

Games of Incomplete Information Played By Statisticians

Annie Liang*

June 27, 2021

Abstract

Players are statistical learners who form beliefs about payoff-relevant parameters based on data. They may update from the same data in different ways, but have common knowledge that all beliefs are formed by updating to this data from a set of “reasonable” learning rules. The strategic predictions that are compatible with this belief restriction depend on the realization of the data, and I define the *plausibility* of a strategic prediction based on the typicality of the data sets that support it. The main results characterize how the plausibility of a given strategic prediction varies depending properties of the learning problem, including the quantity of data that players see. The approach generates new predictions, e.g. that speculative trade is more plausible when players see sparse data and when the learning problem is high-dimensional.

JEL Codes: C70, D83.

1 Introduction

Decision makers form beliefs about many economic unknowns based on existing data. For example, individuals use historical data to evaluate the quality of a new good or to predict the long-run implication of an economic shock. Although it is common to model beliefs

*Department of Economics and Department of Computer Science, Northwestern University. Email: annie.liang@northwestern.edu. I am especially grateful to Drew Fudenberg for his guidance on this paper. This paper also benefitted from useful comments and suggestions by Jetlir Duraj, Siddharth George, Ben Golub, Jerry Green, Philippe Jehiel, Scott Kominers, David Laibson, Jonathan Libgober, Erik Madsen, Stephen Morris, Sendhil Mullainathan, Mariann Ollár, Harry Pei, Andrei Shleifer, Dov Samet, Tomasz Strzalecki, Satoru Takahashi, Anton Tsoy, and Muhamet Yildiz.

in incomplete information games as formed by Bayesian updating to information from the same prior belief, economic actors often use different procedures to learn from data. A recent literature highlights the diversity of these learning procedures, modeling agents as statistical learners who run regressions with different sets of covariates (Mailath and Samuelson, 2020; Olea, Ortoleva, Pai, and Prat, 2020), learn different summary statistics from the data (Jehiel, 2018; Salant and Cherry, 2020), identify different “similarity classes” (Gilboa and Schmeidler, 1995; Jehiel, 2005), or Bayes-update from different priors (Alonso and Câmara, 2016; Sethi and Yildiz, 2016). These approaches all depart from the rigidity of a common prior, whose strong restrictions and conceptual limitations are well summarized in Morris (1995).

The goal of this paper is to propose a general framework in which agents may use different procedures for forming beliefs given data, and to explore which predictions about strategic behaviors are more and less plausible in the presence of such diversity. Use of different models can lead to disagreement across agents, but is also compatible with agents holding similar beliefs. Intuitively, we might expect that diversity in models leads to disagreement when these models are trained on less data, or when the learning problem is so complex that the available historical data is insufficient to coordinate the forecasts of these models. For example, there is substantial public disagreement over when the next recession will be and how serious the risks of a COVID vaccine are, but not over whether December will be colder in Chicago or Miami. This paper offers a framework that can distinguish between these settings, and thus makes new predictions about how strategic behaviors vary depending on the quantity of available data and properties of the learning problem.

In the proposed model, players form beliefs about payoff-relevant parameters based on common data, but potentially disagree on how to interpret that data. Interpretation of data is formalized as a *learning rule*, which is any function that maps data (a sequence of signals) into a belief distribution over payoff-relevant parameters (a first-order belief). I assume that there is a set of learning rules that is common knowledge across players, so that the *reasonable beliefs* given observed data are those that are consistent with the application of some learning rule from this set to the data.

While players may hold different beliefs about an unknown parameter (and know that this is the case), they have common certainty in the event that all players hold a reasonable belief. Such belief restrictions (which follow Battigalli and Siniscalchi (2003)) discipline the higher-order beliefs that players may hold in a way that is intermediate between imposing a common

prior and allowing players to hold any beliefs at all. See [Ollár and Penta \(2017\)](#) and [Ollár and Penta \(2020\)](#) for recent applications of such restrictions to obtain full implementation in mechanism design problems. The key innovation in the present paper is that the set of reasonable beliefs is endogenous to a learning process, so the strategic predictions that are consistent with this belief restriction vary depending on the properties of the learning environment.

Fixing an underlying data-generation process, I measure the “plausibility” of a strategic prediction based on the probability with which data is realized that supports it. Consider, for example, an action that is only rationalizable when players disagree about the value of a parameter, e.g. speculative trade of a common-value good whose quality is unknown. If most realizations of the data lead to similar beliefs under the different learning rules, then we should not expect to see trade occur. On the other hand, if many data sets lead to substantial disagreement, then trade is a reasonable prediction in this learning environment. Formally, I define a *plausibility interval* for each action: The upper bound of the plausibility interval is the probability that the action is rationalizable given *some* belief satisfying the belief restriction, and the lower bound is the probability that the prediction holds for *all* beliefs satisfying the restriction.¹ If both of these probabilities are equal to one, then the analyst has maximal certainty that the action is rationalizable, and if they are both zero, then he has maximal certainty that it is not. In the intermediate cases, there is uncertainty about whether the action is rationalizable, and the plausibility interval presents a way to quantify the extent of that uncertainty.

In Section 3, I show that the proposed approach generates novel comparative statics in two simple, but classic settings—a trade game and a coordination game. In the trade game, I suppose that a buyer and seller form beliefs about the value of a good based on a common data set. The learning environment is parametrized by two parameters: the size of the common data set and the dimensionality of the learning problem. I show that trade is a plausible outcome when players see sparse data and the learning problem is high-dimensional. In contrast, coordination on an action profile with high (but unknown) payoffs is facilitated by larger quantities of data and a simpler learning problem.

The applications in Section 3 have the property that the plausibility interval for the

¹These probabilities reflect a maximally stringent and maximally lenient view towards what constitutes a reasonable prediction. Stronger assumptions on beliefs, beyond those that I have imposed here, would narrow this interval.

prediction of interest can be exactly characterized, but this will not always be possible. The main results in Section 4 demonstrate properties of the plausibility interval that hold more generally. In Section 4.1, I ask what the plausibility interval looks like when players see a large quantity of common data, and whether predictions given a large but finite quantity of data resemble those that we would make with infinite data. I show that if the set of learning rules is too rich, then predictions in this infinite data game can be qualitatively different from those in finite data games with arbitrarily large quantities of data (even if all learning rules eventually converge to the same belief). These are settings where, for any fixed quantity of data, each agent knows that there is another reasonable learning rule which leads to a very different belief from his own. Disagreement is a fundamental property of these settings, and reduction to a common prior leads to misleading predictions.

If, however, the set of learning rules satisfies a uniform convergence property that I describe, then the following statements hold: If an action is strictly rationalizable at the limit, then the analyst’s plausibility interval must converge to $\{1\}$ as the quantity of data gets large, and if an action is *not* rationalizable in the limit, the analyst’s plausibility interval must converge to $\{0\}$. (The intermediate case, in which actions are rationalizable but not strictly rationalizable, is more subtle—see Section 4.1 for details.) Thus we can use the infinite data limiting game as an approximation for the actual game when the number of data points is sufficiently large.

Next, I suppose that players commonly observe a finite, and potentially small, quantity of data, and provide bounds for the plausibility interval in this case. These bounds depend on properties of the learning environment—specifically, how much data the players commonly see, and how fast the different learning rules jointly recover the payoff-relevant parameter—as well as on a cardinal measure of how “strict” the solution is at the limit. They allow us to make statements about plausibility away from the limit of infinite data: A prediction that holds at the limit may fail to hold given small quantities of data if that prediction relies on precise agreement in beliefs across players, and if the learning rules that players use recover the unknown parameter slowly. A prediction that does not hold at the limit may still be a plausible prediction of play in complex learning environments with sparse data. While it is intuitive that the likelihood of different strategic outcomes may depend on the learning environment, these bounds provide (to the best of my knowledge) a first formalization for this reasoning.

1.1 Related Literature

This paper builds on a literature on the role of the common prior assumption in economic theory. (See [Morris \(1995\)](#) for a survey of key conceptual points.) Here I focus on an argument that even if learning does produce common priors in the long run, this does not imply that we should see common priors given a finite quantity of data, especially if that data is complex and hard to interpret. Rather than taking the limiting common prior as one that is already reached, I ask what predictions we can make while data is still being accumulated. The plausibility intervals introduced in this paper provide a quantitative account of when a prediction that is implied by a (limiting) common prior also hold when players’ beliefs are informed by learning from finite (and potentially small) data sets.²

This paper also contributes to a literature on the robustness of strategic predictions to the specification of player beliefs ([Rubinstein, 1989](#); [Dekel et al., 2006](#); [Weinstein and Yildiz, 2007](#); [Chen et al., 2010](#); [Ely and Peski, 2011](#); [Penta, 2012](#)) and equilibrium selection in incomplete information games ([Carlsson and van Damme, 1993](#); [Kajii and Morris, 1997](#); [Morris, Takahashi, and Tercieux, 2012](#)). The permitted types in this paper converge in the *uniform-weak topology*, as proposed and characterized in [Chen, di Tillio, Faingold, and Xiong \(2010\)](#) and [Chen, di Tillio, Faingold, and Xiong \(2017\)](#), and I use results about this topology to prove several of my main results.

Conceptually, the goals of the present paper differ from the previous literature in several respects: First, the focus here is not on equilibrium selection—i.e., choosing one equilibrium from a set of many—but rather on providing a metric for the plausibility of a fixed prediction of interest. Second, in contrast to the many binary or “qualitative” notions of robustness that have been proposed, this paper delivers a quantitative metric, where predictions can be “more or less” robust. Third, while the literature has primarily considered robustness to *perturbations* of beliefs, I am interested here also in predictions that we may make for beliefs that are far from the limiting beliefs. To discipline these beliefs, I endogenize the type space using a statistical learning foundation for belief formation. This aspect of the paper—combining learning foundations with game theoretic implications—connects to papers such as [Dekel et al. \(2004\)](#), [Esponda \(2013\)](#), and [Steiner and Stewart \(2008\)](#), among others.³

²Other reasons that the common prior is tenuous include that the data itself may lead to incomplete learning if it is endogenously acquired, and that convergence of individual beliefs need not imply convergence in beliefs about beliefs ([Cripps et al., 2008](#); [Acemoglu et al., 2015](#)).

³[Brandenburger et al. \(2008\)](#) and [Battigalli and Prestipino \(2011\)](#) also motivate small type structures as

Finally, the modeling of agents as “statisticians” or “machine learners” relates to a growing literature in decision theory (Gilboa and Schmeidler, 2003; Gayer et al., 2007; Al-Najjar, 2009; Al-Najjar and Pai, 2014) and game theory (Jehiel, 2005; Spiegel, 2016; Eliaz and Spiegel, 2020; Olea et al., 2020; Salant and Cherry, 2020; Haghtalab et al., 2020). Of these papers, Steiner and Stewart (2008) and Salant and Cherry (2020) are closest: Steiner and Stewart (2008) characterizes the limiting equilibria of a sequence of games in which players infer payoffs from related games, and Salant and Cherry (2020) models players as statisticians who form beliefs about the action distribution based on statistical inference from a sample of observed players. Different from these papers, the goal of the present paper is to provide a metric of plausibility rather than a new solution concept.

2 Approach

2.1 Preliminaries

Basic Game. There is a finite set \mathcal{I} of players and a finite set of actions A_i for each player $i \in \mathcal{I}$. The set of action profiles is $A = \prod_{i \in \mathcal{I}} A_i$, and the set of possible games is identified with $U := \mathbb{R}^{|\mathcal{I}| \times |A|}$. Agents have beliefs over a set of payoff-relevant parameters Θ , which is a compact and convex subset of finite-dimensional Euclidean space. It is possible to take Θ to be a subset of U , so that each θ is itself a game, or to define beliefs over a lower-dimensional set of payoff-relevant parameters. In either case, the parameters in Θ are assumed to be related to payoffs by a bounded and Lipschitz continuous embedding $g : \Theta \rightarrow U$ (assuming the sup-norm on both spaces).⁴

Beliefs. For each player i , let $X_i^0 = \Theta$, $X_i^1 = X_i^0 \times \prod_{j \neq i} \Delta(X_j^0)$, \dots , $X_i^n = X_i^{n-1} \times \prod_{j \neq i} \Delta(X_j^{n-1})$, etc., so that for each $k \geq 1$, the set $\Delta(X_i^{k-1})$ constitutes the possible k -th order beliefs of player i . Define $T_i^0 = \prod_{n=0}^{\infty} \Delta(X_i^n)$. An element $(t_i^1, t_i^2, \dots) \in T_i^0$ 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 referred to simply as a *belief* or *type*.⁵

emerging from learning, although they do not explicitly model a dynamic learning process.

⁴The map g can be interpreted as capturing the known information about the structure of payoffs.

⁵Types are sometimes modeled as encompassing all uncertainty in the game. In the present paper, types describe players’ structural uncertainty over payoffs, but not their strategic uncertainty over opponent actions.

There is a subset of types T_i^* (that satisfy the property of *coherency*⁶ and common knowledge of coherency) 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 (Mertens and Zamir, 1985; Brandenburger and Dekel, 1993).⁷ The tuple $(T_i^*, \kappa_i^*)_{i \in \mathcal{I}}$ is the *universal type space*. Subsequently I will develop smaller type spaces $(T_i, \kappa_i)_{i \in \mathcal{I}}$ where each $T_i \subseteq T_i^*$ and $\kappa_i : T_i \rightarrow \kappa_i^*(T_i)$, is the restriction of κ_i^* to T_i .

Solution Concept. For every player i and type $t_i \in T_i^*$, 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 is a best reply to some $\pi \in \Delta(\Theta \times T_{-i}^* \times A_{-i})$ satisfying (1) $\text{marg}_{\Theta \times T_{-i}^*} \pi = \kappa_i^*(t_i)$ and (2) $\text{marg}_{A_{-i} \times T_{-i}^*} \pi(\{(a_{-i}, t_{-i}) \mid a_{-i} \in S_{-i}^{k-1}[t_{-i}]\}) = 1$, 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 type t_i 's 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 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 type t_i , or (henceforth) simply *rationalizable* (Dekel et al., 2007; Weinstein and Yildiz, 2017).⁸ Say that a_i is *strictly rationalizable* for type t_i if the best reply conditions above are strengthened to strict best replies.

2.2 Restriction on Beliefs

The proposed approach endogenizes the type space based on two new primitives: a *data-generating process*, and a *set of rules* for how to extrapolate beliefs from realized data.

Formally, let $(Z_t)_{t \in \mathbb{Z}_+}$ be a stochastic process where the random variables Z_t take value in a common set \mathcal{Z} , and the typical sample path is denoted by $\mathbf{z} = (z_1, z_2, \dots)$. The data-generating process is a measure P over the set \mathcal{Z}^∞ of all (infinite) sample paths.⁹ Let P^n denote the induced measure on the first n variables. A *data set* \mathbf{z}_n of size n is the restriction of

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

⁷The notation T_{-i}^* is used throughout the paper to denote the set of profiles of opponent types, $\prod_{j \neq i} T_j^*$.

⁸The set of interim correlated rationalizable actions for a given type t_i is independent of the ambient type space, as shown in Dekel et al. (2007). Thus, the definition given here for the rationalizable actions will apply also for the smaller type spaces $(T_i, \kappa_i)_{i \in \mathcal{I}}$ that I work with in the main text.

⁹In the standard Bayesian framework, the players and analyst would share a common prior belief $\pi \in \Delta(\Theta)$ along with knowledge of a family of data-generating processes $(P_\theta)_{\theta \in \Theta}$. The measure P , as defined here, corresponds to the data-generating process given the true (objective) state. It is possible to extend the current approach so that P is the ex-ante belief over sample paths induced by π and $(P_\theta)_{\theta \in \Theta}$, as discussed in Section 5.

\mathbf{z} to its first n coordinates, and \mathcal{Z}^n is the set of all length- n data sets. I use $Z^n = (Z_1, \dots, Z_n)$ to denote the random initial sequence of length n . Throughout, I assume that this data is publicly observed to isolate the effect of different learning procedures (rather than the more traditional channel of different information).¹⁰

A *learning rule* is any map from data sets into first-order beliefs:

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

There is a set of learning rules \mathcal{M} that are reasonable for the setting, which is commonly known by the players. For example, \mathcal{M} might be any of the following:

Bayesian Updating from a Set of Priors. The set of learning rules \mathcal{M} is indexed to a set of prior beliefs $\Pi \subseteq \Delta(\Theta \times \mathcal{Z}^\infty)$. For any prior $\pi \in \Pi$, the learning rule μ_π maps each realization \mathbf{z}_n into the (marginal) posterior belief on Θ , updating from the prior π .¹¹

A Set of Sample Statistics. Learning rules in \mathcal{M} map the data to different point-estimates for the payoff-relevant parameter. For example, \mathcal{M} might consist of the two learning rules μ_{mean} and μ_{median} , where for any data set \mathbf{z}_n , $\mu_{\text{mean}}(\mathbf{z}_n)$ is a point-mass belief on the mean realization in \mathbf{z}_n (as in [Jehiel \(2018\)](#)), and $\mu_{\text{median}}(\mathbf{z}_n)$ is a point-mass belief on the median realization.

Linear Regression with Different Covariate Sets. Define $\mathcal{Z} = X \times Y$ where $X \subseteq \mathbb{R}^p$ is a set of attribute vectors and Y is a set of outcomes. The payoff-relevant parameter θ is the unknown value of the outcome at a known attribute vector $x^* \in X$. Learning rules in \mathcal{M} are indexed to different subsets of attributes (as in [Olea et al. \(2020\)](#)). For each learning rule $\mu \in \mathcal{M}$, let $I_\mu \subseteq \{1, \dots, p\}$ be the index set of attributes associated with μ , and for any $x \in X$, let $x_\mu = (x_i)_{i \in I_\mu}$ describe the coordinates of x at the indices in I_μ . The linear function of the attributes in I_μ that best fits data $\mathbf{z}_n = \{(x^k, y^k)\}_{k=1}^n$ is $\hat{f}_\mu^{OLS}[\mathbf{z}_n](x) = \beta_\mu^{OLS} \cdot x_\mu$, where

¹⁰These two approaches lead to distinct predictions, since common knowledge disagreement cannot arise from different information under a common prior ([Aumann, 1976](#)), while it is feasible in the proposed framework. The current framework can be further extended to allow for different information in addition to different learning rules, but I do not pursue this here.

¹¹In general, making the mapping deterministic may require choosing a conditional probability if there are multiple ones consistent with Bayes' rule; here and elsewhere in the paper, implicitly assume that the updating rule specifies such a choice when Bayesian rules are mentioned.

$\beta_\mu^{OLS} = \operatorname{argmin}_{\beta \in \mathbb{R}^{|\mu|}} \frac{1}{n} \sum_{i=1}^n (\beta \cdot x_\mu^k - y^k)^2$. Each learning rule μ maps \mathbf{z}_n into a point-mass belief on the prediction $\hat{f}_\mu^{OLS}[\mathbf{z}_n](x^*)$.

For finite n and realized data \mathbf{z}_n , the learning rules in \mathcal{M} induce a set of distributions

$$\mathcal{B}(\mathbf{z}_n) = \{\mu(\mathbf{z}_n) : \mu \in \mathcal{M}\} \subseteq \Delta(\Theta). \quad (1)$$

The key restriction on higher-order beliefs is that players have common certainty in the event that all players have first-order beliefs in $\mathcal{B}(\mathbf{z}_n)$. That is, while players may hold different beliefs over the state, they believe that (all other players believe that. . .) each player's belief belongs to $\mathcal{B}(\mathbf{z}_n)$. Formally, for any set $\mathcal{B} \subseteq \Delta(\Theta)$, and for any player i , define

$$B_i^{1,1}(\mathcal{B}) := \{t_i \in T_i^* : \operatorname{marg}_\Theta \kappa_i^*(t_i) \in \mathcal{B}\}$$

to be the set of player i types whose marginal beliefs on Θ belong to the set \mathcal{B} . For each $k > 1$, and again for each player i , recursively define

$$B_i^{k,1}(\mathcal{B}) = \left\{ t_i \in T_i^* : \kappa_i^*(t_i) \left(\Theta \times \prod_{j \neq i} B_j^{k-1,1}(\mathcal{B}) \right) = 1 \right\}.$$

Then $T_i^{\mathcal{B}} = \bigcap_{k \geq 1} B_i^{k,1}(\mathcal{B})$ is the set of player i types that have common certainty in the event that all players' first-order beliefs belong to \mathcal{B} .¹²

Definition 1. For every \mathbf{z}_n , the induced type space is $\left(T_i^{\mathcal{B}(\mathbf{z}_n)}, \kappa_i^{\mathcal{B}(\mathbf{z}_n)} \right)_{i \in \mathcal{I}}$, where $\kappa_i^{\mathcal{B}(\mathbf{z}_n)} : T_i^{\mathcal{B}(\mathbf{z}_n)} \rightarrow \kappa_i^* \left(T_i^{\mathcal{B}(\mathbf{z}_n)} \right)$ is the restriction of κ_i^* to $T_i^{\mathcal{B}(\mathbf{z}_n)}$.¹³ The type t_i is **permitted** for player i given data \mathbf{z}_n if $t_i \in T_i^{\mathcal{B}(\mathbf{z}_n)}$.

This type space includes all type profiles where each player has common certainty in the event that all players have first-order beliefs in $\mathcal{B}(\mathbf{z}_n)$. In the special case in which

¹²Having defined a type space, we are in a usual Bayesian framework in which players possess a hierarchy of beliefs and maximize expected utility with respect to these beliefs. The novelty of the model is the flexible definition of a set of reasonable first-order beliefs $\mathcal{B}(\mathbf{z}_n)$ given the data \mathbf{z}_n , which includes beliefs formed by non-Bayesian learning procedures. I take this perspective because there is no uncontroversial specification of beliefs based on finite data—frequentist procedures are generally disciplined by their asymptotic properties, and Bayesian inference requires an ad-hoc specification of the prior.

¹³It is straightforward to show that the type sets $T_i^{\mathcal{B}(\mathbf{z}_n)}$ are *belief-closed*; that is, $\kappa_i^{\mathcal{B}(\mathbf{z}_n)}(t_i) \left(\Theta \times T_{-i}^{\mathcal{B}(\mathbf{z}_n)} \right) = 1$ for every $t_i \in T_i^{\mathcal{B}(\mathbf{z}_n)}$.

\mathcal{M} consists of a singleton Bayesian rule, then we return the common prior assumption. In general, some permitted types will be inconsistent with any common prior, e.g. types with common knowledge disagreement.

Remark 1. The assumption that players assign common *certainty* is not crucial. Similar results hold if the players have common p -belief (Monderer and Samet, 1989) in the set for large p . See Section 5.

Remark 2. Restrictions are placed only on the beliefs that players hold on Θ , and not on how they came to form those beliefs. For example, each player may be associated with a fixed learning rule from \mathcal{M} , or randomize over learning rules in \mathcal{M} with a fixed distribution, or (deterministically) use certain learning rules for data sets with certain properties. Indeed, the proposed approach does not require defining a set of learning rules \mathcal{M} , as we can alternatively take as the primitive a map from each possible data set \mathbf{z}_n into a set of beliefs $\mathcal{B}(\mathbf{z}_n)$ (see the subsequent Section 3.2 for such an example).

Remark 3. This framework implicitly assumes that the analyst’s model of the world includes a description of the data-generating process P and the set of learning rules \mathcal{M} . My view is that knowledge of these learning primitives is less demanding than knowledge of players’ subjective beliefs. One reason for this is that data, and the statistical procedures with which agents draw inferences from that data, are “objective” and potentially directly observable,¹⁴ while beliefs are not. Nevertheless, we should prefer predictions that do not rely too sensitively on details of these specifications—for example, I conjecture that the comparative statics in the subsequent Claim 1 holds for any family of classification rules of finite VC dimension.

2.3 Plausibility Interval

Fixing a set of beliefs $\mathcal{B} \subseteq \Delta(\Theta)$ and corresponding interim type space $(T_i^{\mathcal{B}}, \kappa_i^{\mathcal{B}})_{i \in \mathcal{I}}$, say that action a_i is *completely \mathcal{B} -rationalizable* if it is rationalizable for player i of *any* type $t_i \in T_i^{\mathcal{B}}$, and it is *partially \mathcal{B} -rationalizable* if it is rationalizable for player i of *some* type

¹⁴Agents predicting the future spread of a disease may use a known set of epidemiology models for prediction, and agents evaluating the value of a property may run different regressions on the same set of public features of the home.

$t_i \in T_i^{\mathcal{B}}$.¹⁵

Definition 2. For every $n \in \mathbb{Z}_+$, define $\underline{p}^n(a_i)$ to be the probability (over possible datasets \mathbf{z}_n) that action a_i is rationalizable for every type in $T_i^{\mathcal{B}(\mathbf{z}_n)}$; that is,

$$\underline{p}^n(a_i) = P^n(\{\mathbf{z}_n : a_i \text{ is completely } \mathcal{B}(\mathbf{z}_n)\text{-rationalizable}\}). \quad (2)$$

Define $\bar{p}^n(a_i)$ to be the probability (over possible datasets \mathbf{z}_n) that action a_i is rationalizable for some type $t_i \in T_i^{\mathcal{B}(\mathbf{z}_n)}$; that is,

$$\bar{p}^n(a_i) = P^n(\{\mathbf{z}_n : a_i \text{ is partially } \mathcal{B}(\mathbf{z}_n)\text{-rationalizable}\}). \quad (3)$$

The **plausibility interval** for rationalizability of a_i given n observations is $[\underline{p}^n(a_i), \bar{p}^n(a_i)]$.

The probabilities $\underline{p}^n(a_i)$ and $\bar{p}^n(a_i)$ are the ex-ante probabilities that an action is completely or partially $\mathcal{B}(Z_n)$ -rationalizable (prior to the realization of the data), and correspond to maximally stringent and maximally lenient views for whether a_i constitutes a “reasonable” prediction for the type space $(T_i^{\mathcal{B}(Z_n)}, \kappa_i^{\mathcal{B}(Z_n)})_{i \in \mathcal{I}}$. The larger $\underline{p}^n(a_i)$ and $\bar{p}^n(a_i)$ are, the more confident an analyst should be in predicting that a_i is rationalizable. At extremes: If $\bar{p}^n(a_i) = \underline{p}^n(a_i) = 1$, then given observation of n random samples, action a_i is guaranteed to be rationalizable for player i (for all permitted types). If $\bar{p}^n(a_i) = \underline{p}^n(a_i) = 0$, then action a_i is guaranteed to not be rationalizable for player i (for any permitted types). In the intermediate cases, if $0 < \underline{p}^n(a_i) = \bar{p}^n(a_i) < 1$, then rationalizability of the action a_i depends on the realization of the data, and if $\underline{p}^n(a_i) < \bar{p}^n(a_i)$, then rationalizability of a_i varies across permitted types.¹⁶ A similar approach can be used to define plausibility intervals for prediction that an action profile is part of a Bayesian Nash equilibria—see Section 5.

Observation 1. For every player i and action $a_i \in A_i$:

$$(a) \quad \underline{p}^n(a_i) \leq \bar{p}^n(a_i) \text{ for every } n \in \mathbb{Z}_+.$$

¹⁵Given the restriction in Definition 1, the partially \mathcal{B} -rationalizable strategies are the Δ -rationalizable strategies of Battigalli and Siniscalchi (2003), where $\Delta = (\Delta_i)_{i \in \mathcal{I}}$ and each $\Delta_i = \{\nu \in \Delta(\Theta \times T_{-i} \times A_{-i}) \mid \text{marg}_{\Theta} \mu \in \mathcal{B}\}$ encodes the belief restriction that first-order beliefs belong to \mathcal{B} . The concept of complete \mathcal{B} -rationalizability can be interpreted as a “robust” version of Δ -rationalizability.

¹⁶I do not comment on selection methods that may be used to further refine the set of permitted types. For the purposes of this paper, a large plausibility interval means that the prediction depends on details of beliefs that the analyst may not know.

(b) If \mathcal{M} consists of a single learning rule, then $\underline{p}^n(a_i) = \bar{p}^n(a_i)$ for every $n \in \mathbb{Z}_+$.

Note that in the special case in which agents have a common prior, then the common prior determines a distribution over \mathbf{z}_n , and hence over possible interim games. For any player i and action a_i , the probabilities $\underline{p}^n(a_i) = \bar{p}^n(a_i)$, and are equal to 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. This approach is used for example in [Kajii and Morris \(1997\)](#) (if we re-interpret the histories \mathbf{z}_n as the states), where an incomplete information game is “close” to a complete information game if the payoffs of the complete information game occur with high probability under the prior. So the probabilities $\underline{p}^n(a_i)$ and $\bar{p}^n(a_i)$ are a natural generalization of this standard measure of the typicality of a strategic prediction, when a common prior does not exist.

3 Example Applications

I first illustrate the proposed approach in two stylized games, where explicit characterization of the plausibility interval is possible for an action of interest, and yields new predictions about strategic behavior.

3.1 Trade in High-Dimensional Learning Problems

A Seller owns a good of unknown value $v \in \{0, 1\}$. He can *enter* a market at cost c , or *exit* and keep the good. Entering leads to a simultaneous interaction with a Buyer, where the Seller chooses whether to sell the good at a (pre-set) posted price p , and the Buyer chooses whether to purchase the good at that price. The cost c and price p satisfy $0 < c < p < 1$, so the Seller prefers to sell at the low value and prefers to keep the good at the high value. The game and players’ payoffs are described in [Figure 1](#). If players share a common prior about v , then entering is not rationalizable for the Seller in this game, so trade will not occur.¹⁷

Data. Players observe a public data set of past goods and their valuations, where each good is described by a $1 \times m$ vector of attributes $x \in X := [-1, 1]^m$. There is a deterministic

¹⁷If trade does not occur subsequently, then the Seller receives $v - c$ from entering but $v > v - c$ from exiting. Thus, entering can be rationalized only if trade subsequently occurs. But trade can occur only if the Buyer believes that $\mathbb{E}(v) \geq p$ while the Seller believes that $\mathbb{E}(v) \leq p$, implying $\mathbb{E}(v) = p$ under their shared belief. The Seller can improve on his expected payoff of $p - c$ by choosing to exit.

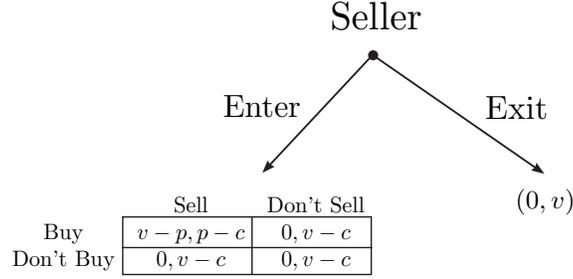


Figure 1: Description of Game

relationship between each good's attributes and its value, where those goods (and only those goods) whose attributes belong to a certain hyper-rectangle R have a high value.¹⁸ Let $f_R = \mathbb{1}(x \in R)$ be the function describing this relationship. Goods x_i are generated uniformly at random from X , and the public data is $\mathbf{z}_n = \{(x_i, f_R(x_i))\}_{i=1}^n$ for some $n \geq 1$. The attributes describing the Seller's good are known to be the zero vector, so players form beliefs about $v := f_R(\vec{\mathbf{0}})$ where f_R is not known.

Restriction on Beliefs. Players know that f_R belongs to the set of *rectangular classification rules* \mathcal{F} , i.e. the set of functions $f_{R'}(x) = \mathbb{1}(x \in R')$ indexed to a hyper-rectangle R' in $[-1, 1]^m$. But players may hold different prior beliefs over this set. Fixing any prior $\pi \in \Delta(\mathcal{F})$, the posterior belief $\pi(f | \mathbf{z}_n)$ given data \mathbf{z}_n is a re-normalization of the prior over all rules consistent with the observed data (see Figure 2), and uniquely determines a belief for the value of the Seller's good. The set of learning rules is $\mathcal{M} = \{\mu_\pi\}_{\pi \in \Delta(\mathcal{F})}$, where each $\mu_\pi(\mathbf{z}_n) \in \Delta(\{0, 1\})$ assigns probability $\pi(\{f : f(\vec{\mathbf{0}}) = 1\} | \mathbf{z}_n)$ to the Seller's good having a high value.

Plausibility Interval. Let $[\underline{p}^n, \bar{p}^n]$ be the plausibility interval for rationalizability of the action *enter* given n observations. An exact characterization of $[\underline{p}^n, \bar{p}^n]$ is given in Lemma 2 in the appendix. This plausibility interval obeys the following comparative statics:

Claim 1. *The probability $\underline{p}^n = 0$ for every $n \in \mathbb{Z}_+$. Additionally:*

(a) *Fixing any number of attributes $m \in \mathbb{Z}_+$, the probability $\bar{p}^n \rightarrow 0$ as $n \rightarrow \infty$.*

¹⁸Such a function may represent, for example, whether all attributes fall into an acceptable range, as judged by a downstream buyer.

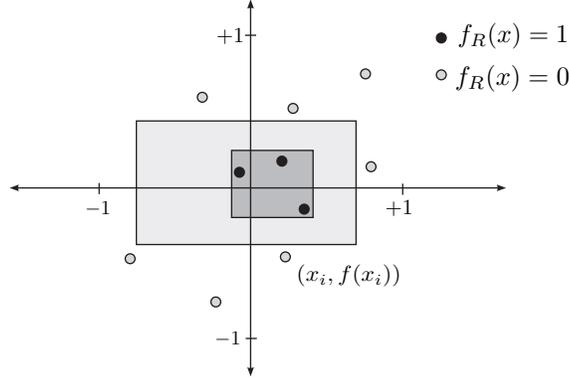


Figure 2: The circles represent the observed data. Each good is described by a vector in $[-1, 1] \times [-1, 1]$. The circle is black if its valuation is 1 and gray if its valuation is low. A rule is consistent with the data if it correctly predicts the valuation for each observation. Two rectangular classification rules are depicted: each predicts ‘1’ for goods in the shaded region and ‘0’ for goods outside. Both are consistent with the observed data.

(b) Fixing any number of data observations $n \in \mathbb{Z}_+$, the probability \bar{p}^n is increasing in the number of attributes m , and $\bar{p}^n \rightarrow 1$ as $m \rightarrow \infty$.

The probability that entering is rationalizable for the Seller for *all* permitted types is zero no matter the parameter values. But the probability that entering is rationalizable for the Seller given *some* permitted type varies depending on n and m . Part (a) says that as the number of observations n grows large, the probability \bar{p}^n vanishes to zero, implying that the plausibility interval converges to a degenerate interval at zero. The infinite-data limit thus returns the prediction of “no trade” consistent with the common prior assumption. But if the quantity of data is finite, and the number of attributes is large, then \bar{p}^n can be substantially greater than zero. Indeed, Part (b) of the claim says that this probability \bar{p}^n can be made arbitrarily close to 1 by increasing the dimensionality of the learning problem via choice of large m . This reflects that in a high-dimensional learning problem, many classification rules are likely to be consistent with the data, including some that yield conflicting predictions. Thus, “rational” disagreement is a distinct possibility.

Figure 3 shows how the plausibility interval varies in the number of observations n for specific examples of m , assuming the true function to be $f_R = \mathbb{1}(x \in [-0.1, 0.1]^m)$. For example, if there are 10 attributes and players observe only 20 goods, then the plausibility interval is $[\underline{p}^n, \bar{p}^n] = [0, 0.99]$.

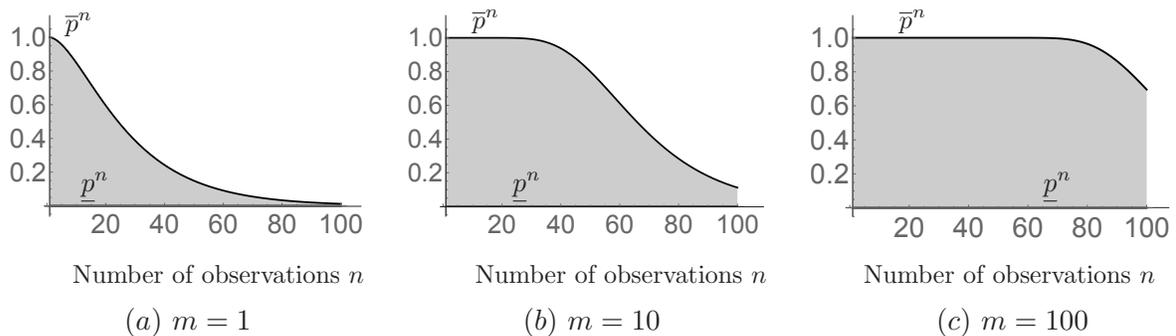


Figure 3: The shaded area depicts plausibility intervals $[p^n, \bar{p}^n]$ for the rationalizability of entering given n common observations.

In this game, prediction of trade fails at the large data limit, but remains plausible for small data sizes. The next section describe a game in which a prediction of interest holds in the infinite-data limit, but may fail when players have beliefs based on a small quantity of data.

3.2 Coordination with Noisy Data

A contagious disease spreads across a population at an unknown speed. Two states are connected by travel, and their governors choose between implementing a *strong* or a *weak* lockdown policy in their states to slow the spread of the disease. Implementation of the strong lockdown policy entails a large economic cost, but if the states coordinate on this policy, then the disease will be suppressed with certainty.

The number of reported cases on days $t = 1, 2, \dots$ grows exponentially according to

$$\log y_t = \beta t + \varepsilon_t \quad (4)$$

where the noise term ε_t is iid across time and follows a normal distribution with known parameters $\mu = 0$ and $\sigma^2 > 0$. The constant β is not known. Payoffs are given by the following matrix:

	<i>Strong</i>	<i>Weak</i>
<i>Strong</i>	$-1, -1$	$-1 - \beta, -\beta$
<i>Weak</i>	$-\beta, -1 - \beta$	$-\beta, -\beta$

The economic cost of the strong lockdown is normalized to 1, and the cost of letting the disease progress without a strong lockdown is given by its growth rate β . A weak lockdown is strictly dominant if $\beta < 1$, but coordination on the strong lockdown is the Pareto-dominant Nash equilibrium if $\beta > 1$.

Data. The two governors form beliefs about the growth rate of the disease based on a public data set $\{(t, y_t)\}_{t=1}^n$, which describes the number of reported cases of the disease, y_t , on days $t = 1, 2, \dots, n$, where each y_t is a random variable governed by (4).

Restriction on Beliefs. Define $\hat{\beta}(\mathbf{z}_n)$ to be the ordinary least-squares estimate of β from the data $\mathbf{z}_n := \{(t, \log y_t)\}_{t=1}^n$, and let ϕ_n be the constant such that $\mathcal{C}(\mathbf{z}_n) = \{\beta : |\beta - \hat{\beta}(\mathbf{z}_n)| \leq \phi_n\}$ is a $(1 - \alpha)$ -confidence interval for β , where $\alpha \in (0, 1)$ is fixed. Suppose players have common certainty in the event that all first-order beliefs belong to the set $\mathcal{B}(\mathbf{z}_n) := \Delta(\mathcal{C}(\mathbf{z}_n))$, i.e. all distributions supported on the confidence interval.

Plausibility Interval. Let $[\underline{p}^n, \bar{p}^n]$ be the plausibility interval for the strong lockdown given n observations. Lemma 3 in the appendix explicitly characterizes \underline{p}^n and \bar{p}^n . This plausibility interval satisfies the following comparative static:

Claim 2. *Suppose the actual growth rate is fast ($\beta > 1$), so that the strong lockdown is rationalizable given complete information of the payoffs. Then, for every $\sigma^2 > 0$, both \bar{p}^n and \underline{p}^n are increasing in n , while for every n , both \bar{p}^n and \underline{p}^n are decreasing in σ^2 .¹⁹*

That is, the analyst has high confidence in predicting that the strong lockdown is rationalizable when the reporting noise σ^2 is small relative to the number of observations n . These comparative statics are complemented by Figure 4, which shows how the plausibility interval varies in n for specific levels of reporting noise σ .

For example, if reporting noise is $\sigma = 10$ and the number of observations is $n = 100$, then the plausibility interval $[0.99, 1]$ is nearly degenerate at certainty. On the other hand, if reporting noise is $\sigma = 100$, then the plausibility interval is $[0.08, 0.99]$ for the same number of observations $n = 100$, suggesting substantial ambiguity regarding whether the strong lockdown is a good prediction of play.

¹⁹If instead $\beta < 1$, then the reverse statements hold; that is, the probabilities \bar{p}^n and \underline{p}^n are decreasing in n and increasing in σ^2 .

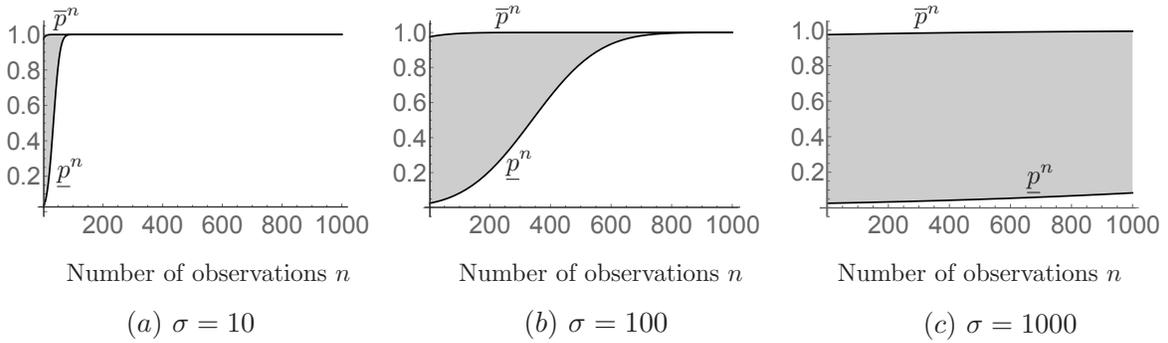


Figure 4: The shaded area depicts plausibility intervals $[\underline{p}^n, \bar{p}^n]$ for the rationalizability of the strong lockdown given n common observations, and allowing the reporting noise σ to vary. (In all panels, $\beta = 2$ and $\alpha = 0.05$.)

4 Main Results

Subsequently, I develop results for the plausibility interval that hold more generally, including when explicit characterization of the plausibility interval is not possible. Section 4.1 characterizes the limiting behavior of the probabilities $\bar{p}^n(a_i)$ and $\underline{p}^n(a_i)$ as the quantity of data n gets large, and Section 4.2 provides bounds that hold for small sample sizes.

4.1 Asymptotic Behavior

I first demonstrate by example that even if each learning rule eventually recovers the true parameter, it is possible that the plausibility interval given infinite data may look very different from plausibility intervals given a large but finite quantity of data. In this case, predictions based on an implicit assumption that agents have completely coordinated their beliefs may be misleading about real environments with finite data.

Example 1 (Rich Sets of Priors and Likelihoods). Consider the trading game described in Section 3.1. Suppose that players commonly observe a sequence of realizations from the set $Z = \{0, 1\}$. Learning rules $\mu_{\pi, q} \in \mathcal{M}$ are indexed to parameters $\pi \in (0, 1)$ and $q \in (1/2, 1)$, where the parameter π is the prior probability of value $v = 1$, and q identifies the following signal structure:

$$\begin{array}{ccc}
 z = 0 & z = 1 & \\
 v = 0 & q & 1 - q \\
 v = 1 & 1 - q & q
 \end{array}$$

Each rule $\mu_{\pi,q}$ is identified with prior π and signal structure q , and maps the observed signal outcomes into the posterior belief over $\{0, 1\}$. Assume that the true data-generating process belongs to this class; that is, there exists some $q^* \in (1/2, 1)$ such that the distribution over the signal set $\{0, 1\}$ is $(q^*, 1 - q^*)$ when $v = 0$, and the distribution is $(1 - q^*, q^*)$ when $v = 1$. Then,

$$\lim_{n \rightarrow \infty} [\underline{p}^n(a_i), \bar{p}^n(a_i)] = [0, 1]$$

while

$$[\underline{p}^\infty(a_i), \bar{p}^\infty(a_i)] = \{0\}$$

so the plausibility interval is discontinuous at $n = \infty$. (See Appendix O.1 for the proof of this statement.) Thus although trade will not occur in the limiting game, it is a plausible outcome when players commonly observe arbitrarily large (but finite) quantities of data.

In the above example, all learning rules lead to the same belief (that is, there is *asymptotic agreement* in the sense of Acemoglu et al. (2015)), so the result does not emerge mechanically from intrinsic long-run differences across the learning rules. But the set of learning rules fails a stronger condition:

Assumption 1 (Uniform Convergence). $\lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{M}} d_P(\mu(Z^n), \mu^\infty) = 0$ *P*-a.s., where d_P is the Prokhorov metric on $\Delta(\Theta)$.²⁰

Assumption 1 holds if for each learning rule $\mu \in \mathcal{M}$, the (random) belief $\mu(Z^n)$ almost surely converges to μ^∞ as the quantity of data n grows large, and moreover the speed of convergence does not vary too much across the different learning rules in \mathcal{M} .²¹ Specifically, the sequence of beliefs $\{\mu(Z^n)\}$ must converge to μ^∞ (as $n \rightarrow \infty$) uniformly across $\mu \in \mathcal{M}$. A sufficient condition for Assumption 1 to hold is that the set of learning rules \mathcal{M} is finite. But failures of Assumption 1 occur for classes of learning rules that we may consider plausible—in particular, if the class \mathcal{M} is too rich as in Example 1.

Under this assumption, each agent i 's type in the idealized infinite-data limit $n = \infty$ is the one with common certainty in the event that all players have first-order belief μ^∞ , denoted

²⁰For any $\nu, \nu' \in \Delta(\Theta)$, the Prokhorov distance between these measures is $d_P(\nu, \nu') = \inf\{\epsilon > 0 : \nu(A) \leq \nu'(A^\epsilon) + \epsilon \text{ for all Borel-measurable } A \subseteq \Theta\}$, where A^ϵ denotes the ϵ -neighborhood of A in the sup-norm.

²¹The limiting belief μ^∞ can be interpreted as a common prior, following what Morris (1995) calls the “frequentist justification” for assumption of a common prior.

t_i^∞ .²² It is not critical for the subsequent results that all differences in beliefs are removed in the limit (see Section 5), but maintaining Assumption 1 allows us to explore more precisely the scope for disagreement in an environment in which agreement is eventually feasible. The first main result says that under this assumption, the limiting plausibility intervals can be tightly linked to predictions in the limiting game.

Theorem 1. *Suppose Assumption 1 is satisfied.*

(a) *If a_i is strictly rationalizable for player i of type t_i^∞ , then*

$$\lim_{n \rightarrow \infty} [\underline{p}^n(a_i), \bar{p}^n(a_i)] = \{1\}.$$

(b) *If a_i is not rationalizable for player i of type t_i^∞ , then*

$$\lim_{n \rightarrow \infty} [\underline{p}^n(a_i), \bar{p}^n(a_i)] = \{0\}.$$

This theorem says that if an action a_i is strictly rationalizable for player i given infinite data, then $\bar{p}^n(a_i)$ and $\underline{p}^n(a_i)$ both converge to 1 as n grows large.²³ Thus, when agents observe sufficiently large quantities of public data, the analyst should be arbitrarily confident in predicting that a_i is rationalizable. On the other hand, if action a_i is not rationalizable given infinite data, then $\bar{p}^n(a_i)$ and $\underline{p}^n(a_i)$ both converge to 0, so for large data sets the analyst should be confident in predicting that a_i is *not* rationalizable.²⁴

This theorem applies to the two examples in Section 3, both of which involved sets of

²²The limiting game can be interpreted as a complete information game where player i 's payoffs are given by $\int_{\Theta} u_i(a, \theta) d\mu^\infty(\theta)$. But note that Assumption 1 does not guarantee that players' beliefs concentrate on a neighborhood of payoffs around this limiting payoff as n grows large. The special case imposing this additional structure on beliefs is considered in Section 4.2.1.

²³In the special case in which the limiting belief μ^∞ is degenerate at a limiting parameter θ^∞ , and players have common certainty that players' first-order beliefs have support in a shrinking neighborhood of θ^∞ (see Section 4.2.1 for a more formal development), then the property that $\bar{p}^n(a_i) \rightarrow 1$ is (nearly) equivalent to the property that a_i is robustly rationalizable, as defined in Morris et al. (2012), with the difference that Morris et al. (2012) consider almost common belief in the exact parameter θ^∞ , while I consider common certainty in a neighborhood of θ^∞ . As Proposition 1 in Morris et al. (2012) shows, strict rationalizability is a sufficient condition for robust rationalizability. See also Kajii and Morris (2020) for related results.

²⁴The intermediate case in which a_i is rationalizable for player i given infinite data, but not strictly rationalizable, is subtle and depends on details of the game. See Online Appendix O.3 for examples in which $\lim_{n \rightarrow \infty} [\underline{p}^n(a_i), \bar{p}^n(a_i)] = \{1\}$ and in which $\lim_{n \rightarrow \infty} [\underline{p}^n(a_i), \bar{p}^n(a_i)] = [0, 1]$. Note that the latter corresponds to a maximally ambiguous outcome—no amount of data is decisive on whether or not the action should be considered rationalizable.

learning rules that satisfy Assumption 1. In the trade game in Section 3.1, entering is not rationalizable for the Seller given infinite data, and correspondingly the plausibility interval for rationalizability of trade shrinks to a degenerate interval at zero as n grows large. In the coordination game in Section 3.2, the strong lockdown is rationalizable given infinite data, and correspondingly the plausibility interval converges to a degenerate interval at 1 as n grows large. Theorem 1 tells us that these properties do not depend on the specific sets of learning rules considered in these examples. We would obtain the same limiting results for any alternative set of learning rules satisfying Assumption 1.

Theorem 1 builds on prior results regarding topologies on the universal type space. Consider any sequence of types $(t_i^n)_{n=1}^\infty$ where each $t_i \in T_i^{\mathcal{B}(\mathbf{z}_n)}$. Under Assumption 1, types t_i^n for large n (almost surely) have common certainty that first-order beliefs lie in a small neighborhood of the limiting belief μ^∞ . Thus, each sequence (t_i^n) can be shown to converge to t_i^∞ in the *uniform-weak topology* (Chen et al., 2010) on the universal type space (see Lemma 3). Part (b) of the theorem then follows from the property of upper hemi-continuity of the rationalizability correspondence in the uniform-weak topology (Chen et al., 2010).

Part (a) of the theorem is related to the lower hemi-continuity of strict rationalizability in the uniform-weak topology (as shown in Chen et al. (2010)),²⁵ which guarantees that for any sequence of types $(t_i^n)_{n=1}^\infty$ from $T_i^{\mathcal{B}(\mathbf{z}_n)}$, the action a_i must eventually be rationalizable along the sequence. But this property is not sufficient to demonstrate Part (a), since the rates of this convergence can differ substantially across different sequences of types from $T_i^{\mathcal{B}(\mathbf{z}_n)}$. (This concern is outside of the scope of Chen et al. (2010), which did not consider the convergence of sets.²⁶) For eventual complete $\mathcal{B}(\mathbf{z}_n)$ -rationalizability, we need that a_i is rationalizable for *all* types from $T_i^{\mathcal{B}(\mathbf{z}_n)}$ when n is sufficiently large.

To establish this, I show that there is a P -measure 1 set of sequences along which the sets $(T_i^{\mathcal{B}(\mathbf{z}_n)})_{n=1}^\infty$ converge to the singleton set $\{t_i^\infty\}$ in the Hausdorff metric induced by the uniform-weak metric. The key lemma underlying this result, Lemma 6, relates the degree of “strictness” of rationalizability of action a_i at the limiting type t_i^∞ to the size of the neighborhood around μ^∞ such that common certainty of that neighborhood implies rationalizability of a_i . The stronger property that types converge uniformly over the set

²⁵It is crucial that convergence occurs in this topology and not simply the product topology, as otherwise the negative results of Weinstein and Yildiz (2007) would apply.

²⁶Besides this more substantial contrast, there is the additional technical difference that Chen et al. (2010) assume a finite state space Θ while I consider a more general environment. For this reason, I need to introduce new arguments utilizing Lipschitz continuity of the mapping g from states to payoffs.

$T_i^{\mathcal{B}(z_n)}$ delivers the desired result.

4.2 (Small) Finite Samples

The previous section characterized plausibility intervals given large numbers n of common observations. I now focus on the setting of small n , and bound the extent to which the agent’s plausibility interval $[\underline{p}^n(a_i), \bar{p}^n(a_i)]$ diverges from its asymptotic limit $[\underline{p}^\infty(a_i), \bar{p}^\infty(a_i)]$. Throughout this section, I impose the simplifying assumptions that observations are i.i.d., and that they take values from a finite set \mathcal{Z} :

Assumption 2. $Z_1, \dots, Z_n \sim_{i.i.d.} Q$.

Assumption 3. $|\mathcal{Z}| < \infty$.

4.2.1 Lower Bound

First fix an action a_i that is strictly rationalizable at the $n = \infty$ limit. By Theorem 1, the analyst’s plausibility interval $[\underline{p}^n(a_i), \bar{p}^n(a_i)]$ converges to a degenerate interval at 1. Proposition 1, below, provides a lower bound on $\underline{p}^n(a_i)$ that explains how fast this convergence occurs.

There are two properties that determine this bound. The first is a strategic property that depends only on limiting expected payoffs: how “strictly” is the action rationalizable in the limiting game? The more slack there is, the less essential it is that players hold exactly the same belief over payoffs. (For example: if coordination on a strict lockdown leads to payoffs for each player that are much higher than his next best action, then it is not critical that these players hold exactly the same beliefs over the growth rate, so long as their beliefs are not too different.) The second is a property of the learning environment: given the set of learning rules \mathcal{M} and the data-generating process P , how much data must there be for agents to have achieved the necessary degree of (common knowledge) agreement about the payoffs? Actions that are more strictly rationalizable, and which feature faster common learning about the payoffs, are more plausible given small quantities of data.

Formally, I define the strictness of rationalizability of action a_i for the limiting type t_i^∞ as follows. Say that a family of sets $(R_i[t_i])_{t_i \in T_i^*}$, where each $R_j[t_j] \subseteq A_j$, has the δ -*strict best reply property* if for each $i \in \mathcal{I}$, type $t_i \in T_i^*$, and action $a_i \in R_i[t_i]$ there is a conjecture

$\sigma_{-i} : \Theta \times T_{-i}^* \rightarrow \Delta(A_{-i})$ to which a_i is a δ -strict best reply for t_i .²⁷ Say that an action a_i is δ -strict rationalizable for type t_i if there exists a family of sets $(R_j[t_j])_{t_j \in T_j^*}$ with the δ -strict best reply property, where $a_i \in R_i[t_i]$.²⁸

Then if the action a_i is strictly rationalizable for type t_i^∞ and players have commonly observed n realizations, the probability that a_i is rationalizable for all permitted types can be lower bounded as follows.

Proposition 1. *Suppose Assumptions 1-3 hold, and suppose action a_i is strictly rationalizable for type t_i^∞ . Define*

$$\delta^\infty := \sup \{ \delta : a_i \text{ is } \delta\text{-strictly rationalizable for type } t_i^\infty \} \quad (5)$$

noting that this quantity is strictly positive. Further define

$$\xi := \sup_{\theta, \theta' \in \Theta} \|\theta - \theta'\|_\infty. \quad (6)$$

Then, for every $n \geq 1$,

$$\underline{p}^n(a_i) \geq 1 - \frac{2K\xi}{\delta^\infty} \mathbb{E} \left(\sup_{\mu \in \mathcal{M}} d_P(\mu(Z^n), \mu^\infty) \right) \quad (7)$$

where K is the Lipschitz constant of the map $g : \Theta \rightarrow U$.

Recalling that $\bar{p}^n(a_i) \geq \underline{p}^n(a_i)$ for every n , this proposition allows us to lower bound the plausibility interval $[\underline{p}^n(a_i), \bar{p}^n(a_i)]$.

This bound formalizes the intuitions given previously. First, the expression in (7) is increasing in δ^∞ , so the “more strictly-rationalizable” the action is for the limiting type, the fewer observations are necessary for the prediction to hold. Second, the bound is decreasing in $\mathbb{E}(\sup_{\mu \in \mathcal{M}} d_P(\mu(Z^n), \mu^\infty))$, which is the expected distance from the limiting belief μ^∞ to the farthest belief in the set $\mathcal{B}(Z^n)$. When Assumption 1 is satisfied, then $\mathbb{E}(\sup_{\mu \in \mathcal{M}} d_P(\mu(Z^n), \mu^\infty)) \rightarrow 0$ as $n \rightarrow \infty$, and the speed of this convergence can be interpreted as the speed at which players *commonly learn* (Cripps et al., 2008). Thus, Proposition 1 says that the more strictly-rationalizable an action is, and the quicker players commonly learn, the fewer observations are necessary for limiting predictions to carry over to small-data settings. In general, we expect common learning to occur more slowly in environments with noisy and high-dimensional data, as well as in environments in which there is substantial

²⁷That is, $\int_\Theta u_i(a_i, \sigma_{-i}(\theta, t_{-i}), \theta) t_i [d\theta \times dt_{-i}] - \int_\Theta u_i(a'_i, \sigma_{-i}(\theta, t_{-i}), \theta) t_i [d\theta \times dt_{-i}] \geq \delta$ for all $a'_i \neq a_i$.

²⁸This is equivalent to γ -rationalizability from Dekel et al. (2007), where $\gamma = -\delta$.

heterogeneity across agents about the “right” model for interpreting the data.

In an important special case, the limiting belief μ^∞ is a point mass at some θ^∞ , and the sets $\mathcal{B}(\mathbf{z}_n)$ consist of beliefs with support on shrinking neighborhoods of θ^∞ . Formally, let

$$\mathcal{C}(\mathbf{z}_n) := \bigcup_{\mu \in \mathcal{M}} \text{supp } \mu(\mathbf{z}_n) \quad \forall \mathbf{z}_n \in \mathcal{Z}^n$$

with the implication that every $\mu(\mathbf{z}_n)$, $\mu \in \mathcal{M}$, assigns probability 1 to $\mathcal{C}(\mathbf{z}_n)$. If $\mathcal{C}(\mathbf{z}_n)$ collapses to the singleton set $\{\theta^\infty\}$ as $n \rightarrow \infty$, then the bound in Proposition 1 can be simplified as follows.

Assumption 4. $\sup_{\theta \in \mathcal{C}(\mathcal{Z}^n)} \|\theta - \theta^\infty\|_\infty$ converges to zero P -almost surely.²⁹

Proposition 2. *Suppose Assumptions 1-4 hold, and the action a_i is strictly rationalizable for type t_i^∞ . Then, for every $n \geq 1$,*

$$\underline{p}^n(a_i) \geq 1 - \frac{2K}{\delta^\infty} \mathbb{E} \left(\sup_{\theta \in \mathcal{C}(\mathcal{Z}^n)} \|\theta - \theta^\infty\|_\infty \right)$$

where K is the Lipschitz constant of the map $g : \Theta \rightarrow U$.

The expressions in Propositions 1 and 2 can be used to derive quantitative bounds for specific sets of learning rules, as in the following example:

Example 2. Consider the payoff matrix from Section 3.2 with unknown parameter $\beta \in \mathbb{R}$. Suppose that players commonly observe n public signals $z_t = \beta + \varepsilon_t$, with standard normal error terms ε_t that are i.i.d. across observations. The set of learning rules is $\mathcal{M} = \{\mu_x\}_{x \in [-\eta, \eta]}$, where each learning rule μ_x is identified with the prior belief $\beta \sim \mathcal{N}(x, 1)$, and maps data into a point mass at the posterior expectation of β . The set $\mathcal{C}(\mathbf{z}_n)$ thus consists of the posterior expectations under the different priors, and players have common certainty in the event that all players have first-order beliefs with support on $\mathcal{C}(\mathbf{z}_n)$. Let the true value of β satisfy $\beta > 1$. Then, the following result is implied by Proposition 2.

Corollary 1. *For each $n \geq 1$,*

$$\underline{p}^n(\text{strong}) \geq 1 - \frac{1}{\beta - 1} \left(\sqrt{\frac{2}{\pi n}} + \frac{\beta + \eta}{n + 1} \right)$$

²⁹That is, there is a P -measure 1 set of (infinite) sequences such that $\sup_{\theta \in \mathcal{C}(\mathbf{z}_n)} \|\theta - \theta^\infty\|_\infty \rightarrow 0$ as $n \rightarrow \infty$ for each sequence \mathbf{z} in this set.

This bound is decreasing in η (the size of the model class), increasing in n (the number of observations), and increasing in $\beta - 1$ (the strictness of the solution at the limit).

The main value of these results is in demonstrating that it is possible to go beyond the asymptotic analysis of Theorem 1 by telling us how plausible a given prediction is when players observe a finite, and potentially small, quantity of data. The bounds in Propositions 1 and 2 depend not only on the basic game but also on properties of the learning environment, and thus are sensitive to changes in these parameters.³⁰ These results justify a view in which the same action may be rationalizable in one context (e.g. industry, time, culture) but not rationalizable in another, if those contexts involve different learning environments. For example: If an economic shock renders most of the past data irrelevant, and players have not reached consensus on the right models for the new regime, then Propositions 1 and 2 suggest that actions which rely on (common knowledge of) agreement are worse predictions, even if players’ payoffs functions have not changed.

4.2.2 Upper Bound

Now fix an action a_i that is *not* rationalizable for player i in the $n = \infty$ limit. By Part (b) of Theorem 1, the analyst’s plausibility interval $[\underline{p}^n(a_i), \bar{p}^n(a_i)]$ must converge to a degenerate interval at zero. Thus, given sufficient data, the analyst should be confident predicting that a_i will not be played. But given small quantities of data n , the action a_i may still constitute a plausible prediction, as we saw in the trading game studied in Section 3.1. Proposition 3, below, provides an upper bound on $\bar{p}^n(a_i)$, which tells us whether the analyst should consider a_i to be a plausible prediction when players have commonly observed only n samples.

The key determinant of this bound is the “typicality” of samples of length n that lead to rationalizability of the action a_i . Let \mathcal{Z}_{a_i} be all data sets \mathbf{z}_n given which the action a_i is partially $\mathcal{B}(\mathbf{z}_n)$ -rationalizable. (This set must be determined on a case-by-case basis.) Let $\hat{Q}_{\mathbf{z}_n} \in \Delta(\mathcal{Z})$ be the empirical measure associated with data set \mathbf{z}_n . The Kullback-Leibler divergence from $\hat{Q}_{\mathbf{z}_n}$ to the actual data-generating distribution Q is $D_{KL}(Q \parallel \hat{Q}_{\mathbf{z}_n}) =$

³⁰For example, Corollary 2 tells us that if players have common knowledge of the set of priors indexed to $\eta = 0.72$, a desired lower bound of 50% on plausibility is achieved with 10 observations, while for the larger set of priors indexed to $\eta = 40$, the same lower bound is achieved with 100 observations.

$\sum_{z \in \mathcal{Z}} Q(z) \log \left(\frac{Q(z)}{\hat{Q}_{\mathbf{z}_n}(z)} \right)$. Define

$$Q_n^* = \operatorname{argmin}_{\hat{Q} \in \{\mathbf{z}_n \in \mathbb{Z}_{a_i}^n\}} D_{KL}(Q \parallel \hat{Q}_{\mathbf{z}_n})$$

to be the empirical measure (associated with a data set in \mathbb{Z}_{a_i}) that minimizes Kullback-Leibler divergence to Q . This is the most likely data set of size n that would lead to a_i being rationalizable. Sanov's theorem yields the following result.

Proposition 3. *Suppose Assumptions 1-3 hold, and action a_i is not rationalizable for type t_i^∞ ; then, for every $n \geq 1$,*

$$\bar{p}_n(a_i) \leq (n+1)^{|\mathcal{Z}|} 2^{-n D_{KL}(Q_n^* \parallel Q)}.$$

Similar to the lower bounds given in Section 4.2.1, this upper bound depends not only on properties of the game but also on the learning environment. Proposition 3 says that actions which are not predicted in the limiting game may nevertheless be plausible for certain values of the sample size n .

5 Extensions

Asymptotic Disagreement. Assumption 1 guarantees that beliefs produced by learning rules in \mathcal{M} uniformly converge to a common limiting belief μ^∞ , so learning eventually removes all differences in beliefs. It is possible to relax this assumption to allow players to have heterogeneous beliefs even in the limit. For any $\epsilon \geq 0$, say that the class of learning rules \mathcal{M} satisfies ϵ -Uniform Convergence if

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

so that the set of expected parameters converges to an ϵ -neighborhood of μ^∞ . Theorem 1 holds as long as the set of learning rules \mathcal{M} satisfies ϵ -Uniform Convergence for some $\epsilon \leq \delta^\infty / (2K\xi)$ (where $\epsilon = 0$ returns the previous result). The rate results also hold without modification.

Approximate Common Certainty. Suppose players have common p -belief in $\mathcal{B}(\mathbf{z}_n)$ for some probability p not necessarily equal to 1. Formally, for any $p \in [0, 1]$, player $i \in \mathcal{I}$, and set $\mathcal{B} \subseteq \Delta(\Theta)$, define

$$B_i^{1,p}(\mathcal{B}) := \{t_i \in T_i^* : \text{marg}_\Theta \kappa_i^*(t_i) \in \mathcal{B}\}.^{31}$$

For each $k > 1$, and again for each player $i \in \mathcal{I}$, recursively define

$$B_i^{k,p}(\mathcal{B}) = \left\{ t_i \in T_i^* : \kappa_i^*(t_i) \left(\Theta \times \prod_{j \neq i} B_j^{k-1,p}(\mathcal{B}) \right) \geq p \right\}.$$

Then $T_i^{\mathcal{B},p} = \bigcap_{k \geq 1} B_i^{k,p}(\mathcal{B})$ is the set of player i types that have common p -belief in the event that all players' first-order beliefs belong to \mathcal{B} . We can develop plausibility intervals as in the main text, replacing $T_i^{\mathcal{B}(\mathbf{z}_n)}$ with $T_i^{p,\mathcal{B}(\mathbf{z}_n)}$. Then, there exists a p^* such that so long as players have common p -belief in the event that all players' first-order beliefs belong to $\mathcal{B}(\mathbf{z}_n)$, where $p > p^*$, then Theorem 1 holds as stated. Rate results similar to those in Section 4.2 can also be obtained (see Online Appendix O.4 for details). Both extensions rely on boundedness of the payoff range.

Plausibility Intervals for Equilibrium. The proposed approach can be paired with solution concepts besides rationalizability. For example, suppose we are interested in evaluating how plausible it is that the action profile $a \in A$ is part of a (pure-strategy) Bayesian Nash equilibrium. The analogous plausibility interval is $[p^n(a), \bar{p}^n(a)]$, where the lower bound $p^n(a)$ is the probability (over possible datasets \mathbf{z}_n) that a_i is a best reply to a_{-i} for every player i of any type $t_i \in T_i^{\mathcal{B}(\mathbf{z}_n)}$. The upper bound \bar{p}^n is the probability that there exists a belief-closed type space $(T_i, \kappa_i)_{i \in \mathcal{I}}$ where each $T_i \subseteq T_i^{\mathcal{B}(\mathbf{z}_n)}$, and the strategy profile σ with $\sigma_i(t_i) = a_i$ for all $i, t_i \in T_i$ is a Bayesian Nash equilibrium. Then, Theorem 1 holds with “strict rationalizability” replaced by “strict equilibrium,” and the rate results provided in Theorem 1 hold when δ^∞ is replaced with an analogous notion for the strictness of the equilibrium in the limiting game.

³¹This set has the same definition as $B_i^{1,1}$ from the main text. It is possible to relax the assumptions further, so that $B_i^{1,p}(\mathcal{B}) := \{t_i \in T_i^* : \sup_{\nu \in \mathcal{B}} d_P(\nu, \text{marg}_\Theta \kappa_i^*(t_i)) \leq 1 - p\}$, but this does not correspond to any standard definitions.

State-Dependent Data Generating Process. In the main text, P denotes the true data-generating process. It is possible to extend the framework to a setting in which the data generating process is state-dependent and the analyst has uncertainty over the state. Specifically, suppose the analyst has a prior belief $\pi \in \Delta(\Theta)$ and knowledge of a family of data-generating processes $(P_\theta)_{\theta \in \Theta}$. The main definitions and results in this paper can be extended to accommodate this perspective in the following way. For each state $\theta \in \Theta$, let μ_θ^∞ be the limiting belief when players observe data drawn according to P_θ . Then extend Assumption 1 to hold across states; that is,

$$\lim_{n \rightarrow \infty} \sup_{\theta \in \Theta} \sup_{\mu \in \mathcal{M}} d_{P_\theta}(\mu(Z^n), \mu_\theta^\infty) = 0 \quad P\text{-a.s.}$$

The natural plausibility interval is then an ex-ante analogue of the definition given in the present paper. That is, for each state θ let $[\bar{p}_\theta^n, \underline{p}_\theta^n]$ denote the plausibility interval given process P_θ , and define the new ex-ante plausibility interval to be the one that averages across these intervals according to the prior π . Applying Theorem 1 state-by-state, we can demonstrate a similar result in this ex-ante environment, while the rate results in Propositions 1-3 can be modified accordingly by taking an average of the bounds with respect to π .

6 Conclusion

Economists make predictions in incomplete information games based on models of unobservable beliefs. A large literature on the robustness of strategic predictions to the specification of agent beliefs provides guidance regarding whether these predictions should be trusted. These robustness notions tend to be qualitative—we learn whether the prediction is or isn't robust to perturbations in the agents' beliefs. Here I offer a different perspective, namely a quantitative metric for how robust the prediction is. The metric depends on the quantity of data that agents get to see. Predictions that hold given infinite quantities of data may not hold given large quantities of data, and those that hold given large quantities of data may not hold in environments where agents see only a few observations. Likewise, predictions that don't hold at the limit may nevertheless be plausible when agents' beliefs are coordinated by a small number of observations. The proposed framework provides a way of formalizing this, generating new comparative statics for how the plausibility of a strategic prediction varies

with primitives of the learning environment.

References

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2015): “Fragility of Asymptotic Agreement under Bayesian Learning,” *Theoretical Economics*, 11, 187–225.
- AL-NAJJAR, N. (2009): “Decisionmakers as Statisticians: Diversity, Ambiguity, and Learning,” *Econometrica*, 77, 1371–1401.
- AL-NAJJAR, N. AND M. PAI (2014): “Coarse Decision Making and Overfitting,” *Journal of Economic Theory*, 150, 467–486.
- ALONSO, R. AND O. CÂMARA (2016): “Bayesian Persuasion with Heterogeneous Priors,” *Journal of Economic Theory*, 165, 672–706.
- AUMANN, R. J. (1976): “Agreeing to Disagree,” *The Annals of Statistics*, 4, 1236–1239.
- BATTIGALLI, P. AND A. PRESTIPINO (2011): “Transparent Restrictions on Beliefs and Forward Induction Reasoning in Games with Asymmetric Information,” *The B.E. Journal of Theoretical Economics*, 13, 79–130.
- BATTIGALLI, P. AND M. SINISCALCHI (2003): “Rationalization and Incomplete Information,” *Advances in Theoretical Economics*, 3, 1534–5963.
- BRANDENBURGER, A. AND E. DEKEL (1993): “Hierarchies of Belief and Common Knowledge,” *Journal of Economic Theory*, 59, 189–198.
- BRANDENBURGER, A., A. FRIEDENBERG, AND J. KIESLER (2008): “Admissibility in Games,” *Econometrica*, 76, 307–352.
- 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*, 5, 445–478.
- (2017): “Characterizing the Strategic Impact of Misspecified Beliefs,” *Review of Economic Studies*, 84, 1424–1471.
- CHEN, Y.-C. AND S. TAKAHASHI (2020): “On Robust Selection and Robust Rationalizability,” Working Paper.
- CRIPPS, M., J. ELY, G. MAILATH, AND L. SAMUELSON (2008): “Common Learning,” *Econometrica*, 76, 909–933.

- DEKEL, E., D. FUDENBERG, AND D. LEVINE (2004): “Learning to Play Bayesian Games,” *Games and Economic Behavior*, 46, 282–303.
- DEKEL, E., D. FUDENBERG, AND S. MORRIS (2006): “Topologies on Types,” *Theoretical Economics*, 1, 275–309.
- (2007): “Interim Correlated Rationalizability,” *Theoretical Economics*, 2, 15–40.
- ELIAZ, K. AND R. SPIEGLER (2020): “A Model of Competing Narratives,” *American Economic Review*, 110, 3786–3816.
- ELY, J. AND M. PESKI (2011): “Critical Types,” *Review of Economic Studies*, 78, 907–937.
- ESPONDA, I. (2013): “Rationalizable Conjectural Equilibrium: A Framework for Robust Predictions,” *Theoretical Economics*, 8, 467–501.
- GAYER, G., I. GILBOA, AND O. LIEBERMAN (2007): “Rule-Based and Case-Based Reasoning in Housing Prices,” *The B.E. Journal of Theoretical Economics*, 7, 1–37.
- GEARY, R. C. (1935): “The Ratio of the Mean Deviation to the Standard Deviation as a Test for Normality,” *Biometrika*, 27, 310–332.
- GIBBS, A. AND F. SU (2002): “On Choosing and Bounding Probability Metrics,” *International Statistic Review*.
- GILBOA, I. AND D. SCHMEIDLER (1995): “Case-Based Decision Theory,” *The Quarterly Journal of Economics*, 110, 605–639.
- (2003): “Inductive Inference: An Axiomatic Approach,” *Econometrica*, 71, 1–26.
- HAGHTALAB, N., M. O. JACKSON, AND A. D. PROCACCIA (2020): “Belief Polarization in a Complex World: A Learning Theory Perspective,” Working Paper.
- HASTIE, T., R. TIBSHIRANI, AND J. FRIEDMAN (2009): *The Elements of Statistical Learning*, Springer.
- JEHIEL, P. (2005): “Analogy-Based Expectation Equilibrium,” *Journal of Economic Theory*, 123, 81–104.
- (2018): “Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism,” *American Economic Review*, 108, 1582–1597.
- KAJII, A. AND S. MORRIS (1997): “The Robustness of Equilibria to Incomplete Information,” *Econometrica*, 65, 1283–1309.
- (2020): “Refinements and Higher-Order Beliefs: A Unified Survey,” *The Japanese Economic Review*, 71, 7–34.
- MAILATH, G. AND L. SAMUELSON (2020): “Learning Under Diverse World Views: Model-

- Based Inference,” *American Economic Review*, 110, 1464–1501.
- MERTENS, J.-F. AND S. ZAMIR (1985): “Formulation of Bayesian Analysis for Games with Incomplete Information,” *International Journal of Game Theory*, 14, 1–29.
- MONDERER, D. AND D. SAMET (1989): “Approximating Common Knowledge with Common Beliefs,” *Games and Economic Behavior*, 1, 170–190.
- MORRIS, S. (1995): “The Common Prior Assumption in Economic Theory,” *Economics and Philosophy*, 11, 227–253.
- MORRIS, S., S. TAKAHASHI, AND O. TERCIEUX (2012): “Robust Rationalizability Under Almost Common Certainty of Payoffs,” *The Japanese Economic Review*, 63, 57–67.
- OLEA, J. L. M., P. ORTOLEVA, M. PAI, AND A. PRAT (2020): “Competing Models,” Working Paper.
- OLLÁR, M. AND A. PENTA (2017): “Full Implementation and Belief Restrictions,” *American Economic Review*, 107, 2243–2277.
- (2020): “A Network Solution to Robust Implementation: the Case of Identical but Unknown Distributions,” .
- PENTA, A. (2012): “Higher Order Uncertainty and Information: Static and Dynamic Games,” *Econometrica*, 80, 631–660.
- RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior Under Almost Common Knowledge,” *American Economic Review*, 79, 385–391.
- SALANT, Y. AND J. CHERRY (2020): “Statistical Inference in Games,” Working Paper.
- SETHI, R. AND M. YILDIZ (2016): “Communication with Unknown Perspectives,” *Econometrica*, 84, 2029–2069.
- SPIEGLER, R. (2016): “Bayesian Networks and Boundedly Rational Expectations,” *Quarterly Journal of Economics*, 131, 1243–1290.
- STEINER, J. AND C. STEWART (2008): “Contagion through Learning,” *Theoretical Economics*, 3, 431–458.
- WEINSTEIN, J. AND M. YILDIZ (2007): “A Structure Theorem for Rationalizability with Application to Robust Prediction of Refinements,” *Econometrica*, 75, 365–400.
- (2017): “Interim Correlated Rationalizability in Infinite Games,” *Journal of Mathematical Economics*, 72, 82–87.

Appendix

A Proofs for Section 3

A.1 Proof of Claim 1

Suppose that $f_R(\vec{\mathbf{0}}) = 1$, so that the value of the Seller's good is high. (The proof follows along similar lines in the other case.) I will first show that $\underline{p}^n = 0$ for every n . Let π be a point mass at f_R . An agent with this prior assigns probability 1 to $v = 1$ no matter the outcome of the data. Hence, the degenerate distribution at 1 belongs to $\mathcal{B}(\mathbf{z}_n)$ for every \mathbf{z}_n , so the type with common certainty in $v = 1$ is a permitted type for the Seller with probability 1.³² But entering is not rationalizable for the Seller of this type, implying $\underline{p}^n = 0$ as desired.

To prove Parts (a) and (b) of the claim, which refer to the probability \bar{p}^n , I first show that entering is rationalizable for some permitted type if and only if there exist $\tilde{f}, \tilde{f}' \in \mathcal{F}$ that are consistent with the data, and which make conflicting predictions at the zero vector (Lemma 1). I characterize the probability of this event in Lemma 2, from which the comparative statics for \bar{p}^n follow directly.

Lemma 1. *Fix an arbitrary data set $\mathbf{z}_n = \{(x_i, f(x_i))\}_{i=1}^n$. Entering is rationalizable for the Seller of some permitted type if and only if there exist $\tilde{f}, \tilde{f}' \in \mathcal{F}$ where*

- (1) $\tilde{f}(x_i) = \tilde{f}'(x_i) = f(x_i)$ for each observation $i = 1, \dots, n$
- (2) $\tilde{f}(\vec{\mathbf{0}}) = 1$ while $\tilde{f}'(\vec{\mathbf{0}}) = 0$

Proof. Suppose there exists a pair \tilde{f}, \tilde{f}' satisfying (1) and (2), and define $\pi_{\tilde{f}}, \pi_{\tilde{f}'} \in \Delta(\mathcal{F})$ to be point masses on \tilde{f} and \tilde{f}' . Since these rules are consistent with the data by (1), the posterior beliefs updated to \mathbf{z}_n are likewise degenerate at \tilde{f} and \tilde{f}' , and thus assign (respectively) probability 1 to $v = 1$ and probability 1 to $v = 0$. This implies that degenerate distributions at 1 and 0 both belong to $\mathcal{B}(\mathbf{z}_n)$. Entering is rationalizable for the Seller who believes that $v = 1$ with probability 1, and who believes with probability 1 that the Buyer believes $v = 0$ with probability 1. This belief is consistent with common certainty that all players have first-order beliefs in $\mathcal{B}(\mathbf{z}_n)$.

Now suppose that no such pair \tilde{f}, \tilde{f}' exists, implying either that every $\tilde{f} \in \mathcal{F}$ consistent with the data predicts $f(\vec{\mathbf{0}}) = 0$, or that every $\tilde{f} \in \mathcal{F}$ consistent with the data predicts $f(\vec{\mathbf{0}}) = 1$. Then either $\mathcal{B}(\mathbf{z}_n)$ is the singleton set consisting of a degenerate distribution at 1, or it is the singleton set consisting of a degenerate distribution at 0. If the former, the only permitted type is the one with common certainty in $v = 1$, and if the latter, the only permitted type is the one with common certainty in $v = 0$. Entering is not rationalizable for the Seller given either of these types. \square

³²Here, and elsewhere in the proof, type t_i has *common certainty in $v = 1$* if the type has common certainty in the event $\{f \mid f(\vec{\mathbf{0}}) = 1\} \times (X \times \{0, 1\})^\infty \times T_{-i}^*$.

Lemma 2. *Suppose the true function is $f(x) = \mathbb{1}(x \in R)$ where $R = [-r_1, \bar{r}_1] \times [-r_2, \bar{r}_2] \times \dots \times [-r_m, \bar{r}_m]$ for a sequence of constants $r_1, \bar{r}_1, \dots, r_m, \bar{r}_m \in (0, 1)$. Then*

$$\bar{p}^n(a_i) = 1 - \prod_{k=1}^m \left(1 - \left(\frac{1}{2} \right)^n [(2 - r_k)^n + (2 - \bar{r}_k)^n - (2 - (r_k + \bar{r}_k))^n] \right).$$

Proof. From Lemma 1, the probability \bar{p}^n is equal to the measure of data sets \mathbf{z}_n given which there exist $\tilde{f}, \tilde{f}' \in \mathcal{F}$ that are consistent with \mathbf{z}_n , and which make conflicting predictions at the input $\vec{\mathbf{0}}$. The true classification rule f is always consistent with the data, and predicts $f(\vec{\mathbf{0}}) = 1$, so a pair of such rules exists if we can additionally find a rule $\tilde{f} \in \mathcal{F}$ consistent with the data that predicts $\tilde{f}(\vec{\mathbf{0}}) = 0$.

A necessary and sufficient condition for existence of such a rule is that there is some dimension k on which either every observation x_i satisfies $x_i^k < 0$, or every x_i satisfies $0 < x_i^k$. This allows some $\tilde{f} \in \mathcal{F}$ to be consistent with the data, but to predict 0 at the zero vector.

For each dimension k , the probability that there is at least one observation x_i with $x_i^k \in [-r_k, 0)$ and at least one observation x_j with $x_j^k \in (0, \bar{r}_k]$ is

$$1 - \left(\frac{1}{2} \right)^n [(2 - r_k)^n + (2 - \bar{r}_k)^n - (2 - (r_k + \bar{r}_k))^n].$$

Observe that attribute values are independent across dimensions. So the probability that for every dimension k , there is at least one observation $x_i^k \in [-r_k, 0)$ and at least one observation x_j with $x_j^k \in (0, \bar{r}_k]$ is

$$\prod_{k=1}^m \left(1 - \left(\frac{1}{2} \right)^n [(2 - r_k)^n + (2 - \bar{r}_k)^n - (2 - (r_k + \bar{r}_k))^n] \right).$$

The desired probability is the complement of this event, which yields the expression in the lemma. \square

The following functional form is used in the main text:

Corollary 2. *In the special case in which the true function is $f(x) = \mathbb{1}(x \in R)$ where $R = [-a, a]^m$ for some $a \in (0, 1)$, then $\bar{p}^n(a_i) = 1 - [1 - (2 \left(\frac{2-a}{2} \right)^n - (1-a)^n)]^m$.*

A.2 Proof of Claim 2

I first demonstrate the following lemma, which characterizes the probabilities \underline{p}^n and \bar{p}^n .

Lemma 3. *For every $n \geq 1$,*

$$\underline{p}^n = 1 - \Phi \left(z_\alpha - \frac{\beta - 1}{\sigma} \sqrt{\frac{n^2 - 1}{12}} \right)$$

while

$$\bar{p}^n = 1 - \Phi \left(-z_\alpha - \frac{\beta - 1}{\sigma} \sqrt{\frac{n^2 - 1}{12}} \right)$$

where $z_\alpha = -\Phi^{-1}(\alpha/2)$ with Φ denoting the CDF of the standard normal distribution.

Since $\beta > 1$ by assumption, both expressions are decreasing in σ and increasing in n . Thus Claim 2 follows.

Towards this lemma, I first prove the following intermediate result:

Lemma 4. Write $T_i^\mathcal{C}$ for the set of player i types with common certainty in the event that all players have first-order beliefs that assign probability 1 to \mathcal{C} .

- (a) The strong policy is rationalizable for all types $t_i \in T_i^\mathcal{C}$ if and only if $\mathcal{C} \subseteq [1, \infty)$.
- (b) The strong policy is rationalizable for some type $t_i \in T_i^\mathcal{C}$ if and only if $\mathcal{C} \cap [1, \infty) \neq \emptyset$.

Proof. (a) Necessity of $\mathcal{C} \subseteq [1, \infty)$ is straightforward, as otherwise there exists some $\beta' \in \mathcal{C} \setminus [1, \infty)$ and the strong policy is not rationalizable for the type with common certainty that $\beta = \beta'$. Suppose $\mathcal{C} \subseteq [1, \infty)$ and choose any $t_i \in T_i^\mathcal{C}$. For each $\beta \in \mathcal{C}$, $u_i(\text{strong}, \text{strong}, \beta) = -1$ while $u_i(\text{weak}, \text{strong}, \beta) = -\beta \leq -1$. So

$$\int u_i(\text{strong}, \text{strong}, \beta) t_i^1(\beta) d\beta = -1 \geq \int u_i(\text{weak}, \text{strong}, \beta) t_i^1(\beta) d\beta$$

where t_i^1 denotes the first-order belief of type t_i . Thus the family of sets (R_1, R_2) with $R_1 = R_2 = \{\text{strong}\}$ are closed under best reply, and rationalizability of the strong policy follows.

(b) Suppose $\mathcal{C} \cap [1, \infty) = \emptyset$. Then for every $\beta \in \mathcal{C}$,

$$u_i(\text{strong}, \text{strong}, \beta) = -1 \leq -\beta = u_i(\text{weak}, \text{strong}, \beta).$$

So the strong policy is strictly dominated (and hence not rationalizable) for player i given any type $t_i \in T_i^\mathcal{C}$. If instead $\mathcal{C} \cap [1, \infty) \neq \emptyset$, then the strong policy is rationalizable for any type with common certainty in some β in this intersection. So the strong policy is rationalizable for at least one type $t_i \in T_i^\mathcal{C}$, as desired. \square

I now prove Lemma 3.

Proof. Using standard results for ordinary least-squares (Hastie et al., 2009), the distribution of the OLS estimator $\hat{\beta}$ is

$$\hat{\beta} \sim \mathcal{N}\left(\beta, \frac{\sigma^2}{\frac{1}{n} \sum_{t=1}^n (t - \bar{t})^2}\right)$$

where $\bar{t} = \frac{1}{n} \sum_{t=1}^n t$. Since

$$\begin{aligned} \frac{1}{n} \sum_{t=1}^n (t - \bar{t})^2 &= \frac{1}{n} \left(\sum_{t=1}^n t^2 - 2\bar{t} \sum_{t=1}^n t + \sum_{t=1}^n \bar{t}^2 \right) \\ &= \frac{(n+1)(2n+1)}{6} - \frac{(n+1)^2}{2} + \left(\frac{n+1}{2}\right)^2 = \frac{(n^2-1)}{12} \end{aligned}$$

we can simplify the variance of $\hat{\beta}$ to $\frac{12\sigma^2}{n^2-1}$. The $(1 - \alpha)$ -confidence interval for β given data \mathbf{z}_n is thus

$$\mathcal{C}(\mathbf{z}_n) = \left[\hat{\beta}(\mathbf{z}_n) - z_\alpha \cdot \sigma \cdot \sqrt{\frac{12}{n^2-1}}, \hat{\beta}(\mathbf{z}_n) + z_\alpha \cdot \sigma \cdot \sqrt{\frac{12}{n^2-1}} \right] \quad (8)$$

where $\hat{\beta}(\mathbf{z}_n)$ is the OLS estimate of β given the data \mathbf{z}_n , and $z_\alpha = -\Phi^{-1}(\alpha/2)$ is the critical value associated with the $(1 - \alpha)$ -confidence level. The probability that the interval in (8) is contained in $[1, \infty)$ is

$$\Pr \left(\hat{\beta}(\mathbf{z}_n) > 1 + z_\alpha \cdot \sigma \cdot \sqrt{\frac{12}{n^2-1}} \right).$$

which is in turn equal to

$$1 - \Phi \left(1.96 - \frac{\beta - 1}{\sigma} \sqrt{\frac{n^2-1}{12}} \right). \quad (9)$$

By Part (a) of Lemma 4, \underline{p}^n is equal to (9), delivering the first part of the lemma.

The probability that the interval in (8) has nonempty intersection with $[1, \infty)$ is given by

$$\Pr \left(\hat{\beta}(\mathbf{z}_n) > 1 - z_\alpha \cdot \sigma \cdot \sqrt{\frac{12}{n^2-1}} \right)$$

which is equal to

$$1 - \Phi \left(-z_\alpha - \frac{\beta - 1}{\sigma} \sqrt{\frac{n^2-1}{12}} \right) \quad (10)$$

By Part (b) of Lemma 4, \bar{p}^n is equal to (10), concluding the proof. \square

B Proofs for Main Results (Sections 4.1 and 4.2)

B.1 Proof of Theorem 1 Part (a)

Recall from Section 2.1 that Θ and U are endowed with the sup-norm, and the map $g : \Theta \rightarrow U$ is Lipschitz continuous with Lipschitz constant K . The set of probability measures $\Delta(\Theta)$ is endowed with the Prokhorov metric d_P . The *Wasserstein distance* on $\Delta(\Theta)$ is

$$d_W(\nu, \nu') = \sup \left\{ \int h d\nu - \int h d\nu' : \|h\|_L \leq 1 \right\}$$

where $\|h\|_L$ is the Lipschitz constant of the function $h : \Theta \rightarrow \mathbb{R}$.

Lemma 5. *Fix any player i , action $a_i \in A_i$, mixed strategy $\alpha_i \in \Delta(A_i)$, and set $R_{-i} \subseteq A_{-i}$. Let $a_{-i}(\theta) : \Theta \rightarrow \Delta(A_{-i})$ be any function satisfying*

$$a_{-i}(\theta) \in \operatorname{argmax}_{a_{-i} \in R_{-i}} (u_i(a_i, a_{-i}, \theta) - u_i(\alpha_i, a_{-i}, \theta)) \quad \forall \theta \in \Theta$$

and define $h : \Theta \rightarrow \mathbb{R}$ by

$$h(\theta) = u_i(a_i, a_{-i}(\theta), \theta) - u_i(\alpha_i, a_{-i}(\theta), \theta).$$

Then, the function h is Lipschitz continuous with Lipschitz constant $2K$.

Proof. Choose any $\theta, \theta' \in \Theta$, and without loss of generality, suppose $h(\theta) \geq h(\theta')$. Then

$$\begin{aligned} |h(\theta) - h(\theta')| &= |(u(a_i, a_{-i}(\theta), \theta) - u(\alpha_i, a_{-i}(\theta), \theta)) - \\ &\quad (u(a_i, a_{-i}(\theta'), \theta') - u(\alpha_i, a_{-i}(\theta'), \theta'))| \\ &\leq |(u(a_i, a_{-i}(\theta), \theta) - u(\alpha_i, a_{-i}(\theta), \theta)) - \\ &\quad (u(a_i, a_{-i}(\theta), \theta') - u(\alpha_i, a_{-i}(\theta), \theta'))| \\ &\leq |u(a_i, a_{-i}(\theta), \theta) - u(a_i, a_{-i}(\theta), \theta')| + \\ &\quad |u(\alpha_i, a_{-i}(\theta), \theta) - u(\alpha_i, a_{-i}(\theta), \theta')| \\ &\leq 2\|g(\theta) - g(\theta')\|_\infty \leq 2K\|\theta - \theta'\|_\infty \end{aligned}$$

using in the final inequality that $g : \Theta \rightarrow U$ has Lipschitz constant K . □

Below, let F^ϵ denote the ϵ -neighborhood of the set F .

Lemma 6. *Suppose a_i is δ -strictly rationalizable for player i of type t_i^∞ , where $\delta > 0$. Let \mathcal{B} be any subset of $\{\mu^\infty\}^{\delta/(2K\xi)}$, where K is the Lipschitz constant of $g : \Theta \rightarrow U$, and ξ is as defined in (6). Then, a_i is rationalizable for all types $t_i \in T_i^{\mathcal{B}}$.³³*

Proof. Fix $\epsilon > 0$, and consider an arbitrary set $\mathcal{B} \subseteq \{\mu^\infty\}^\epsilon$. I will show that a_i is rationalizable for all types $t_i \in T_i^{\mathcal{B}}$ when ϵ is sufficiently small.

To show this, I use Proposition 1 from [Chen et al. \(2010\)](#):³⁴

Proposition 4 ([Chen et al. \(2010\)](#)). *For each $k \geq 1$, player $i \in \mathcal{I}$, type $t_i \in T_i$, and action $a_i \in A_i$, we have $a_i \in S_i^k[t_i]$ if and only if for each $\alpha_i \in \Delta(A_i \setminus \{a_i\})$, there exists a measurable $\sigma_{-i} : \Theta \times T_{-i} \rightarrow \Delta(A_{-i})$ with*

$$\text{supp } \sigma_{-i}(\theta, t_{-i}) \subseteq S_{-i}^{k-1}[t_{-i}] \quad \forall (\theta, t_{-i}) \in \Theta \times T_{-i}$$

such that

$$\int_{\Theta \times T_{-i}} [u_i(a_i, \sigma_{-i}(\theta, t_{-i}), \theta) - u_i(\alpha_i, \sigma_{-i}(\theta, t_{-i}), \theta)] t_i [d\theta \times dt_{-i}] \geq 0$$

By assumption, there is a $\delta \in \mathbb{R}_{++}$ such that a_i is δ -strictly rationalizable for player i of type t_i^∞ . This implies that there exists a family of sets $(R_j)_{j \in \mathcal{I}} \subseteq \prod_{j \in \mathcal{I}} A_j$, where $a_i \in R_i$, and for every

³³[Chen et al. \(2010\)](#) demonstrate a similar result for finite state spaces Θ (see their Proposition 2). I use ideas from their proof here, but consider a more general environment, replacing finiteness of Θ with Lipschitz continuity on $g : \Theta \rightarrow U$.

³⁴Proposition 1 from [Chen et al. \(2010\)](#) characterizes γ -rationalizability for arbitrary $\gamma \in \mathbb{R}$. For the purposes of this proof, it is sufficient to set $\gamma = 0$.

$a_j \in R_j$ there exists a $\sigma_{-j}^\infty : \Theta \rightarrow \Delta(A_{-j})$ satisfying

$$\text{supp } \sigma_{-j}^\infty(\theta) \subseteq R_{-j} \quad \forall \theta \in \Theta$$

and

$$\int_{\Theta} u_j(a_j, \sigma_{-j}^\infty(\theta), \theta) d\mu^\infty - \int_{\Theta} u_j(a'_j, \sigma_{-j}^\infty(\theta), \theta) d\mu^\infty \geq \delta \quad \forall a'_j \neq a_j \quad (11)$$

I will show that for each $k \geq 1$, player j , type $t_j \in T_j^{\mathcal{B}}$, action $a_j \in R_j$, and mixed strategy $\alpha_j \in \Delta(A_j \setminus \{a_j\})$, there exists a measurable $\sigma_{-j} : \Theta \times T_{-j}^{\mathcal{B}} \rightarrow \Delta(A_{-j})$ with

$$\text{supp } \sigma_{-j}(\theta, t_{-j}) \subseteq R_{-j} \quad \forall (\theta, t_{-j}) \in \Theta \times T_{-j}^{\mathcal{B}}$$

and

$$\int_{\Theta \times T_{-j}^{\mathcal{B}}} [u_j(a_j, \sigma_{-j}(\theta, t_{-j}), \theta) - u_j(\alpha_j, \sigma_{-j}(\theta, t_{-j}), \theta)] t_j [d\theta \times dt_{-j}] \geq 0. \quad (12)$$

Since $a_i \in R_i$ by design, it follows from Proposition 4 that for any type $t_i \in T_i^{\mathcal{B}}$, the action $a_i \in S_i^k[t_i]$ for every k , and hence $a_i \in S_i^\infty[t_i]$, as desired.

Fix an arbitrary player j , $a_j \in R_j$, type $t_j \in T_j^{\mathcal{B}}$, and $\alpha_j \in \Delta(A_j \setminus \{a_j\})$. Define $a_{-j} : \Theta \rightarrow A_{-j}$ to satisfy

$$a_{-j}(\theta) \in \underset{a_{-j} \in R_{-j}}{\text{argmax}} (u_j(a_j, a_{-j}, \theta) - u_j(\alpha_j, a_{-j}, \theta)) \quad \forall \theta \in \Theta$$

and define $\sigma_{-j} : \Theta \times T_{-j}^{\mathcal{B}} \rightarrow \Delta(A_{-j})$ so that each $\sigma_{-j}(\theta, t_{-j})$ is a point mass at $a_{-j}(\theta)$. Then by definition

$$\text{supp } \sigma_{-j}(\theta, t_{-j}) \subseteq R_{-j} \quad \forall (\theta, t_{-j}) \in \Theta \times T_{-j}^{\mathcal{B}}.$$

Further define

$$h(\theta) := u_j(a_j, a_{-j}(\theta), \theta) - u_j(\alpha_j, a_{-j}(\theta), \theta) \quad \forall \theta \in \Theta.$$

For notational ease, write $\nu \in \Delta(\Theta)$ for the first-order belief of type t_j . Then

$$\begin{aligned} & \int_{\Theta \times T_{-j}^{\mathcal{B}}} u_j(a_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] - \int_{\Theta \times T_{-j}^{\mathcal{B}}} u_j(\alpha_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] \\ &= \int_{\Theta} u_j(a_j, a_{-j}(\theta), \theta) \nu [d\theta] - \int_{\Theta} u_j(\alpha_j, a_{-j}(\theta), \theta) \nu [d\theta] = \int_{\Theta} h(\theta) \nu [d\theta] \end{aligned}$$

so the desired condition in (12) follows if we can show that $\int_{\Theta} h(\theta) \nu [d\theta] \geq 0$.

By Lemma 5, the function $h : \Theta \rightarrow \mathbb{R}$ has Lipschitz constant $2K$, so

$$\left| \int_{\Theta} h(\theta) d\nu - \int_{\Theta} h(\theta) d\mu^\infty \right| \leq 2K \cdot d_W(\nu, \mu^\infty)$$

where d_W is the Wasserstein distance on $\Delta(\Theta)$. This implies

$$\int_{\Theta} h(\theta) d\nu \geq \int_{\Theta} h(\theta) d\mu^\infty - 2K \cdot d_W(\nu, \mu^\infty).$$

Applying Theorem 2 in [Gibbs and Su \(2002\)](#), $d_W(\nu, \mu^\infty) \leq \xi \cdot d_P(\nu, \mu^\infty)$, where d_P is the Prokhorov distance on $\Delta(\Theta)$ and ξ is as defined in (6). So

$$\int_{\Theta} h(\theta) d\nu \geq \int_{\Theta} h(\theta) d\mu^\infty - 2K\xi \cdot d_P(\nu, \mu^\infty) \quad (13)$$

It follows from the inequality in (11) that

$$\int_{\Theta} h(\theta) d\mu^\infty = \int_{\Theta} u_j(a_j, \sigma_{-j}^\infty(\theta), \theta) d\mu^\infty - \int_{\Theta} u_j(\alpha_j, \sigma_{-j}^\infty(\theta), \theta) d\mu^\infty \geq \delta,$$

so (13) implies

$$\int_{\Theta} h(\theta) d\nu \geq \delta - 2K\xi \cdot d_P(\nu, \mu^\infty).$$

Finally, by assumption that $t_j \in T_j^{\mathcal{B}}$ for some $\mathcal{B} \subseteq \{\mu^\infty\}^\epsilon$, the Prokhorov distance between the first-order belief of type t_j and the limiting belief ν^∞ is $d_P(\nu, \mu^\infty) \leq \epsilon$. So

$$\int_{\Theta} h(\theta) d\nu \geq \delta - 2K\xi\epsilon.$$

It follows that $\epsilon \leq \delta/(2K\xi)$ is a sufficient condition for the constructed σ_{-j} to satisfy the desired condition in (12). \square

Since a_i is strictly rationalizable for type t_i^∞ (by assumption), there exists a $\delta \in \mathbb{R}_{++}$ for which a_i is δ -strictly rationalizable. Assumption 1 implies that

$$\lim_{n \rightarrow \infty} P^n \left(\left\{ \mathbf{z}_n \mid \sup_{\mu \in \mathcal{M}} d_P(\mu(\mathbf{z}_n), \mu^\infty) \leq \epsilon \right\} \right) = 0 \quad \forall \epsilon > 0.$$

which further implies

$$\lim_{n \rightarrow \infty} P^n \left(\left\{ \mathbf{z}_n \mid \mathcal{B}(\mathbf{z}_n) \subseteq \{\theta^\infty\}^{\delta/2K\xi} \right\} \right) = 0 \quad (14)$$

By Lemma 6,

$$\underline{p}^n(a_i) \geq P^n \left(\left\{ \mathbf{z}_n \mid \mathcal{B}(\mathbf{z}_n) \subseteq \{\theta^\infty\}^{\delta/2K\xi} \right\} \right) \quad \forall n \geq 1$$

so from (14) we can directly conclude that $\underline{p}^n(a_i) \rightarrow 1$. Theorem 1 Part (a) follows.

B.2 Proof of Theorem 1 Part (b)

I begin by reviewing definitions from [Chen et al. \(2010\)](#) that will be used in the proof. For each player i , let $X_i^0 = \Theta$, and recursively for $k \geq 1$, define $X_i^k = \Theta \times \prod_{j \neq i} \Delta(X_j^{k-1})$.³⁵ The space of k -th order beliefs for player i is defined $T_i^k := \Delta(X_i^{k-1})$, noting that each $T_i^k = \Delta(\Theta \times T_{-i}^{k-1})$. The

³⁵The sets X^k defined in Section 2.1 can be identified with the sets X^k defined in this way.

uniform-weak metric on the universal type space T_i^* is

$$d_i^{UW}(s_i, t_i) = \sup_{k \geq 1} d_i^k(s_i, t_i) \quad \forall s_i, t_i \in T_i^*$$

where d^0 is the supremum norm on Θ and recursively for $k \geq 1$, d_i^k is the Prokhorov distance on $\Delta(\Theta \times T_{-i}^{k-1})$ induced by the metric $\max\{d^0, d_{-i}^{k-1}\}$ on $\Theta \times T_{-i}^{k-1}$.³⁶ The uniform-weak topology on the universal type space is the metric topology induced by d_i^{UW} .

Lemma 7. *Let \mathcal{B} be a subset of $\{\mu^\infty\}^\epsilon$ and choose any $s_i \in T_i^\mathcal{B}$. Then $d_i^{UW}(t_i^\infty, s_i) \leq \epsilon$.*

Proof. For simplicity of notation, write t_i for t_i^∞ . It will be useful to define

$$T_i^{\mathcal{B},k} = \left\{ s_i^k \in T_i^k \mid s_i \in T_i^\mathcal{B} \right\}$$

for the set of all k -th order beliefs that are consistent with some type $s_i \in T_i^\mathcal{B}$.³⁷ I will show that

$$d_P\left(T_i^{\mathcal{B},k}, t_i^k\right) := \sup_{s_i^k \in T_i^{\mathcal{B},k}} d_P(s_i^k, t_i^k) \leq \epsilon \quad \forall k \geq 1 \quad (15)$$

from which the desired lemma directly follows.

By construction, $T_i^{\mathcal{B},1} = \mathcal{B}$, so the assumption $\mathcal{B} \subseteq \{\mu^\infty\}^\epsilon$ immediately implies (15) for $k = 1$. Proceed by induction. Suppose $d_P\left(T_i^{\mathcal{B},k}, t_i^k\right) \leq \epsilon$, and consider any measurable set $E \subseteq T^k$. If $t_i^k \in E$, then $t_i^{k+1}(E) = 1$ by definition of t_i . Also,

$$s_i^{k+1}(E^\epsilon) \geq s_i^{k+1}\left(\{t_i^k\}^\epsilon\right) \geq s_i^{k+1}(T_i^{\mathcal{B},k}) = 1$$

where the second inequality follows from the inductive hypothesis, and the final inequality follows by assumption that $s_i \in T_i^\mathcal{B}$. So

$$t_i^{k+1}(E) \leq s_i^{k+1}(E^\epsilon) + \epsilon. \quad (16)$$

If $t_i^k \notin E$, then $t_i^{k+1}(E) = 0$ (again by definition of t_i), so (16) follows trivially. Thus

$$d_i^{k+1}(t_i, s_i) = \inf\{\delta \mid t_i^{k+1}(E) \leq s_i^{k+1}(E^\delta) + \delta \quad \forall \text{ measurable } E \subseteq T_i^k\} \leq \epsilon$$

and so $d_i^{UW}(t_i, s_i) = \sup_{k \geq 1} d_i^k(t_i, s_i) \leq \epsilon$ as desired. \square

Lemma 7 implies the subsequent corollary.

Corollary 3. *Suppose Assumption 1 holds. Consider any sequence $\mathbf{z} \in \mathcal{Z}^\infty$ satisfying*

$$\lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{M}} d_P(\mu(\mathbf{z}_n), \mu^\infty) = 0 \quad (17)$$

³⁶This definition is slightly modified from Chen et al. (2010), where d^0 was the discrete metric on Θ . The change reflects the difference that Θ was taken to be a finite set in Chen et al. (2010), while it is a compact and convex subset of Euclidean space here.

³⁷Here and elsewhere, t_i^k denotes the k -th order belief of type t_i .

and choose any sequence of types $(s_i^n)_{n=1}^\infty$ with $s_i^n \in T_i^{\mathcal{B}(\mathbf{z}_n)}$ for each $n \geq 1$. Then

$$\lim_{n \rightarrow \infty} d_i^{UW}(t_i^\infty, s_i^n) = 0.$$

Now we will complete the proof of Theorem 1 Part (b). By Assumption 1, there is a set $\mathcal{Z}^* \subseteq \mathcal{Z}^\infty$ of P -measure 1 such that

$$\lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{M}} d_P(\mu(\mathbf{z}_n), \mu^\infty) = 0 \quad \forall \mathbf{z} \in \mathcal{Z}^* \quad (18)$$

Suppose towards contradiction that $\bar{p}^n(a_i) \not\rightarrow 0$. Then, there is a set $\hat{\mathcal{Z}} \subseteq \mathcal{Z}^\infty$ with strictly positive P -measure such that for every $\mathbf{z} \in \hat{\mathcal{Z}}$, there is a sequence of types $(t_i^n(\mathbf{z}))_{n=1}^\infty$ where $t_i^n(\mathbf{z}) \in T_i^{\mathcal{B}(\mathbf{z}_n)}$ for every $n \geq 1$, and $a_i \in S_i^\infty[t_i^n(\mathbf{z})]$ for all n sufficiently large.

But since \mathcal{Z}^* has P -measure 1, it must be that $\hat{\mathcal{Z}} \cap \mathcal{Z}^* \neq \emptyset$. Choose any \mathbf{z} from this intersection. Then, Lemma 3 and the display in (18) imply that $t_i^n(\mathbf{z}) \rightarrow t_i^\infty$ in the uniform-weak topology. But rationalizability is upper hemi-continuous in the uniform-weak topology (Theorem 1, Chen et al. (2010)). So $a_i \notin S_i^\infty[t_i^\infty]$ implies $a_i \notin S_i^\infty[t_i^n(\mathbf{z})]$ for infinitely many n , a contradiction.

B.3 Proof of Proposition 1

By assumption, a_i is strictly rationalizable for type t_i^∞ , so $\delta^\infty > 0$. Applying Lemma 6,

$$\begin{aligned} \underline{p}^n(a_i) &\geq P^n(\{\mathbf{z}_n : \mathcal{B}(\mathbf{z}_n) \subseteq \{\mu^\infty\}^{\delta^\infty/(2K\xi)}\}) \\ &= P^n\left(\left\{\mathbf{z}_n : \sup_{\mu \in \mathcal{M}} d_P(\mu(\mathbf{z}_n), \mu^\infty) \leq \delta^\infty/(2K\xi)\right\}\right) \\ &\geq 1 - \frac{2K\xi}{\delta^\infty} \mathbb{E}\left(\sup_{\mu \in \mathcal{M}} d_P(\mu(Z^n), \mu^\infty)\right) \end{aligned}$$

using Markov's inequality in the final line.

B.4 Proof of Proposition 2

Suppose a_i is strictly rationalizable for player i in the complete information game θ^∞ , and let δ^∞ be as defined in (5). Then, there exists a family of sets $(R_j)_{j \in \mathcal{I}}$ with $a_i \in R_i$, where for each player j and action $a_j \in R_j$, there is a mixed strategy $\sigma_{-j} \in \Delta(A_{-j})$ satisfying $\sigma_{-j}[R_{-j}] = 1$, and

$$u_i(a_i, \sigma_{-j}, \theta^\infty) - u_i(a'_i, \sigma_{-j}, \theta^\infty) \geq \delta \quad \forall a'_i \neq a_i.$$

Now consider an arbitrary set $\mathcal{C} \subseteq \{\theta^\infty\}^\epsilon$ and a type t_i with common certainty in the event that every player's first-order belief assigns probability 1 to \mathcal{C} . Write $\nu \in \Delta(\Theta)$ for the first-order

belief of type t_i . For any action $a'_j \neq a_j$,

$$\begin{aligned}
\int u_j(a_j, \sigma_{-j}, \theta) d\nu - \int u_j(a'_j, \sigma_j, \theta) d\nu &= \int u_j(a_j, \sigma_{-j}, \theta) d\nu - \int u_j(a_j, \sigma_{-j}, \theta) d\mu^\infty \\
&\quad + \int u_j(a_j, \sigma_{-j}, \theta) d\mu^\infty - \int u_j(a'_j, \sigma_{-j}, \theta) d\mu^\infty \\
&\quad + \int u_j(a'_j, \sigma_{-j}, \theta) d\mu^\infty - \int u_j(a'_j, \sigma_j, \theta) d\nu \\
&\geq \int u_j(a_j, \sigma_{-j}, \theta) d\mu^\infty - \int u_j(a'_j, \sigma_{-j}, \theta) d\mu^\infty \\
&\quad - \left| \int u_j(a_j, \sigma_{-j}, \theta) d\nu - \int u_j(a_j, \sigma_{-j}, \theta) d\mu^\infty \right| \\
&\quad - \left| \int u_j(a'_j, \sigma_{-j}, \theta) d\nu - \int u_j(a'_j, \sigma_j, \theta) d\mu^\infty \right| \\
&\geq \delta - 2K \cdot d_P(\nu, \mu^\infty) \geq \delta - 2K\epsilon
\end{aligned}$$

using in the penultimate inequality that $g : \Theta \rightarrow U$ has Lipschitz constant K . Since this bound on the payoff difference holds across all actions $a'_j \neq a_j$, the action a_j is a best reply to belief ν whenever $\epsilon \leq \delta/(2K)$.

This allows us to construct the lower bound

$$\begin{aligned}
\underline{p}^n(a_i) &\geq Q^n \left(\left\{ \mathbf{z}_n : \mathcal{C}(\mathbf{z}_n) \subseteq \{\theta^\infty\}^{\delta^\infty/(2K)} \right\} \right) \\
&= Q^n \left(\left\{ \mathbf{z}_n : \sup_{\theta \in \mathcal{C}(\mathbf{z}_n)} \|\theta - \theta^\infty\|_\infty \leq \delta^\infty/(2K) \right\} \right) \\
&\geq 1 - \frac{2K}{\delta^\infty} \mathbb{E} \left(\sup_{\theta \in \mathcal{C}(\mathbf{z}_n)} \|\theta - \theta^\infty\|_\infty \right)
\end{aligned}$$

using Markov's inequality in the final line.

O For Online Publication

O.1 Example 1

Fix an arbitrary $(\pi, q) \in (0, 1) \times (1/2, 1)$. Given data \mathbf{z}_n , the posterior belief $\mu_{\pi, q}(\mathbf{z}_n)$ assigns probability

$$\hat{v}(\pi, q, \mathbf{z}_n) := 1 / \left(1 + \frac{1 - \pi}{\pi} \left(\frac{1 - q}{q} \right)^{n(2\bar{z}_n - 1)} \right) \quad (19)$$

to $v = 1$, where $\bar{z}_n = \frac{1}{n} \sum_{n'=1}^n z_{n'}$ denotes the average realization in the sequence \mathbf{z}_n .

Suppose without loss that $v = 1$, and let $q^* \in (1/2, 1)$ be the true frequency of $z = 1$. By the strong Law of Large numbers, there is a measure 1 set of sequences \mathcal{Z}^* satisfying

$$\lim_{n \rightarrow \infty} \left(\frac{1}{n} \sum_{n'=1}^n z_{n'} \right) = q^*$$

for every $\mathbf{z} = (z_1, z_2, \dots) \in \mathcal{Z}^*$. The expression in (19) converges to 1 on this set for every $(\pi, q) \in (0, 1) \times (1/2, 1)$. So at $n = \infty$, players share the common limiting belief μ^∞ , which assigns probability 1 to $v = 1$. Since entering is not rationalizable for the Seller given common certainty in the event in that all players assign probability 1 to $v = 1$, it follows that $\underline{p}(\infty) = \bar{p}(\infty) = 0$.

I show next that the probability $\bar{p}(n)$ converges to 1 as $n \rightarrow \infty$. Fix an arbitrary n , and define $\mathcal{Z}_n^\dagger = \{\mathbf{z}_n \mid \bar{z}_n > 1/2\}$ to be the set of length- n sequences with majority realizations of $z = 1$. For every $\mathbf{z}_n \in \mathcal{Z}_n^\dagger$, the expression $((1 - q)/q)^{n(2z_n - 1)}$ is bounded between $1/2$ and 1 on the domain $q \in (1/2, 1)$, while the image of $(1 - \pi)/\pi$ is all of \mathbb{R}_+ . Thus, the display in (19) ranges from zero to 1; that is,

$$\{\hat{v}(\pi, q, \mathbf{z}_n) : \pi \in (0, 1), q \in (1/2, 1)\} = (0, 1) \quad \forall \mathbf{z}_n \in \mathcal{Z}_n^\dagger.$$

It follows that for every $\mathbf{z}_n \in \mathcal{Z}_n^\dagger$, there exist pairs $(\pi, q), (\pi', q') \in (0, 1) \times (1/2, 1)$ satisfying $\hat{v}(\pi, q, \mathbf{z}_n) < p < \hat{v}(\pi', q', \mathbf{z}_n)$. Entering is rationalizable for the Seller with a type that assigns probability $\hat{v}(\pi', q', \mathbf{z}_n)$ to the high value, and which assigns probability 1 to the Buyer assigning probability $\hat{v}(\pi, q, \mathbf{z}_n)$ to the high value. So entering is partially $\mathcal{B}(\mathbf{z}_n)$ -rationalizable for every $\mathbf{z}_n \in \mathcal{Z}_n^\dagger$, implying $\bar{p}^n(a_i) \geq P^n(\mathcal{Z}_n^\dagger)$.

Again by the law of large numbers, the measure of datasets with majority realizations of $z = 1$ converges to 1 as $n \rightarrow \infty$; that is, $P^n(\mathcal{Z}_n^\dagger) \rightarrow 1$. So $\lim_{n \rightarrow \infty} \bar{p}^n(a_i) = 1$, as desired.

O.2 Proof of Corollary 1

First observe that $\delta^\infty = \beta - 1$, since the action *Strong* is δ -strictly rationalizable for every $\delta < \beta - 1$ and not for any $\delta \geq \beta - 1$. It remains to determine $\mathbb{E} \left[\sup_{\theta' \in \mathcal{C}(Z^n)} \|\theta' - \theta^\infty\|_\infty \right]$. Write \bar{Z}_n for the

(random) empirical mean of n signal realizations, and $\hat{\beta}_x(\mathbf{z}_n)$ for the expectation of β given signals \mathbf{z}_n and prior $\beta \sim \mathcal{N}(x, 1)$. Then, using standard formulas for updating to Gaussian signals:

$$\mathbb{E} \left(\sup_{x \in [-\eta, \eta]} |\beta - \hat{\beta}_x(Z^n)| \right) = \mathbb{E} \left[\max_{x \in [-\eta, \eta]} \left(\left| \beta - \frac{x + n\bar{Z}_n}{n+1} \right| \right) \right]$$

We can further bound the RHS as follows:

$$\begin{aligned} \mathbb{E} \left[\max_{x \in [-\eta, \eta]} \left(\left| \beta - \frac{x + n\bar{Z}_n}{n+1} \right| \right) \right] &\leq \mathbb{E} \left(\left| \beta - \frac{n\bar{Z}_n}{n+1} \right| \right) + \max_{x \in [-\eta, \eta]} \left| \frac{x}{n+1} \right| \\ &= \mathbb{E} \left(\left| \beta - \frac{n\bar{Z}_n}{n+1} \right| \right) + \eta/(n+1) \\ &\leq \mathbb{E} (|\beta - \bar{Z}_n|) + \mathbb{E} \left(\frac{\bar{Z}_n}{n+1} \right) + \eta/(n+1) \\ &= \sqrt{\frac{2}{n\pi}} + \frac{\beta + \eta}{n+1} \end{aligned}$$

using in the final line the expected absolute deviation of the empirical mean of n observations from a Gaussian distribution (Geary, 1935). Finally, the map $g : \Theta \rightarrow U$ has Lipschitz constant 1. Applying Proposition 2, we have the desired bound.

O.3 Examples Related to Theorem 1

Part (a) of Theorem 1 provides a sufficient condition for the plausibility interval $[\underline{p}^n(a_i), \bar{p}^n(a_i)]$ to converge to certainty— a_i is strictly rationalizable for type t_i^∞ —and Part (b) of Theorem 1 provides a necessary condition— a_i is rationalizable for type t_i^∞ . The condition that a_i is strictly rationalizable is not necessary, as I demonstrate in Section O.3.1, and the condition that a_i is rationalizable is not sufficient, as I demonstrate in Section O.3.2.

In each of these examples, I assume (as in Section 4.2.1) that the limiting belief μ^∞ is degenerate at a limiting parameter θ^∞ , and players have common certainty of shrinking neighborhoods of this parameter. That is, for every realization \mathbf{z}_n , players have common certainty in the event that players have first-order beliefs with support on $\mathcal{C}(\mathbf{z}_n)$, where the support sets $\mathcal{C}(\mathbf{z}_n)$ satisfy Assumption 4.

O.3.1 Strict Rationalizability is Not Necessary

Consider the following complete information game

$$\begin{array}{cc} & a_3 & a_4 \\ a_1 & \theta, 0 & \theta, 0 \\ a_2 & 0, 0 & 0, 0 \end{array}$$

and suppose that the limiting belief is degenerate at $\theta^\infty = 1$. Then, the action a_1 is *strictly dominant* for player 1 in the limiting complete information game, and also for all types with common certainty in the event that players have first-order beliefs with support on a small enough neighborhood of θ^∞ . So Assumption 1 implies $\lim_{n \rightarrow \infty} [\underline{p}^n(a_i), \bar{p}^n(a_i)] = \{1\}$. But action a_1 is not *strictly rationalizable* for type t_i^∞ .

O.3.2 Rationalizability is not Sufficient

I show next that rationalizability of a_i for type t_i^∞ is not sufficient for the analyst's plausibility interval for a_i to converge to certainty. Section O.3.3 provides a simple example to this effect. Define Θ^{a_i} to be the set of parameter values θ such that a_i is rationalizable for player i in the complete information game indexed to θ . If a_i is on the boundary of Θ^{a_i} , then common certainty of shrinking neighborhoods around θ^∞ does not guarantee rationalizability of a_i . More surprisingly, common certainty in arbitrarily small open sets within the *interior* of Θ^{a_i} also does not guarantee rationalizability of a_i , and I provide an example of this in Section O.3.4. (See also the working paper of [Chen and Takahashi \(2020\)](#) for a nice two-player example to this effect.)

O.3.3 θ^∞ is on the Boundary of Θ^{a_i}

Consider the following two-player game, parametrized by $\theta \in [\underline{\theta}, \bar{\theta}]$ for some $\underline{\theta} < 0 < \bar{\theta}$:

	a	b	
a	θ, θ	$0, 0$	
b	$0, 0$	$1, 1$	

Suppose that the limiting parameter $\theta^\infty = 0$, so that a is rationalizable in the limiting complete information game, but not strictly rationalizable. It is straightforward to see that common certainty of shrinking neighborhoods of θ^∞ does not guarantee rationalizability of action a , as the type with common certainty of any $\theta' < 0$ considers a to be strictly dominated.

O.3.4 θ^∞ is in the Interior of Θ^{a_i}

But even if θ^∞ is not on the boundary of the set Θ^{a_i} , it may be that common certainty of a shrinking neighborhood of θ^∞ does not guarantee rationalizability of a_i . Consider the following four-player game. Players 1 and 2 choose between actions in $\{a, b\}$, and player 3 chooses between matrices

from $\{l, r\}$. Their payoffs are:

$$\begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 1, 1, 0 & 0, 0, 0 \\ 0, 0, 0 & 0, 0, 0 \end{array} \\
 & (l)
 \end{array}
 \qquad
 \begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 0, 0, 0 & 0, 0, 0 \\ 0, 0, 0 & 1, 1, 0 \end{array} \\
 & (r)
 \end{array}
 \tag{20}$$

A fourth player predicts whether players 1 and 2 chose matching actions or mis-matching actions. He receives a payoff of 1 if he predicts correctly (and 0 otherwise).³⁸ Player 4's action does not affect the payoffs of the other three players.

Let the state space $\Theta = \mathbb{R}^{64}$ be the set of all payoff matrices given these actions, where the payoffs described above are a particular θ . Match is clearly rationalizable for player 4 at θ ; it is also rationalizable for player 4 on a neighborhood of θ (in the Euclidean metric).³⁹

Nevertheless, I will show existence of a sequence of types for player 4 with common certainty in increasingly small neighborhoods of θ , given which Match fails to be rationalizable. Along this sequence, player 4 believes that a is uniquely rationalizable for player 1, while b is uniquely rationalizable for player 2, so the action Match is strictly dominated.

Define θ_ε^1 to be the following perturbation of the payoff matrix θ (with player 4's payoffs unchanged):

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

Let θ_ε^2 correspond to the following payoff matrix (again with player 4's payoffs unchanged):

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

³⁸In more detail: player 4 chooses between $\{\text{Match}, \text{Mismatch}\}$. His payoff from Match is 1 if players 1 and 2 choose the same action (both a or both b) and 0 otherwise; his payoff from Mismatch is 1 if players 1 and 2 chose different actions (a and b or flipped), and 0 otherwise.

³⁹Suppose neither l nor r are strictly dominated for player 1; then, all actions are rationalizable for player 1-3, so Match is rationalizable for player 4. If either l or r is strictly dominated for player 1, then one of the following will be a rationalizable family: $\{l\} \times \{a\} \times \{a\} \times \{\text{Match}\}$, $\{l\} \times \{a, b\} \times \{a, b\} \times \{\text{Match}\}$, $\{r\} \times \{b\} \times \{b\} \times \{\text{Match}\}$, or $\{r\} \times \{a, b\} \times \{a, b\} \times \{\text{Match}\}$. Thus, Match is rationalizable for player 4.

Let $\varepsilon > 0$. If player 1 has common certainty in the state θ_ε^1 , then a is his uniquely rationalizable action: l strictly dominates r for player 3, given which a strictly dominates b for player 1. By a similar argument, if player 2 has common certainty in the state θ_ε^2 , then b is his uniquely rationalizable action. These statements hold for ε arbitrarily small. Construct a sequence of types $(t_4^{\varepsilon_n})$ for player 4, where each type $t_4^{\varepsilon_n}$ has common certainty that player 1 has common certainty in the state $\theta_{\varepsilon_n}^1$ and player 2 has common certainty in the state $\theta_{\varepsilon_n}^2$. Then, player 4 of type $t_4^{\varepsilon_n}$ has common certainty in an ε -neighborhood of θ , but only one rationalizable action: Mismatch. Take $\varepsilon_n \rightarrow 0$ (with each $\varepsilon_n > 0$) and the desired conclusion obtain: rationalizability of Match holds at $\lim_{n \rightarrow \infty} \varepsilon_n$ but fails to hold arbitrarily far out along the sequence ε_n .

O.4 Extension to Common p -Belief

For each $q \in [0, 1]$, define:

$$\underline{p}^{n,q}(a_i) = P^n \left(\left\{ \mathbf{z}_n : a_i \in S_i^\infty[t_i] \quad \forall t_i \in T_i^{\mathcal{B},q} \right\} \right).$$

where $q = 1$ returns the definition of $\underline{p}^n(a_i)$ given in the main text.

Proposition 5. *Suppose a_i is strictly rationalizable for type t_i^∞ , and define*

$$\delta^\infty := \sup \{ \delta : a_i \text{ is } \delta\text{-strictly rationalizable for type } t_i^\infty \}$$

noting that this quantity is strictly positive. Define

$$M := \sup_{a, a' \in A, \theta, \theta' \in \Theta, j \in \mathcal{I}} |u_j(a, \theta) - u_j(a', \theta')|. \quad (23)$$

Then, for every $n \geq 1$, and $q > M/(\delta^\infty + M)$,

$$\underline{p}^{n,q}(a_i) \geq 1 - \frac{2M\xi q}{\delta^\infty q - (1-q)M} \mathbb{E} \left(\sup_{\mu \in \mathcal{M}} d_P(\mu(Z^n), \mu^\infty) \right)$$

Proof. I first demonstrate a lemma analogous to Lemma 6.

Lemma 8. *Suppose a_i is δ^∞ -strictly rationalizable for player i of type t_i^∞ . Let $\mathcal{B} \subseteq \Delta(\Theta)$ by any set satisfying*

$$\sup_{\nu \in \mathcal{B}} d_P(\nu, \mu^\infty) \leq \frac{\delta^\infty q - (1-q)M}{2M\xi q}$$

where M is as defined in (23) and ξ is as defined in (6). Then, a_i is rationalizable for all types $t_i \in T_i^{\mathcal{B},q}$.

Proof. The proof follows along similar lines to the proof of Lemma 6. Fix $\epsilon > 0$, and consider an arbitrary set $\mathcal{B} \subseteq \{\mu^\infty\}^\epsilon$. I will show that a_i is rationalizable for all types $t_i \in T_i^{\mathcal{B},q}$ when ϵ is sufficiently small and q is sufficiently large.

By assumption, action a_i is δ^∞ -strictly rationalizable for player i of type t_i^∞ . This implies that there exists a family of sets $(R_j)_{j \in \mathcal{I}} \subseteq \prod_{j \in \mathcal{I}} A_j$, where $a_i \in R_i$, and for every $a_j \in R_j$ there exists a $\sigma_{-j}^\infty : \Theta \rightarrow \Delta(A_{-j})$ satisfying

$$\text{supp } \sigma_{-j}^\infty(\theta) \subseteq R_{-j} \quad \forall \theta \in \Theta$$

and

$$\int_{\Theta} u_j(a_j, \sigma_{-j}^\infty(\theta), \theta) d\mu^\infty - \int_{\Theta} u_j(a'_j, \sigma_{-j}^\infty(\theta), \theta) d\mu^\infty \geq \delta^\infty \quad \forall a'_j \neq a_j \quad (24)$$

Partition the set of types $T_j^{\mathcal{B},q}$ into those types whose first-order beliefs belong to \mathcal{B}

$$T_j^1 := \left\{ t_j \in T_j^{\mathcal{B},q} \mid t_j^1 \in \mathcal{B} \right\} \quad \forall j \in \mathcal{I}$$

and all remaining types $\bar{T}_j^1 := T_j^{\mathcal{B},q} \setminus T_j^1$. By construction, every type in $T_j^{\mathcal{B},q}$ assigns probability at least q to T_{-j}^1 . I will now show that there exists a family of sets $(V_j[t_j])_{j \in \mathcal{I}, t_j \in T_j^{\mathcal{B},q}}$ with the property that for each $k \geq 1$, player j , type $t_j \in T_j^{\mathcal{B},q}$, action $a_j \in R_j$, and mixed strategy $\alpha_j \in \Delta(A_j \setminus \{a_j\})$, there exists a measurable $\sigma_{-j} : \Theta \times T_{-j}^{\mathcal{B},q} \rightarrow \Delta(A_{-j})$ with

- (1) $\text{supp } \sigma_{-j}(\theta, t_{-j}) \subseteq V_{-j}[t_{-j}] \quad \forall (\theta, t_{-j}) \in \Theta \times T_{-j}^{\mathcal{B},q}$
- (2) $\int_{\Theta \times T_{-j}^{\mathcal{B},q}} [u_j(a_j, \sigma_{-j}(\theta, t_{-j}), \theta) - u_j(\alpha_j, \sigma_{-j}(\theta, t_{-j}), \theta)] t_j [d\theta \times dt_{-j}] \geq 0$
- (3) $V_j[t_j] = R_j$ for every player j and type $t_j \in T_j^1$.

Since $a_i \in R_i$ by design, it follows from Proposition 4 that for any type $t_i \in T_i^{\mathcal{B},q}$, the action $a_i \in S_i^k[t_i]$ for every k , and hence $a_i \in S_i^\infty[t_i]$, as desired.

Fix an arbitrary player j , $a_j \in R_j$, type $t_j \in T_j^{\mathcal{B},q}$, and $\alpha_j \in \Delta(A_j \setminus \{a_j\})$. Define $a_{-j} : \Theta \rightarrow A_{-j}$ to satisfy

$$a_{-j}(\theta) \in \underset{a_{-j} \in R_{-j}}{\text{argmax}} (u_j(a_j, a_{-j}, \theta) - u_j(\alpha_j, a_{-j}, \theta)) \quad \forall \theta \in \Theta$$

Further define

$$h(\theta) := u_j(a_j, a_{-j}(\theta), \theta) - u_j(\alpha_j, a_{-j}(\theta), \theta) \quad \forall \theta \in \Theta. \quad (25)$$

Let $\sigma_{-j} : \Theta \times T_{-j}^{\mathcal{B},q} \rightarrow \Delta(A_{-j})$ be any conjecture with the property that $\sigma_{-j}(\theta, t_{-j})$ is a point mass at $a_{-j}(\theta)$ for every $(\theta, t_{-j}) \in \Theta \times T_{-j}^{\mathcal{B},1}$. The conjectures $\sigma_{-j}(\theta, t_{-j})$ for $(\theta, t_{-j}) \notin \Theta \times T_{-j}^{\mathcal{B},1}$ are not explicitly specified. By definition,

$$\text{supp } \sigma_{-j}(\theta, t_{-j}) \subseteq R_{-j} \quad \forall (\theta, t_{-j}) \in \Theta \times T_{-j}^{\mathcal{B},1}.$$

Then

$$\begin{aligned}
& \int_{\Theta \times T_{-j}^1} u_j(a_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] - \int_{\Theta \times T_{-j}^1} u_j(\alpha_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] \\
&= \left(\int_{\Theta \times T_{-j}^1} u_j(a_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] + \int_{\Theta \times \bar{T}_{-j}^1} u_j(a_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] \right) - \\
&\quad \left(\int_{\Theta \times T_{-j}^1} u_j(\alpha_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] + \int_{\Theta \times \bar{T}_{-j}^1} u_j(\alpha_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] \right) \\
&= \left(\int_{\Theta \times T_{-j}^1} u_j(a_j, a_{-j}(\theta), \theta) t_j [d\theta \times dt_{-j}] - \int_{\Theta \times T_{-j}^1} u_j(\alpha_j, a_{-j}(\theta), \theta) t_j [d\theta \times dt_{-j}] \right) + \\
&\quad \left(\int_{\Theta \times \bar{T}_{-j}^1} u_j(a_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] - \int_{\Theta \times \bar{T}_{-j}^1} u_j(\alpha_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] \right) \\
&= \int_{\Theta \times T_{-j}^1} h(\theta) t_j [d\theta \times dt_{-j}] + \\
&\quad \left(\int_{\Theta \times \bar{T}_{-j}^1} u_j(a_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] - \int_{\Theta \times \bar{T}_{-j}^1} u_j(\alpha_j, \sigma_{-j}(\theta, t_{-j}), \theta) t_j [d\theta \times dt_{-j}] \right) \\
&\geq \int_{\Theta \times T_{-j}^1} h(\theta) t_j [d\theta \times dt_{-j}] - M \int_{\Theta \times \bar{T}_{-j}^1} t_j [d\theta \times dt_{-j}]
\end{aligned}$$

where the final inequality follows from the definitions of h (as given in (25)) and M (as given in (23)). In the proof of Lemma 6, we showed that the inequality in (24) implies $\int_{\Theta} h(\theta) t_j^1 [d\theta] \geq \delta^\infty - 2M\xi\epsilon$. Since moreover t_j assigns probability at least p to the set T_{-j}^1 , we can further bound

$$\int_{\Theta \times T_{-j}^{\emptyset, 1}} h(\theta) t_j [d\theta \times dt_{-j}] + M \int_{\Theta \times \bar{T}_{-j}^1} t_j [d\theta \times dt_{-j}] \geq q(\delta^\infty - 2M\xi\epsilon) - (1-q)M$$

Thus, a_j is a best reply for type t_j so long as

$$q(\delta^\infty - 2M\xi\epsilon) - (1-q)M \geq 0$$

which simplifies to $\epsilon \leq \frac{\delta^\infty q - (1-q)M}{2M\xi q}$. This bound holds across all players j and actions $a_j \in T_j^1$. \square

Thus

$$\begin{aligned}
\underline{p}^n(a_i) &\geq P^n \left(\left\{ \mathbf{z}_n : \sup_{\mu \in \mathcal{M}} d_P(\mu(\mathbf{z}_n), \mu^\infty) \leq \frac{\delta^\infty q - (1-q)M}{2M\xi q} \right\} \right) \\
&\geq 1 - \frac{2M\xi p}{\delta^\infty q - (1-q)M} \mathbb{E} \left(\sup_{\mu \in \mathcal{M}} d_P(\mu(Z^n), \mu^\infty) \right)
\end{aligned}$$

using Markov's inequality in the final line. \square