, a_{-i})$ for every i , implying that \mathbf{a} is a Nash equilibrium in the limiting game with common prior P . Conversely, suppose that \mathbf{a} is not a Nash equilibrium in the limiting game with common prior P . Then, there is some player i for whom $\mathbb{E}_P(\theta) \notin B_i(a_i, a_{-i})$. This further implies that

$$\{\mathbb{E}_P(\theta)\}^\epsilon \not\subseteq B_i(a_i, a_{-i}) \text{ for some } \epsilon > 0$$

where $\{\mathbb{E}_P(\theta)\}^\epsilon$ is the ϵ -neighborhood of the payoff matrix $\mathbb{E}_P(\theta)$. Since $\Theta(Z^n) \subseteq \mathbb{E}_P(\theta)$ with probability arbitrarily close to 1 for n sufficiently large, it follows that condition (9) does not hold.

B.5 Example

Consider the following four player game. Players 1 and 2 choose between actions in $\{a, b\}$, and player 3 chooses between matrices from $\{l, r\}$. The expected payoffs for players 1-3 are given (in order) by:

$$\begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 1, 1, 0 & 0, 0, 0 \\ 0, 0, 0 & 0, 0, 0 \end{array} \\
 & (l)
 \end{array}
 \qquad
 \begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 0, 0, 0 & 0, 0, 0 \\ 0, 0, 0 & 1, 1, 0 \end{array} \\
 & (r)
 \end{array}
 \tag{21}$$

A fourth player chooses between $\{\text{Match}, \text{Mismatch}\}$, where his action is not relevant to the other players' payoffs. Player 4 receives a payoff of 1 from Match if players 2 and 3 choose the same action (both choose a or both choose b), and he receives 0 otherwise. He receives a payoff of 1 from Mismatch if players 2 and 3 choose different actions, and 0 otherwise.

In every game with payoffs sufficiently close³⁰ to those above, Match is rationalizable for player 4, because one (or both of) actions a and b is simultaneously rationalizable for player 2 and 3.³¹ Nevertheless, I will show that Match fails to be rationalizable for a sequence of types with common certainty in increasingly small neighborhoods of the payoffs described above. The basic idea is that player 4 can believe that a is uniquely rationalizable for player 1, while b is uniquely rationalizable for player 2.

To see this, consider the following two perturbed versions of the payoff matrix above where $\epsilon > 0$:

$$\begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 1, 1, 0 & 0, 0, 0 \\ 0, 0, 0 & -\epsilon, 0, 0 \end{array} \\
 & (l)
 \end{array}
 \qquad
 \begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 0, 0, -\epsilon & 0, 0, -\epsilon \\ 0, 0, -\epsilon & 1, 1, -\epsilon \end{array} \\
 & (r)
 \end{array}
 \tag{22}$$

and

$$\begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 1, 1, -\epsilon & 0, 0, -\epsilon \\ 0, 0, -\epsilon & 0, 0, -\epsilon \end{array} \\
 & (l)
 \end{array}
 \qquad
 \begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} -\epsilon, 0, 0 & 0, 0, 0 \\ 0, 0, 0 & 1, 1, 0 \end{array} \\
 & (r)
 \end{array}
 \tag{23}$$

³⁰Distance is measured in the Euclidean metric.

³¹If neither l nor r are strictly dominated for Player 1, then all actions are rationalizable for player 1-3, and if either l or r is strictly dominated for player 1, then one of the following will be a rationalizable family for player 1-3: $\{l\} \times \{a\} \times \{a\}$, $\{l\} \times \{a, b\} \times \{a, b\}$, $\{r\} \times \{b\} \times \{b\}$, or $\{r\} \times \{a, b\} \times \{a, b\}$.

If player 1 has common certainty in the payoffs in (24), then a is his uniquely rationalizable action: r is strictly dominated by l for player 4, and in the surviving game, b is strictly dominated by a for player 1. By a similar argument, if player 2 has common certainty in the payoffs in (25), then b is his uniquely rationalizable action. These statements hold for ϵ arbitrarily small. Construct a sequence of types t_ϵ , where each type t_ϵ has common certainty that player 1 has common certainty in (24) and player 2 has common certainty in (25). Player 4 of type t_ϵ has only one rationalizable action, Mismatch. Take $\epsilon \rightarrow 0$ and the desired conclusion obtains.

B.6 Proof of Lemma 3

To show sufficiency of the condition in (10), suppose that for all payoffs in the set V , action a_i is rationalizable using the same chain of best responses. Then there exists a family $(R_j)_{j \in \mathcal{I}}$ closed under *strict* best reply for all games in V , where action $a_i \in R_i$. Clearly these relations are preserved when all players have first-order beliefs with support in V . Thus, for any sequence of types (t_1, \dots, t_I) with common certainty in V , there exists a family of sets $R_j \subseteq A_j$ where every action $a_j \in R_j$ is a best reply to a distribution $\pi \in \Delta(\Theta \times T_{-j} \times A_{-j})$ satisfying $\text{marg}_{\Theta \times T_{-j}} \pi = g(t_j)$ and $\pi(a_{-j} \in R_{-j}[t_{-j}]) = 1$, and also $a_i \in R_i$. The desired conclusion that action a_i is rationalizable for player i thus follows from the following result from Dekel et al. (2007):

Proposition 12 (Dekel et al. (2007)). *Fix any type profile $(t_j)_{j \in \mathcal{I}}$. Consider any family of sets $R_j \subseteq A_j$ such that every action $a_j \in R_j$ is a best reply to a distribution $\pi \in \Delta(\Theta \times T_{-j} \times A_{-j})$ that satisfies $\text{marg}_{\Theta \times T_{-j}} \pi = \kappa_j(t_j)$ and $\pi(a_{-j} \in R_{-j}[t_{-j}]) = 1$. Then, $R_j \subseteq S_j^\infty[t_j]$ for every player j .*

The condition in (10) thus guarantees that $p_n^R(a_i) \rightarrow 1$, so rationalizability of action a_i is robust to inference.

To show necessity of the condition in (11), consider the event in which data \mathbf{z}_n is realized and $\Theta(\mathbf{z}_n) \not\subseteq \Theta_{a_i}^R$. Then, there is a plausible payoff matrix $\theta \notin \Theta_{a_i}^R$, and a plausible belief with common certainty in θ . Clearly action a_i is not rationalizable for player i with this belief, so we are done.

B.7 Proof of Proposition 4

Suppose a_i is strictly rationalizable in the game θ . Then, there exists a family $(R_j)_{j \in \mathcal{I}}$ with the property that each $R_j \subseteq A_j$, action $a_i \in A_i$, and for every player i , each action $a_j \in R_j$ is a strict best response to some distribution over R_{-j} .

Players will eventually have common certainty of the payoffs in $\Theta_{a_i}^R$ only if the limiting game θ is in the interior of $\Theta_{a_i}^R$. Next I show that θ is in the interior of $\Theta_{a_i}^R$ only if a_i is weakly strict-rationalizable in the limiting game θ . This follows from

the stronger result below, which says that weak strict rationalizability characterizes the interior of $\Theta_{a_i}^R$.

Lemma 5. $\theta \in \text{Int}(\Theta_{a_i}^R)$ if and only if a_i is weakly strict-rationalizable in the complete information game with payoffs θ .

Proof. If: Suppose the game θ is not in the interior of $\Theta_{a_i}^R$. There must then exist a sequence $\theta^n \rightarrow \theta$ (in the sup-metric), where for large n , the game θ^n is also not in the interior of $\Theta_{a_i}^R$. Thus in each late game θ^n , there is an order of elimination of strictly dominated strategies that removes a_i . Moreover, since action sets are finite, there is a finite number of possible such orders of elimination. This implies existence of a subsequence along which the same order of iterated elimination of strategies removes a_i . At the limiting payoffs θ , action a_i must fail to survive elimination of *weakly* dominated strategies along this order, and is therefore not weakly strict-rationalizable.

Only if: Suppose a_i is not weakly strict-rationalizable. Then, there exists a sequence of sets $(W_j^k)_{k \geq 1}$ for every player j satisfying the recursive description in Section 6.2, such that $a_i \notin W_i^K$ for some $K < \infty$. To show that θ is not in the interior of $\Theta_{a_i}^R$, I construct a sequence of payoff functions θ^n with $\theta^n \rightarrow \theta$ such that a_i is not rationalizable in any late game along this sequence.

For every $n \geq 1$, construct the payoff function u^n according to the following procedure. First, for every player j , let $\theta^{n,0} = \theta$. Then, for every $l \geq 1$, define $\theta^{n,l}$ such that

$$u_j(a_j, a_{-j}, \theta^{n,l}) = \begin{cases} u_j(a_j, a_{-j}, \theta^{n,l-1}) + \epsilon/n & \forall a_j \in W_j^l, a_{-j} \in A_{-j} \\ u_j(a_j, a_{-j}, \theta^{n,l-1}) & \forall a_j \notin W_j^l, a_{-j} \in A_{-j} \end{cases}$$

That is, we iteratively increase the payoffs of the surviving strategies at each round of elimination (according to $(W_j^k)_{k \geq 1}$) by ϵ/n . Finally, let $\theta^n = \theta^{n,K}$.

I claim that a_i is not rationalizable in any complete information game θ^n , when n is sufficiently large. Let $(S_j^{k,n})_{k \geq 1}$ be the sets of player j actions surviving k rounds of iterated elimination of strictly dominated strategies in game θ^n . Define $N := \frac{\epsilon K}{\gamma}$, where

$$\gamma = \frac{1}{2} \min_i \min_{a_i \in A_i} |u_i(a_i, a_{-i}) - \max_{a'_i \neq a_i} u_i(a'_i, a_i)|.$$

I will show that $S_j^{k,n} = W_j^k$ for all k and every player j in the game θ^n where $n \geq N$.

Claim 2. Fix any game θ^n where $n \geq N$. If a_j is a best reply to any α_{-j} in game θ , then it is also a strict best reply to α_{-j} in game θ^n .

Proof. Define $a_j^* = \operatorname{argmax}_{a'_j \neq a_j} u(a'_j, a_{-j}, \theta^n)$. Then,

$$\begin{aligned}
& u_j(a_j, a_{-j}, \theta^n) - u_j(a_j^*, a_{-j}, \theta^n) \\
&= u_j(a_j, a_{-j}, \theta^n) - u_j(a_j, a_{-j}, \theta) \\
&\quad + u_j(a_j, a_{-j}, \theta) - u_j(a_j^*, a_{-j}, \theta) \\
&\quad + u_j(a_j^*, a_{-j}, \theta) - u_j(a_j^*, a_{-j}, \theta^n) \\
&\geq -(\epsilon K)/n + 2\gamma - (\epsilon K)/n \\
&= 2\gamma - (2\epsilon K)/N > 0
\end{aligned}$$

using in the last line that $n > N$. \square

Proceed by induction. Trivially, $S_j^{0,n} = W_j^0 = A_j$ for every j and n . Suppose $S_j^{k,n} = W_j^k$ for every player j , game $n \geq N$, and round $k \leq L$. Now consider any action $a_j \in S_j^{L,n}$. If a_j is a strict best response to some strategy $\alpha_{-j} \in \Delta(S_{-j}^{k,n})$, then $a_j \in W_j^{k+1}$, and by Claim 2, also $a_j \in S_j^{k+1,n}$ when $n > N$. Suppose a_j is weak best response to any $\alpha_{-j} \in \Delta(S_{-j}^{k,n})$. 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, clearly by construction $\theta^n \rightarrow \theta$. So $u \notin \operatorname{Int}(U_{a_i^*}^R)$, as desired. \square

B.8 Proof of Lemma 4

Consider the event in which there is a family $(R_i)_{i \in \mathcal{I}}$ such that $a_i \in BR_i(\nu)$ for some belief $\nu_i \in \Delta(\Theta(\mathbf{z}_n) \times R_{-i})$ for each player i . Suppose players have common certainty that each player i 's first-order belief is $\operatorname{marg}_\Theta \nu_i$. Then, action a_i is rationalizable for player i , so the condition in (12) implies $\underline{p}_n^R(i, a_i) \rightarrow 0$ as desired.

Conversely, suppose that a_i is rationalizable for player i of type $t_i \in T(\mathbf{z}_n)$. By definition of rationalizability, there exists a family of sets $R_j \subseteq A_j$ such that every action $a_j \in R_j$ is a best reply to a distribution $\pi \in \Delta(\Theta \times T_{-j} \times A_{-j})$ that satisfies $\operatorname{marg}_{\Theta \times T_{-j}} \pi = \kappa(t_j)$, and $\pi(a_{-j} \in R_{-j}[t_{-j}]) = 1$; moreover, since $t_i \in T(\mathbf{z}_n)$ by assumption, we have that $\operatorname{marg}_{T_{-i}} \pi$ assigns probability 1 to the event that each $t_j \in T(\mathbf{z}_n)$. Now let each $\nu_i = \operatorname{marg}_{\Theta \times A_{-j}} t_i$. Since the types in $T(\mathbf{z}_n)$ have common certainty in $\Theta(\mathbf{z}_n)$, we have that $\nu_i \in \Delta(\Theta(\mathbf{z}_n) \times R_{-i})$ as desired. Thus, the condition in (12) is necessary for rationalizability of a_i to be plausible under inference.

B.9 Proof of Proposition 5

This follows directly from Lemma 4, noting that expected payoffs are continuous in first-order beliefs (with respect to the Prokhorov metric) on a compact set.

C Supplementary Material to Section 7

C.1 Proof of Proposition 11

(a) To simplify notation, set $\delta := \delta_{\mathbf{a}}^{NE}$. Clearly, the expected payoff matrix $E_Q(\theta) \in \Theta_{\mathbf{a}}^{NE}$ for every distribution $Q \in B_\delta(P)$. Applying Lemma 1, if $\mu(\mathbf{z}_n) \in B_\delta(P)$ for every $\mu \in \mathcal{M}$, then the strategy profile $\sigma_{\mathbf{a}}$ (as defined in Appendix B) is a Bayesian Nash equilibrium. Write P^n for the induced measure over n -length sequences in \mathcal{Z}_n . Then,

$$\begin{aligned} p_n^{NE}(a^*) &\geq P^n(\{\mathbf{z}_n : \mu(\mathbf{z}_n) \in B_\delta(P) \forall \mu \in \mathcal{M}\}) \\ &= P^n\left(\left\{\mathbf{z}_n : \sup_{\mu \in \mathcal{M}} d(\mu(\mathbf{z}_n), P) \leq \delta\right\}\right) \\ &= 1 - P^n\left(\left\{\mathbf{z}_n : \sup_{\mu \in \mathcal{M}} d(\mu(\mathbf{z}_n), P) > \delta\right\}\right) \\ &\geq 1 - \frac{1}{\delta} \mathbb{E}_{P^n}\left(\sup_{\mu \in \mathcal{M}} d(\mu(\mathbf{z}_n), P)\right) \end{aligned}$$

using Markov's inequality in the final line.

(b) To simplify notation, set $\delta := \delta_{a_i^*}^R$. Since a_i is strictly rationalizable for player i , $\delta \geq 0$. Set $V = B_\delta(P)$ and apply Lemma 3 to conclude that a_i is rationalizable for any type with common certainty in the set

$$\{E_Q(\theta) : Q \in B_\delta(P)\}.$$

So if $\mu(\mathbf{z}_n) \in B_\delta(P)$ for every $\mu \in \mathcal{M}$, then the strategy a_i is rationalizable for all plausible types of player i . This allows us to construct the lower bound

$$\begin{aligned} p_n^R(a_i) &\geq P^n(\{\mathbf{z}_n : \mu(\mathbf{z}_n) \in B_\delta(P) \forall \mu \in \mathcal{M}\}) \\ &\geq 1 - \frac{1}{\delta} \mathbb{E}_{P^n}\left(\sup_{\mu \in \mathcal{M}} d(\mu(\mathbf{z}_n), P)\right) \end{aligned}$$

as in part (a).

D Supplementary Material to Section 8

D.1 Proof of Proposition 8

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

$$\Theta(\mathbf{z}_n) = \{\phi(\mathbf{x}^*) : \phi(\mathbf{x}_k) = y_k \quad \forall k = 1, \dots, n\}$$

and let $T_{\mathbf{z}_n}$ be the set of hierarchies of belief with common certainty in $\Theta(\mathbf{z}_n)$. First, I will show that ‘Join’ is rationalizable for all types in $T_{\mathbf{z}_n}$ if and only if $\Theta(\mathbf{z}_n) = \{1\}$. Suppose towards contradiction that $\Theta(\mathbf{z}_n) \neq \{1\}$ so that $0 \in \Theta(\mathbf{z}_n)$; then, the set $T_{\mathbf{z}_n}$ includes the type with common certainty in $\theta = 0$, but Join is not rationalizable for player 1 of this type.

In the other direction, suppose $\Theta(\mathbf{z}_n) = \{1\}$. 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 $\Theta(\mathbf{z}_n) = \{1\}$ if and only if every rectangular classification rule ϕ that exactly fits the data predicts $\phi(\mathbf{x}^*) = 1$. We can reduce this problem to whether the smallest hyper-rectangle that contains every observed vector \mathbf{x}_k also contains \mathbf{x}^* . Specifically, the probability $p_n^R(i, \text{Join})$ is equal to the probability that on every dimension k ,

$$\exists \text{ observations } (\mathbf{x}^i, 1) \text{ and } (\mathbf{x}^j, 1) \text{ such that } x_k^i < \mathbf{x}_k^* \text{ and } x_k^j > \mathbf{x}_k^*, \quad (24)$$

that is, a “successful” observation lies on either side of \mathbf{x}^* in dimension k .

If $k \in \{1, \dots, p^*\}$, then (26) is satisfied on dimension k only if some \mathbf{x}^i satisfying $x_k^i \in [-c', \mathbf{x}_k^*)$, and also some \mathbf{x}^j satisfying $x_k^j \in (\mathbf{x}_k^*, c']$, are 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 $k \in \{p^* + 1, \dots, \bar{p}\}$, then (26) is satisfied on dimension k only if some \mathbf{x}^i satisfying $x_k^i < \mathbf{x}_k^*$ is sampled, and additionally some \mathbf{x}^j satisfying $x_k^j > \mathbf{x}_k^*$ is sampled. The probability that this occurs is

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

Now, observe that realizations of characteristics are independent across dimensions. So the probability that (26) 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.

D.2 Extended Discussion of Footnote 25.

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

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

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 .

D.3 Proof of Proposition 9

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

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

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

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

Claim 3. For any vector $\mathbf{u} = (\mathbf{w} \quad \mathbf{z})$ where $\mathbf{w} \in \mathbb{R}^{1 \times p'}$ and $\mathbf{z} \in \mathbb{R}$,

$$\text{Var}(\mathbf{w}\hat{\beta}) \geq \text{Var}(\mathbf{u}\tilde{\beta}).^{32} \tag{25}$$

Proof. Write X for the $n \times p$ matrix stacking row vectors $(\mathbf{x}_1^i, \dots, \mathbf{x}_{p'}^i)$, where $i \in \{1, \dots, n\}$, 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,

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

³²Iosif Pinelis contributed to this proof.

Observe that

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

and similar manipulations yield that

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

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 (27) holds if and only if the matrix

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

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 $(\mathbf{c} \ \mathbf{d})$. Algebraic manipulations yield

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

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

References

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2015): “Fragility of Asymptotic Agreement under Bayesian Learning,” *Theoretical Economics*.
- AL-NAJJAR, N. (2009): “Decisionmakers as Statisticians: Diversity, Ambiguity, and Learning,” *Econometrica*.
- AL-NAJJAR, N. AND M. PAI (2014): “Coarse Decision Making and Overfitting,” *Journal of Economic Theory*.
- AUMANN, R. J. (1976): “Agreeing to Disagree,” *The Annals of Statistics*.
- BATTIGALLI, P. AND M. SINISCHALCHI (2003): “Rationalization and Incomplete Information,” *Advances in Theoretical Economics*.
- BILLOT, A., I. GILBOA, D. SAMET, AND D. SCHMEIDLER (2005): “Probabilities as Similarity-Weighted Frequencies,” *Econometrica*.
- BRANDENBURGER, A. AND E. DEKEL (1993): “Hierarchies of Belief and Common Knowledge,” *Journal of Economic Theory*.
- CARLIN, B. I., S. KOGAN, AND R. LOWERY (2013): “Trading Complex Assets,” *The Journal of Finance*.
- CARLIN, B. I., F. A. LONGSTAFF, AND K. MATOBA (2014): “Disagreement and asset prices,” *Journal of Financial Economics*.
- CARLSSON, H. AND E. VAN DAMME (1993): “Global Games and Equilibrium Selection,” *Econometrica*, 61, 989–1018.
- CHEN, Y.-C., A. DI TILLIO, E. FAINGOLD, AND S. XIONG (2010): “Uniform topologies on types,” *Theoretical Economics*.
- (2017): “Characterizing the Strategic Impact of Misspecified Beliefs,” *Review of Economic Studies*.
- CHEN, Y.-C. AND S. TAKAHASHI (2017): “On Robust Selection and Robust Rationalizability,” Working Paper.
- CRIPPS, M., J. ELY, G. MAILATH, AND L. SAMUELSON (2008): “Common Learning,” *Econometrica*.
- DEKEL, E., D. FUDENBERG, AND D. LEVINE (2004): “Learning to Play Bayesian Games,” *Games and Economic Behavior*.

- DEKEL, E., D. FUDENBERG, AND S. MORRIS (2006): “Topologies on Types,” *Theoretical Economics*.
- (2007): “Interim Correlated Rationalizability,” *Theoretical Economics*.
- EINMAHL, U. AND D. M. MASON (2005): “Uniform in Bandwidth Consistency of Kernel-Type Function Estimators,” *The Annals of Statistics*.
- ELY, J. AND M. PESKI (2011): “Critical Types,” *Review of Economic Studies*.
- ESPONDA, I. (2013): “Rationalizable conjectural equilibrium: A framework for robust predictions,” *Theoretical Economics*.
- GAYER, G., I. GILBOA, AND O. LIEBERMAN (2007): “Rule-Based and Case-Based Reasoning in Housing Prices,” *The B.E. Journal of Theoretical Economics*.
- GEANAKOPOLOS, J. AND H. POLEMARCHAKIS (1982): “We Can’t Disagree Forever,” *Journal of Economic Theory*.
- GILBOA, I., O. LIEBERMAN, AND D. SCHMEIDLER (2006): “Empirical Similarity,” *Review of Economics and Statistics*.
- GILBOA, I., A. POSTLEWAITE, AND D. SCHMEIDLER (2008): “Probability and Uncertainty in Economic Modeling,” *Journal of Economic Perspectives*.
- GILBOA, I., L. SAMUELSON, AND D. SCHMEIDLER (2013): “Dynamics of inductive inference in a unified framework,” *Journal of Economic Theory*.
- GILBOA, I. AND D. SCHMEIDLER (1995): “Case-Based Decision Theory,” *The Quarterly Journal of Economics*.
- (2003): “Inductive Inference: An Axiomatic Approach,” *Econometrica*.
- HASTIE, T., R. TIBSHIRANI, AND J. FRIEDMAN (2009): *The Elements of Statistical Learning*, Springer.
- KAJII, A. AND S. MORRIS (1997): “The Robustness of Equilibria to Incomplete Information,” *Econometrica*.
- KANDEL, E. AND N. PEARSON (1995): “Differential Interpretation of Information and Trade in Speculative Markets,” *Journal of Political Economy*.
- MANKIW, G., R. REIS, AND J. WOLFERS (2004): “Disagreement about inflation expectations,” *NBER Macroeconomics Annual 2003*.

- MARLON, J., A. LEISEROWITZ, AND G. FEINBERG (2013): “Scientific and Public Perspectives on Climate Change,” Tech. rep., Yale Project on Climate Change Communication.
- MERTENS, J.-F. AND S. ZAMIR (1985): “Formulation of Bayesian Analysis for Games with Incomplete Information,” *International Journal of Game Theory*.
- MILGROM, P. AND N. STOKEY (1982): “Information, Trade, and Common Knowledge,” *Journal of Economic Theory*.
- MONDERER, D. AND D. SAMET (1989): “Approximating Common Knowledge with Common Beliefs,” *Games and Economic Behavior*.
- MORRIS, S., S. TAKAHASHI, AND O. TERCIEUX (2012): “Robust Rationalizability Under Almost Common Certainty of Payoffs,” *The Japanese Economic Review*.
- OLEA, J. M., P. ORTOLEVA, M. PAI, AND A. PRAT (2017): “Competing Models,” Working Paper.
- RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior Under “Almost Common Knowledge”,” *American Economic Review*, 79.
- SPIEGLER, R. (2016): “Bayesian Networks and Boundedly Rational Expectations,” *Quarterly Journal of Economics*.
- STEINER, J. AND C. STEWART (2008): “Contagion through Learning,” *Theoretical Economics*.
- TAKAHASHI, S. (2017): “Strategic Implications of Almost Common Certainty of Payoffs,” Working Paper.
- WEINSTEIN, J. AND M. YILDIZ (2007): “A Structure Theorem for Rationalizability with Application to Robust Prediction of Refinements,” *Econometrica*.
- WIEGEL, D. (2009): “How many Southern whites believe Obama was born in America?” *Washington Independent*.
- YU, J. (2011): “Disagreement and return predictability of stock portfolios,” *Journal of Financial Economics*.