FINITE-ORDER IMPLICATIONS OF ANY EQUILIBRIUM

JONATHAN WEINSTEIN AND MUHAMET YILDIZ

ABSTRACT. Fix an arbitrary equilibrium refinement. Assume that, when we need to check whether a prediction of this refinement applies to a particular situation, we only know finitely many orders of players' interim beliefs. Many models may be consistent with this limited knowledge. In this paper, we characterize the predictions that are robust to these alternative specifications of the model, in the sense that they remain true independent of which of these models is chosen, given that we apply the fixed equilibrium refinement to each model. Assuming that the set of underlying payoff parameters is rich enough, for generic payoffs, we show that a prediction is robust if and only if it is true for all rationalizable strategies. Therefore, equilibrium refinements will be useful in generating stronger predictions than those of rationalizability only when we have information about the entire infinite hierarchy of beliefs, which is unlikely in practice. We also show that our result remains intact when we restrict our attention to finite type spaces, or impose the common-prior assumption.

Key words: robustness, equilibrium refinement, rationalizability, incomplete information, higher-order beliefs.

JEL Numbers: C72, C73.

Date: First Version: September, 2003; This Version: January, 2004.

We thank Daron Acemoglu, Pierpaolo Battigalli, Adam Brandenburger, Glenn Ellison, Jeffrey Ely, Drew Fudenberg, and the seminar participants at Berkeley, Caltech, Econometric Society Winter Meeting 2003, Harvard, MIT, Northwestern, Penn State, UPenn, Stanford, Washington University, Wisconsin-Madison, and Yale for helpful comments. We thank Eddie Dekel and two anonymous referees for detailed comments. We are especially grateful to Stephen Morris for extensive discussions that led to numerous improvements in the paper. This paper was written when Yildiz visited Cowles Foundation of Yale University, and he thanks them for their generous support and hospitality.

1. Introduction

In economic models, there are often many Nash equilibria. In order to be able to make sharp predictions, game theorists have therefore developed stronger solution concepts, which lead to a multitude of refinements of equilibrium, such as perfect and robust equilibrium. In applications, researchers typically use these refinements, often applying them to Bayesian games in which specific type spaces are chosen to model the players' incomplete information. In this paper, we characterize the predictions of such refinements that retain their validity when we actually have only a partial knowledge of the players' incomplete information.

To explain our framework, imagine a researcher who subscribes to some refinement of Nash equilibrium. For each possible incomplete-information model, represented by a Bayesian game, the researcher can compute a set of possible strategies using the refinement. He would like to be able to make predictions in a case of genuine incomplete information. By this we mean that when he is asked to make his prediction, the players already have some beliefs about what the payoffs are, which are called the first-order beliefs, beliefs about the other players' beliefs, which are called the second-order beliefs, and so on. In a standard type space, where a type of a player is a signal about payoff-relevant fundamentals, players form these beliefs at the interim stage upon learning their types, and we can compute an entire hierarchy of beliefs for a given type using the joint distribution of signals and the fundamentals (see Section 3).

Now we introduce the key assumption that the researcher is restricted to observing only finitely many orders of beliefs. This is especially plausible because common sense suggests that the players themselves will have their beliefs only partially articulated in their own minds. In particular, we assume that the researcher observes the first k orders of beliefs and nothing more. Many types that come from different type spaces are consistent with this observation, i.e., their first k orders of beliefs are as observed. He applies his refinement to obtain a set of strategies in each of these type spaces, specifying a set of actions for each of these types. The researcher cannot rule out any of these actions based on his observation and his refinement.

Now consider the researcher analyzing a particular type space, using his refinement. Under our restriction on what he can observe in the interim stage, he can only make predictions that are robust to alternative specifications of higher-order beliefs, in the sense that they will remain true for each of the models consistent with his observation on the first k orders of beliefs, given that he would have used the same equilibrium refinement in those alternative models as well. These are the predictions that can be verified by knowing only finitely many orders of beliefs.

Assuming that the space of underlying payoff parameters is rich enough, for generic payoffs we will characterize the robust predictions: *a prediction is robust if and only if it is true for all rationalizable strategies in the model.* That is, no equilibrium refinement produces any more predictive power than rationalizability, unless we also make a very precise assumption about the infinite hierarchy of beliefs. In other words, we cannot verify the predictions that are driven by the equilibrium refinement without observing the entire infinite hierarchy of beliefs.

In particular, we show that, if the above researcher observes that the first k orders of beliefs are consistent with a type t_i, then he cannot rule out any action that survives k rounds of iterated elimination of actions that are never a *strict* best reply for type t_i. On the other hand, by a result of Dekel, Fudenberg, and Morris (2003), he can rule out the actions that are eliminated in k rounds of iterated elimination actions that are never a *weak* best reply for type t_i. When there are no ties for best response (e.g., in generic games and nice games[1]), these elimination procedures lead to the same outcome, yielding the above characterization. This characterization is extended to all cases by allowing slight uncertainty in lower-order beliefs.

One may think that the above lack of robustness may be due to some large type spaces that do not satisfy certain standard assumptions, such as the common-prior assumption. Our conclusions, however, remain virtually unchanged when we restrict the possible type spaces to be finite and generated by a common prior. This is mainly because the common-prior assumption does not put any significant restriction on finite-order beliefs (Lipman (2003)).

A key element of our perspective on robustness is that we consider hierarchies of beliefs about payoff-relevant parameters to be the basic objects to be perturbed.[2]

[1]These are games where the action spaces are compact intervals and the utility functions are strictly concave in own action and continuous, as in many classical economic models.

[2]We define a perturbation as a mapping to a new type space such that the image of each type has the same first k orders of beliefs as the original type (see Section 5).

Alternatively, one could focus on perturbations of prior beliefs on type profiles (e.g. Kajii and Morris (1997)). That is, we focus on the interim stage (after types are realized, but before actions are taken), rather than the ex-ante stage. This approach is attractive because the ex-ante stage is often a hypothetical construct used to model the beliefs players hold in the interim stage (see Dekel and Gul (1987) and Battigalli (2003) for a detailed discussion). It is therefore appropriate to consider two situations with similar interim beliefs to be close, regardless of what prior beliefs led to the existing situation. Since distinct ex-ante models can lead to similar interim beliefs, our approach deals with the modeler's problem of selecting an ex-ante model that is appropriate for a given interim situation. We want to make predictions that are independent of selection among the models that are consistent with the data.

In the next section, using a variant of the e-mail game, we illustrate our results and compare them to the major existing results in the robustness literature. We present the basic definitions in Section 3. In Section 4, we prove a key result about the sensitivity of equilibria in universal type space, which will lead to the results in Section 5. In Section 5, we formulate our robustness notion and characterize the robust predictions of arbitrary equilibrium refinements. In Section 6, we present a modified version of our result without the richness assumption. In Section 7, we present three applications of our results, obtaining stark results on Cournot oligopoly and continuity and robustness of equilibrium strategies. Section 8 concludes. Some of the proofs are relegated to the appendix.

2. E-mail Game and Literature Review

To fix our ideas consider the coordinated attack game with payoff matrix

	Attack	No Attack
Attack	θ,θ	$\theta-5$,0
No Attack	0,$\theta-5$	0,0

where $\theta \in \Theta = \{-2, 2, 6\}$. First, consider the model in which it is common knowledge that $\theta = 2$. This case is modeled by a type space T with only one type for each player i, which will be denoted by $t_i^{CK}(2)$. In this game, there are two Nash equilibria in pure strategies, namely, (Attack, Attack) and (No Attack, No Attack), and an equilibrium in mixed strategies. Now imagine an incomplete information game in which the players may find it possible that $\theta = -2$. Ex ante, players assign probability 1/2 to each of the values -2 and 2. Player 1 observes the value of θ and automatically sends a message if $\theta = 2$. Each player automatically sends a message back whenever he receives one, and each message is lost with probability 1/2. When a message is lost the process automatically stops, and each player is to take one of the actions of Attack or No Attack. This game can be modeled by the type space $\tilde{T} = \{-1, 1, 3, 5, \ldots\} \times \{0, 2, 4, 6, \ldots\}$, where the type t_i is the total number of messages sent or received by player i (except for type $t_1 = -1$ who knows that $\theta = -2$), and the common prior p on $\Theta \times \tilde{T}$ where $p(\theta = -2, t_1 = -1, t_2 = 0) = 1/2$ and for each integer $m \geq 1$, $p(\theta = 2, t_1 = 2m-1, t_2 = 2m-2) = 1/2^{2m}$ and $p(\theta = 2, t_1 = 2m-1, t_2 = 2m) = 1/2^{2m+1}$. Here, for $k \geq 1$, type k knows that $\theta = 2$, knows that the other player knows $\theta = 2$, and so on through k orders of belief. The new incomplete-information game is dominance-solvable, and the unique equilibrium action for each player is No Attack.

Recall the researcher from the introduction, who could observe only k orders of beliefs for each player and has no knowledge of how players arrived at these beliefs. Suppose that he observes that the players mutually believe that $\theta = 2$ through k orders of beliefs. He cannot know whether the true model is T, or whether the true model is $\tilde{T}$ and the players have types greater than or equal to k. Now suppose that he subscribes to a non-empty equilibrium refinement that selects the (Attack,Attack) equilibrium in the complete information game T, as Pareto-dominance does. Since his refinement must assign the outcome (No Attack, No Attack) for each possible type in the alternative type space $\tilde{T}$, the researcher cannot know whether the outcome will be (Attack, Attack) or (No Attack, No Attack) according to his solution concept. To put it differently, even if the researcher believed that the right model is T and it is commonly accepted that his solution concept is the right solution concept, he could not prove his prediction that there will be an attack with the available data, which is insufficient to confirm that the correct model is T and not $\tilde{T}$. Consequently, we say that his prediction that there will be an attack is not robust at order k, for any k.

The above argument is based on Rubinstein's (1989) result that an equilibrium refinement that selects (Attack,Attack) must be sensitive to specifications of higher-order beliefs. A refinement that selects (No Attack,No Attack) is not sensitive in this example. Subsequently, in their seminal paper, Carlsson and van Damme (1993) presented a class of similar perturbations, where players observe noisy signals, according to which nearby types would always play the risk-dominant equilibrium, which is (No Attack, No Attack) in the above game. They then proposed that we should instead select the risk-dominant equilibrium of (No Attack, No Attack). As we will discuss later, Kajii and Morris (1997) also proposed a notion of robustness to incomplete information, according to which (No Attack, No Attack) is a robust equilibrium. Nevertheless, as we will now show, an equilibrium refinement that selects (No Attack, No Attack) must be sensitive to specifications of higher-order beliefs in precisely the same way as refinements that select (Attack, Attack). Indeed, if in the above example we replace $\theta = -2$ with $\theta = 6$, we obtain another, equally natural model, for which (Attack, Attack) is the unique equilibrium outcome for each type profile.[3] That is, we consider a new model with type space $\check{T} = \{-1, 1, 3, 5, \ldots\} \times \{0, 2, 4, 6, \ldots\}$ and the common prior q on $\Theta \times \tilde{T}$ where $q(\theta = 6, t_1 = -1, t_2 = 0) = 1/2$ and for each integer $m \geq 1$, $q(\theta = 2, t_1 = 2m - 1, t_2 = 2m - 2) = 1/2^{2m}$ and $q(\theta = 2, t_1 = 2m - 1, t_2 = 2m) = 1/2^{2m+1}$. One can easily check that this game is dominance-solvable, and all types play Attack. In our formulation, then, *both* of the possible predictions are non-robust, and the predictions of risk-dominance are no more robust than those of Pareto-dominance or other refinements. Indeed, our result will establish for arbitrary games and equilibrium refinements that a prediction will be robust only if it is true for all strategies that survive our elimination process—in the above example, all actions survive this process.

With the above example in mind, we now compare our result to important earlier results on robustness. Kajii and Morris (1997) introduced a notion of robustness of a given equilibrium of a given complete-information game to incomplete information. Their definition requires that for any incomplete-information game with a common prior that puts high probability on the original game, the original equilibrium action of the complete information game is played by most of the types in an equilibrium of the incomplete information game.[4] This concept of robustness

[3] At the end of Section 5, relaxing their assumptions on the noise structure, we also find a similar example in the framework of Carlsson and van Damme (1993).

[4] In Subsection 7.3, we give the formal definition of robust equilibrium and show how their results change when we switch to an interim notion of perturbation or drop the common-prior assumption.

rules out incomplete information games that involve large changes in prior but may lead to interim beliefs that are similar to the actual situation. They thus exclude the e-mail game, if the probability of $\theta \neq 2$ is not close to zero. Thus, they would exclude the constructions above, in which that probability is $1/2$. Kajii and Morris (1997) show, nevertheless, that this construction would work for the risk-dominant equilibrium (as in the original e-mail game), even if that probability is small.[5] Our latter example would not work if that probability is small. As we have shown above, however, a researcher could not distinguish these probabilities without having the knowledge of the entire infinite-hierarchy of beliefs. Then, the key difference between our notions of perturbation is that they focus on small changes to *prior* beliefs, without regard to the size of changes to *interim* beliefs, while our focus is the reverse. Their approach is appropriate when there is an ex-ante stage with well-understood inference rules and we know the prior to some degree. As we discussed in the introduction, however, in genuine incomplete-information situations, the type spaces and ex-ante stage are just tools for modeling interim beliefs. In that case, it is appropriate to consider types with similar interim beliefs, even if they come from models that assign small prior probability to the actual situation. Another distinction is that we ask if a predicted behavior is true in the perturbed game in every equilibrium that satisfies a given refinement, while they ask whether the predicted behavior is true in some equilibrium of the perturbed game. Our approach allows us to check directly if a theorem of the form "for all equilibria s that satisfy (a given) refinement, $Q(s)$ is true" remains valid if the modeling assumptions are slightly altered in the sense of this paper. The robust equilibrium notion of Kajii and Morris is silent as to whether such a theorem would remain valid if we modify the model using their perturbation.

Brandenburger and Dekel (1987) have shown that given a distribution on rationalizable strategy profiles of a given complete information game, we can enrich type space by adding payoff-irrelevant types and find an equilibrium in the new game that yields the same distribution on the strategy profiles of the original game.[6]

[5]Using the construction of Kajii and Morris (1997), for complete-information games, one could show that any robust prediction of a refinement must be true for $(p_1, \ldots, p_n)$-dominant equilibrium with $p_1 + \cdots + p_n < 1$ if such an equilibrium exists.

[6]This result has been extended to incomplete-information games by Dekel, Fudenberg, and Morris (2003) and also by Battigalli and Siniscalchi (2003), who also consider common-knowledge restriction on first-order beliefs. Our discussion applies to these extensions as well.

That is, in order for a prediction to be robust with respect to the entire set of equilibria *without any refinements*, it must be true for all rationalizable strategies. In contrast, we show that *for any refinement* of equilibrium, when we allow other payoff-relevant types from alternative models that lead to similar interim beliefs, in order for a prediction gained by the refinement to be robust, it must be true for all strategies that survive our elimination process. Notice that these two results point to somewhat contradictory properties for equilibrium correspondence. Brandenburger and Dekel suggest a large multiplicity of equilibria. Our result suggests that it is hard to refine rationalizability without precise information about types. This in turn suggests that there are relatively few rationalizability actions for a large set of types. Following our framework, Yildiz (2005) shows that generically there is indeed a unique rationalizable action.

Considering only payoff-irrelevant types (as Brandenburger and Dekel do), one cannot address the problem of robustness for arbitrary refinements, addressed in this paper. To see this, consider a non-dominance-solvable complete-information game with a pure Nash equilibrium, such as the complete-information case of the coordination game above. Following Brandenburger and Dekel, consider incomplete information games in which players have private information about some payoff-irrelevant parameter. By Tan and Werlang (1988), there is such an incomplete information game with an equilibrium such that each rationalizable strategy is played by some type. Since all types have the same belief hierarchy, this equilibrium is highly sensitive to higher-order beliefs. Because of this equilibrium, we cannot rule out any rationalizable action without invoking a refinement. Nevertheless, in those models, there always exists another equilibrium in which all types play according to a fixed pure strategy equilibrium of the original complete-information game. In contrast to extreme discontinuity of the former equilibrium with respect to players' hierarchy of beliefs about payoffs,[7] this equilibrium is constant. Then, following Carlsson and van Damme (1998) and others, such as Morris and Shin (1998), we can invoke a continuity argument to eliminate the discontinuous equilibrium. This would take us back to a subset of the equilibria of the original game. In contrast, if one considers the payoff-relevant types of our paper, it will be hard to invoke such a continuity argument. We show that, under a rich set of parameters, every rationalizable strategy will be discontinuous at every type with multiple

[7]It is not even a function of these beliefs.

rationalizable strategies; refinements based on continuity would lead to the empty set when we consider all possible types.

It is crucial for the result of Brandenburger and Dekel that they drop the common-prior assumption. By Aumann (1987), under the common-prior assumption, all types in their type spaces must play a correlated equilibrium strategy. Since economists commonly work under the common-prior assumption, they may ignore such a result. In contrast, each of the type spaces we constructed above has a common prior. Indeed, our results will remain virtually intact when we impose the common-prior assumption.

Our approach is closest to that of Fudenberg, Kreps, and Levine (1988) and Dekel and Fudenberg (1990). Fudenberg, Kreps, and Levine (1988) have shown that any equilibrium of any complete-information game can be made strict by perturbing the payoffs, showing that one cannot obtain any more predictions than those of all equilibria by considering refinements that do not eliminate any strict equilibria. Our result covers refinements that do eliminate some strict equilibria, such as the popular risk-dominance, and compares them to the larger set of all rationalizable strategies for arbitrary information structures.

In the same vein, Dekel and Fudenberg (1990) analyze the robustness of predictions based on iterated elimination of weakly dominated strategies when one allows payoff uncertainty as in this paper. They show that with uncertainty about players' beliefs at all orders, the robust predictions gained from this procedure for a complete-information game is equivalent to those of iterated strict dominance after one round of eliminating weakly dominated strategies.[8] They also show that even if we know that each player's prior put high probability to original payoffs, we could not rule out the possibility that a strategy that survive the latter elimination process is a strict equilibrium action. Hence, under this limited knowledge, the "robust" predictions of a refinement that does not eliminate any

[8]Borgers (1994) shows that the latter solution concept charcterizes the strategies that are consistent with almost common knowledge of players not playing weakly dominated strategies. Here, almost common knowledge is in the sense of common p-belief by Monderer and Samet (1989). Monderer and Samet show that an equilibrium remains as an approximate equilibrium (similar to robust equilibrium of Kajii and Morris) if there is common p-belief of the original game for high p, but we cannot check this condition without knowledge of infinite hierarchy of beliefs.

strict equilibrium are no more than the predictions of the latter solution concept. Nevertheless, as they note, their construction crucially relies on their departure from the common-prior assumption (as in Brandenburger and Dekel), and without common-prior assumption, such restrictions on priors do not put any restriction on the second-order beliefs and higher (see Section 7.3). Hence, in our formulation, they in effect assume that the researcher has information only about the first-order beliefs. In contrast, we consider arbitrary refinements, arbitrary information structures (as opposed to the complete information games) for the original game, impose common-prior assumption, allow the researcher to observe players' beliefs at arbitrarily high orders (as opposed to just observing the first-order beliefs), and yet we conclude that the robust predictions are just those of rationalizability.[9]

Following the critique of Wilson (1987), a sizeable literature has established that some central findings in economics, such as the full surplus extraction property of Cremer and Mclean (1988) in mechanism design (Neeman (2004) and Heifetz and Neeman (2003)) and the Coase conjecture in bargaining (Feinberg and Skrzypacz (2002)), crucially rely on the assumptions on the second-order beliefs and higher. In this paper, considering general games, under a richness assumption, we show that every equilibrium is highly sensitive to the way higher-order beliefs are specified. When all the common-knowledge assumptions are dropped, one cannot make any prediction that is stronger than rationalizability, no matter how sophisticated the refinements one uses. Of course, some may want to make explicit common-knowledge restrictions on players' payoffs and beliefs. In that case, it seems that a similar analysis to ours would show that the predictions of any refinement that remain valid with only partial knowledge of interim beliefs will be equivalent to that of Δ-rationalizability of Battigalli and Siniscalchi (2003), which corresponds to common-knowledge of these assumptions and rationality.[10]

[9]We show that, if we just know the first-order beliefs, then robust predictions are just those of rationality.

[10]Battigalli and Siniscalchi (2003) justify their solution concept following the framework of Brandenburger and Dekel; the above discussion of Brandenburger and Dekel applies to their analysis as well.

3. Basic Definitions

We consider a finite set of players $N = \{1, 2, \ldots, n\}$. There is a possibly unknown payoff-relevant parameter $\theta \in \Theta$ where Θ is a compact (and hence complete and separable) metric space. Each player i has action space A_i and utility function $u_i : \Theta \times A \to \mathbb{R}$, where $A = \prod_i A_i$.[11] We consider the set of games that differ in their specifications of the belief structure on θ, i.e. their type spaces, which we also call models. A type space is a set $T = T_1 \times \cdots \times T_n$ associated with beliefs $\kappa_{t_i} \in \Delta(\Theta \times T_{-i})$ for each $t_i \in T_i$.

Given any type t_i in a type space T, we can compute the belief of t_i on Θ by

$$t_i^1 = \operatorname{marg}_\Theta \kappa_{t_i},$$

which is called the *first-order belief of* t_i. We can compute the *second-order belief of* t_i, i.e. his belief about $(\theta, t_1^1, \ldots, t_n^1)$, by setting

$$t_i^2(F) = \kappa_{t_i}\left(\left\{(\theta, t_{-i}) \mid \left(\theta, t_i^1, t_{-i}^1\right) \in F\right\}\right)$$

for each measurable $F \subseteq \Theta \times \Delta(\Theta)^n$. We can compute an entire hierarchy of beliefs $\left(t_i^1, t_i^2, \ldots, t_i^k, \ldots\right)$ by proceeding in this way. We say that a type space T *does not have redundant types* if

$$t_i \neq \tilde{t}_i \Rightarrow \left(t_i^1, t_i^2, \ldots, t_i^k, \ldots\right) \neq \left(\tilde{t}_i^1, \tilde{t}_i^2, \ldots, \tilde{t}_i^k, \ldots\right) \qquad \left(\forall t_i, \tilde{t}_i \in T_i\right).$$

Mertens and Zamir (1985) showed that any type space without redundant types can be embedded in the universal type space, which we proceed to define following Brandenburger and Dekel (1993).[12]

Our definition incorporates an additional assumption that the players' beliefs at each finite order have countable (or finite) support. This assumption is made to

[11]**Notation:** Given any list $Y_1, \ldots, Y_n$ of sets, write $Y = \prod_i Y_i$, $Y_{-i} = \prod_{j \neq i} Y_j$, $y_{-i} = (y_1, \ldots, y_{i-1}, y_{i+1}, \ldots, y_n) \in Y_{-i}$, and $(y_i, y_{-i}) = (y_1, \ldots, y_{i-1}, y_i, y_{i+1}, \ldots, y_n)$. Likewise, for any family of functions $f_j : Y_j \to Z_j$, we define $f_{-i} : Y_{-i} \to Z_{-i}$ by $f_{-i}(y_{-i}) = (f_j(y_j))_{j \neq i}$. Given any metric space (Y, d), we write $\Delta(Y)$ for the space of probability distributions on Y, endowed with Borel σ-algebra and the weak topology. We use the product σ-algebra in product spaces. We write δ_x for the probability distribution that puts probability 1 on $\{x\}$. We also write $\operatorname{supp}(\pi)$ for the support of a probability distribution π, $\operatorname{marg}_Y \pi$ for the marginal of π on Y, and proj_Y for the projection mapping to Y.

[12]If there are redundant types, one needs to consider a larger type space (Ely and Peski (2004)) in order to analyze the robustness of predictions. The results of such an analysis will, if anything, show more sensitivity to the assumptions about higher-order beliefs.

avoid technical issues related to measurability (see Remark 1.) Our type space is dense in universal type space, and any countable type space with no redundant type is embedded in our space.

Define a sequence of sets X_k inductively by $X_0 = \Theta$ and $X_k = \left[\hat{\Delta}\left(X_{k-1}\right)\right]^n \times X_{k-1}$ for each $k > 0$, where $\hat{\Delta}\left(X_{k-1}\right) \subseteq \Delta\left(X_{k-1}\right)$ is the set of probability distributions on X_{k-1} that have countable support. Universal type space $T^u = \prod_{i \in N} T_i^u$ is the subset of $\left(\prod_{k=1}^{\infty} \hat{\Delta}\left(X_{k-1}\right)\right)^n$ in which it is common knowledge that the players' beliefs are coherent, i.e., the players know their own beliefs and their marginals from different orders agree. Notice that a type t_i in universal type space is simply the infinite hierarchy of beliefs $(t_i^1, t_i^2, \ldots) \in \prod_{k=1}^{\infty} \hat{\Delta}\left(X_{k-1}\right)$. We will use the variables $t_i, \tilde{t}_i \in T_i^u$ as generic types of any player i and $t, \tilde{t} \in T^u$ as generic type profiles.

From now on we will focus on type spaces with no redundant types; allowing redundant types would not change our results. As we mentioned, each such type space is isomorphic to a belief-closed subset of universal type space.[13] We will therefore represent all such type spaces as belief-closed subsets of universal type space, and write each type t_i as its belief hierarchy:

$$t_i = \left(t_i^1, t_i^2, \ldots, t_i^k, \ldots\right).$$

A type space (or a belief-closed subset) T is said to be *finite* iff T contains finitely many members and t_i^1 has finite support for each $t_i \in T_i$. Members of finite type spaces are referred to as *finite types*.

A *strategy* of a player i w.r.t. T_i is any measurable function $s_i : T_i \to A_i$. Given any type t_i and any profile s_{-i} of strategies, we write $\pi\left(\cdot|t_i, s_{-i}\right) \in \Delta\left(\Theta \times A_{-i}\right)$ for the joint distribution of the underlying uncertainty and the other players' actions induced by t_i and s_{-i}; $\pi\left(\cdot|t_i, \sigma_{-i}\right)$ is similarly defined for correlated mixed strategy profile σ_{-i}. For each $i \in N$ and for each belief $\pi \in \Delta\left(\Theta \times A_{-i}\right)$, we write $BR_i\left(\pi\right)$ for the set of actions $a_i \in A_i$ that maximize the expected value of $u_i\left(\theta, a_i, a_{-i}\right)$ under the probability distribution π. A strategy profile $s^* = (s_1^*, s_2^*, \ldots)$ is a *Bayesian Nash equilibrium* iff at each t_i,

$$s_i^*\left(t_i\right) \in BR_i\left(\pi\left(\cdot|t_i, s_{-i}^*\right)\right).$$

[13]A set $T \subset T^u$ is said to be *belief-closed* iff $\mathrm{supp}(\kappa_{t_i}) \subset \Theta \times T_{-i}$ for each $t_i \in T_i$.

An equilibrium s^* on universal type space T^u is said to have *full range* iff

$$s^*(T^u) = A. \tag{FR}$$

We prove a modified version of our results without this full range assumption in Section 6. Moreover, the assumption is without loss of generality, as we can restrict the action space to $s^*(T^u)$, by eliminating the actions that are never played in equilibrium s^*. Finally, full range is implied by the following assumption on the richness of the set Θ, commonly used in the global games literature. We will mention this assumption explicitly for those results for which it is used.

Assumption 1 (Richness Assumption)**.** *For each i and each a_i, there exists $\theta^{a_i} \in \Theta$ such that*

$$u_i(\theta^{a_i}, a_i, a_{-i}) > u_i(\theta^{a_i}, a_i', a_{-i}) \qquad (\forall a_i' \neq a_i, \forall a_{-i}).$$

This assumption corresponds to allowing the broadest possible set of beliefs about other players' payoffs, which should be allowed when we drop all common-knowledge assumptions.

Lemma 1. *Under Assumption 1, every equilibrium s^* on universal type space has full range.*

Elimination Processes. We will use interim rationalizability of Dekel, Fudenberg, and Morris (2003); see also Battigalli (2003), Battigalli and Siniscalchi (2003). Interim rationalizability allows correlations not only within players' strategies but also between their strategies and θ. Clearly, allowing such correlation only makes our sets larger. Since our main result is a lower bound in terms of these sets, this only strengthens our result. Moreover, our characterization provides yet another justification for this correlated rationalizability. For each i and t_i, set $S_i^0[t_i] = A_i$, and define sets $S_i^k[t_i]$ for $k > 0$ iteratively, by letting $a_i \in S_i^k[t_i]$ if and only if $a_i \in BR_i\left(\mathrm{marg}_{\Theta\times A_{-i}}\pi\right)$ for some $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that $\mathrm{marg}_{\Theta\times T_{-i}}\pi = \kappa_{t_i}$ and $\pi\left(a_{-i} \in S_{-i}^{k-1}[t_{-i}]\right) = 1$, i.e., a_i is a best reply to a belief that assigns positive probability only to the actions that survive $k-1$ orders of elimination. The set of all rationalizable actions for player i (with type t_i) is

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

By a *rationalizable strategy* (w.r.t. a model T), we mean a strategy $s_i : T_i \to A_i$ with $s_i(t_i) \in S_i^\infty[t_i]$ for each $t_i \in T_i$.

Next we define the set of strategies that survive iterative elimination of strategies that are never strict best reply, denoted by $W^\infty[t_i]$, similarly. We set $W_i^0[t_i] = A_i$ and let $a_i \in W_i^k[t_i]$ if and only if $BR_i\left(\text{marg}_{\Theta \times A_{-i}}\pi\right) = \{a_i\}$ for some $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that $\text{marg}_{\Theta \times T_{-i}}\pi = \kappa_{t_i}$ and $\pi\left(a_{-i} \in W_{-i}^{k-1}[t_{-i}]\right) = 1$. Finally, we set

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

Notice that we eliminate a strategy if it is not a strict best-response to any belief on the remaining strategies of the other players. Clearly, this yields a smaller set than the result of iterative elimination of weakly dominated strategies.[14] In some games, the latter may yield strong predictions. For example, in finite perfect information games it leads to backwards induction outcomes. Nevertheless, in generic normal-form games of complete information, all these concepts are equivalent and usually have weak predictive power.

4. Sensitivity to higher-order beliefs

In Section 5, we will analyze the robustness of predictions according to arbitrary equilibrium refinements to the assumptions about higher-order beliefs in arbitrary type spaces. Robustness of such predictions is closely related to the sensitivity of an equilibrium on universal type space with respect to the changes in higher-order beliefs. In this section, we will focus on the latter problem.

We fix an equilibrium s^* on universal type space and a type t_i of a player i. According to equilibrium, he will play $s_i^*(t_i)$. Now imagine a researcher who only knows the first k orders of beliefs of player i and knows that equilibrium s^* is played. All the researcher can conclude from this information is that i will play one of the actions in

$$A_i^k[s^*, t_i] \equiv \left\{ s_i^*\left(\tilde{t}_i\right) | \tilde{t}_i \in T_i^u, \quad \tilde{t}_i^m = t_i^m \quad \forall m \leq k \right\}.$$

[14] In particular, if we use non-reduced normal-form of an extensive-form game, many strategies will be outcome equivalent, in which case our procedure will eliminate all of these strategies. To avoid such over-elimination, we can use reduced-form, by representing all outcome-equivalent strategies by only one strategy.

Assuming, plausibly, that a researcher can verify only finitely many orders of a player's beliefs, all a researcher can ever know is that player i will play one of the actions in

$$A_i^{\infty}[s^*, t_i] = \bigcap_{k=0}^{\infty} A_i^k[s^*, t_i].$$

Our next result finds tight bounds for the sets $A_i^k[s^*, t_i]$ for games with countable action spaces.

Proposition 1. *For any countable (or finite) action space A, any equilibrium s^* with full range, any $k \in \mathbb{N}$, $i \in N$, and any t_i,*

$$W_i^k[t_i] \subseteq A_i^k[s^*, t_i] \subseteq S_i^k[t_i];$$

in particular,

$$W_i^{\infty}[t_i] \subseteq A_i^{\infty}[s^*, t_i] \subseteq S_i^{\infty}[t_i].$$

The conclusion that $W_i^k[t_i] \subseteq A_i^k[s^*, t_i]$ can be spelled out as follows. Suppose that we know a player's beliefs up to the kth order and do not have any further information. Suppose also that he has an action a_i that survives k rounds of iterated elimination of strategies that cannot be a strict best reply—for some type whose first k orders of beliefs match what we know. Then, we cannot rule out that a_i will be played in equilibrium s^*. Now, consider the equilibrium refinement that selects the restriction of s^* to T at each type space T. If an action a_i survives k rounds of our elimination procedure for a type t_i within a type space T, then we can find another type, possibly from another type space, whose first k orders of beliefs are as specified by t_i but plays a_i according to the unique solution of the refinement. Hence, we cannot rule out action a_i as a solution of our equilibrium refinement when we only know that the first k orders of beliefs are as specified by t_i. A prediction of our refinement then can only be robust if it remains true when we assign a_i as the solution at t_i. In the next section we will formalize this for arbitrary equilibrium refinements. We now proceed to the proof of the result, which highlights the logic behind this sensitivity to higher-order beliefs.

Proof. We first show that $W_i^k[t_i] \subseteq A_i^k[s^*, t_i]$. For $k = 0$, the statement is equivalent to the full-range assumption. For any given k and any player i, write each t_{-i} as $t_{-i} = (l, h)$ where $l = \left(t_{-i}^1, t_{-i}^2, \ldots, t_{-i}^{k-1}\right)$ and $h = \left(t_{-i}^k, t_{-i}^{k+1}, \ldots\right)$ are the lower and higher-order beliefs, respectively. Let $L = \left\{l | \exists h : (l,h) \in T^u_{-i}\right\}$. The induction hypothesis is that

$$W^{k-1}_{-i}[l] \equiv \bigcup_{h'} W^{k-1}_{-i}[(l,h')] \subseteq A^{k-1}_{-i}[s^*,(l,h)] \quad (\forall (l,h) \in T^u_{-i}).$$

Fix any type t_i and any $a_i \in W^k_i[t_i]$. We will construct a type $\tilde{t}_i$ such that $s^*_i(\tilde{t}_i) = a_i$ and the first k orders of beliefs are same under t_i and $\tilde{t}_i$, showing that $a_i \in A^k_i[s^*,t_i]$. Now, by definition of $W^k_i[t_i]$, $BR_i(\mathrm{marg}_{\Theta\times A_{-i}}\pi) = \{a_i\}$ for some $\pi \in \Delta\left(\Theta \times T^u_{-i} \times A_{-i}\right)$ such that $\mathrm{marg}_{\Theta\times T_{-i}}\pi = \kappa_{t_i}$ and $\pi\left(a_{-i} \in W^{k-1}_{-i}[t_{-i}]\right) = 1$. By the induction hypothesis, for each $(\theta,l,a_{-i}) \in \mathrm{supp}(\mathrm{marg}_{\Theta\times L\times A_{-i}}\pi)$, $a_{-i} \in W^{k-1}_{-i}[l] \subseteq A^{k-1}_{-i}[s^*,(l,h)]$ for some h. Hence, there exists a mapping $\mu : \mathrm{supp}(\mathrm{marg}_{\Theta\times L\times A_{-i}}\pi) \to \Theta \times T^u_{-i}$,

$$\mu : (\theta,l,a_{-i}) \mapsto \left(\theta,l,\tilde{h}(a_{-i},\theta,l)\right), \tag{4.1}$$

such that

$$s^*_{-i}\left(l,\tilde{h}(a_{-i},\theta,l)\right) = a_{-i}. \tag{4.2}$$

We define $\tilde{t}_i$ by

$$\kappa_{\tilde{t}_i} \equiv \left(\mathrm{marg}_{\Theta\times L\times A_{-i}}\pi\right) \circ \mu^{-1},$$

the probability distribution induced on $\Theta \times T^u_{-i}$ by the mapping μ and the probability distribution π. Notice that, since t^k_i has countable support and the action spaces are countable, the set $\mathrm{supp}(\mathrm{marg}_{\Theta\times L\times A_{-i}}\pi)$ is countable, in which case μ is trivially measurable. Hence $\kappa_{\tilde{t}_i}$ is well-defined. By a well-known isomorphism by Mertens and Zamir (1985), $\kappa_{\tilde{t}_i}$ is the belief of a (unique) type $\tilde{t}_i$, such that

$$\tilde{t}^m_i = \delta_{\tilde{t}^{m-1}_i} \times \mathrm{marg}_{\Theta\times[\Delta(X_{m-2})]^{N\setminus\{i\}}}\kappa_{\tilde{t}_i}, \qquad (\forall m > 1) \tag{4.3}$$

and $\tilde{t}^1_i = \mathrm{marg}_\Theta \kappa_{\tilde{t}_i}$. Since $\mathrm{supp}(\kappa_{\tilde{t}_i})$ is countable, each $\tilde{t}^m_i$ has countable support. By construction of μ, the first k orders of beliefs (about (θ,l)) are identical under t_i and $\tilde{t}_i$:

$$\begin{aligned}
\mathrm{marg}_{\Theta\times L}\kappa_{\tilde{t}_i} &= \mathrm{marg}_{\Theta\times L}\left[\left(\mathrm{marg}_{\Theta\times L\times A_{-i}}\pi\right)\circ\mu^{-1}\right] = \mathrm{marg}_{\Theta\times L}\left(\mathrm{marg}_{\Theta\times L\times A_{-i}}\pi\right)\\
&= \mathrm{marg}_{\Theta\times L}\pi = \mathrm{marg}_{\Theta\times L}\left(\mathrm{marg}_{\Theta\times T^u_{-i}}\pi\right) = \mathrm{marg}_{\Theta\times L}\kappa_{t_i},
\end{aligned}$$

where the second equality is by (4.1). Together with (4.3) and identical equality for t_i, this shows that $\tilde{t}^m_i = t^m_i$ for each $m \le k$. Towards showing that $s^*_i(\tilde{t}_i) = a_i$, let $\tilde{\pi} = \kappa_{\tilde{t}_i} \circ \gamma^{-1} \in \Delta\left(\Theta \times T^u_{-i} \times A_{-i}\right)$ be the equilibrium belief of type $\tilde{t}_i$, where $\gamma : (\theta, l, h) \mapsto \left(\theta, l, s^*_{-i}(l, h)\right)$. By construction,

$$\begin{aligned}
\operatorname{marg}_{\Theta \times L \times A_{-i}} \tilde{\pi} &= \kappa_{\tilde{t}_i} \circ \gamma^{-1} \circ \operatorname{proj}^{-1}_{\Theta \times L \times A_{-i}} \\
&= \left(\operatorname{marg}_{\Theta \times L \times A_{-i}} \pi\right) \circ \mu^{-1} \circ \gamma^{-1} \circ \operatorname{proj}^{-1}_{\Theta \times L \times A_{-i}} = \operatorname{marg}_{\Theta \times L \times A_{-i}} \pi.
\end{aligned}$$

[By (4.2) and the definition of γ, $\operatorname{proj}_{\Theta \times L \times A_{-i}} \circ \gamma \circ \mu$ is the identity mapping, yielding the last equality.] Therefore,

$$\pi\left(\cdot|\tilde{t}_i, s^*_{-i}\right) = \operatorname{marg}_{\Theta \times A_{-i}} \tilde{\pi} = \operatorname{marg}_{\Theta \times A_{-i}} \pi.$$

Since a_i is the only best reply to these beliefs, $\tilde{t}_i$ must play a_i in equilibrium s^*:

$$s^*_i\left(\tilde{t}_i\right) \in BR_i\left(\pi\left(\cdot|\tilde{t}_i, s^*_{-i}\right)\right) = BR_i\left(\operatorname{marg}_{\Theta \times A_{-i}} \pi\right) = \{a_i\}. \tag{4.4}$$

To see the inclusion $A^k_i[s^*, t_i] \subseteq S^k_i[t_i]$, observe that for any $\tilde{t}_i$ with $\tilde{t}^m_i = t^m_i$ for each $m \leq k$, we have

$$s^*_i\left(\tilde{t}_i\right) \in S^\infty_i\left[\tilde{t}_i\right] \subseteq S^k_i\left[\tilde{t}_i\right] = S^k_i[t_i],$$

where the last equality is due to a result by Dekel, Fudenberg, and Morris (2003) that $S^k_i[t_i]$ depends only on the first k orders of beliefs $\left(t^1_i, \ldots, t^k_i\right)$, completing the proof. [Dekel, Fudenberg, and Morris (2003) make some further finiteness assumptions. For a constructive but much longer proof of the last part under our relaxed assumptions, see our earlier working paper.] □

Remark 1. Notice that the countability assumptions about the finite-order beliefs and the action spaces are used only to make sure that $\kappa_{\tilde{t}_i}$ is a well-defined probability distribution, or μ is measurable. In fact, whenever μ is measurable, our proof is valid. Below, we present another class of games in which μ is measurable; μ may not be measurable in general. These assumptions are not needed at all for the inclusion $A^k_i[s^*, t_i] \subseteq S^k_i[t_i]$.

The next example shows that either of the inclusions in Proposition 1 may be strict in general.

Example 1. Take $N = \{1,2\}$, $\Theta = \{\theta_0, \theta_1\}$, and let the action spaces and the payoff functions for each θ be given by

	a^0	a^1
a^0	0,0	0,0
a^1	0,0	1,1

(Note that θ is not payoff relevant.) Define s^* by

$$s_i^*(t_i) = \begin{cases} a^0 & \text{if } t_i = t_i^{CK}(\theta_0); \\ a^1 & \text{otherwise.} \end{cases}$$

Clearly, for each $k \geq 1$, we have $W_i^k[t_i] = \{a^1\}$ and $S_i^k[t_i] = \{a^0, a^1\}$ for each t_i, while $A_i^k\left[s^*; t_i^{CK}(\theta_0)\right] = \{a^0, a^1\}$, and $A_i^k\left[s^*; t_i^{CK}(\theta_1)\right] = \{a^1\}$.

On the other hand, in many cases the two elimination processes are equivalent, and Proposition 1 yields a precise characterization of the set A^k. We now present two such important cases. First, let $t_i^{CK}\left(\bar{\theta}\right)$ denote the type who believes that it is common knowledge that $\theta = \bar{\theta}$, for some $\bar{\theta}$ at which the payoffs are generic.[15] In this case, any action that is not strictly dominated will be a strict best reply against some belief (at each round), and hence the two elimination processes will be equivalent. Therefore, Proposition 1 yields the following characterization.

Corollary 1. *For any finite-action game and any equilibrium s^* with full range, if the payoffs are generic at some θ, then for each i and k,*

$$A_i^k\left[s^*, t_i^{CK}(\theta)\right] = S_i^k\left[t_i^{CK}(\theta)\right].$$

In particular, letting $k = \infty$, we find that the set of actions we cannot exclude with only finite-order knowledge of players' beliefs is precisely the set of rationalizable actions.

The second case is the class of "nice" games (Moulin (1984)), which are widely used in economic theory, such as imperfect competition, spatial competition, provision of public goods, theory of firm, etc.

[15] We say that *the payoffs are generic at θ* iff there do not exist i, non-zero $\alpha \in \mathbb{R}^{A_i}$, and distinct a_i, a_i', a_{-i}, and a_{-i}' such that (i) $u_i(\theta, a_i, a_{-i}) = u_i(\theta, a_i', a_{-i})$ or (ii) $\sum_{a_i} \alpha(a_i) u_i(\theta, a_i, a_{-i}) = \sum_{a_i} \alpha(a_i) u_i\left(\theta, a_i, a_{-i}'\right) = 0$.

Definition 1. A game is said to be *nice* iff for each i, $A_i = [0,1]$ and $u_i(\theta, a_i, a_{-i})$ is continuous in $a = (a_i, a_{-i})$ and strictly concave[16] in a_i.

In a nice game, since players always have unique best reply, our elimination processes will be equivalent, yielding the functional equation

$$W = S. \tag{4.5}$$

Moreover, a lemma by Moulin (1984) and Battigalli (2003) ensures that we can focus on degenerate beliefs, allowing us to circumvent the measurability issue discussed in Remark 1. We then obtain the same characterization of A^k for nice games that we found above under the conditions of Corollary 1. (The formal statement of the lemma and the proof of the proposition are in the appendix.)

Proposition 2. *For any nice game, any equilibrium s^* with full range, any countable, belief-closed $T \subset T^u$, for any $k \leq \infty$, $i \in N$, and $t_i \in T_i$,*

$$S_i^k[t_i] = A_i^k[s^*, t_i].$$

Under a stability condition similar to that of Nyarko (1996), Weinstein and Yildiz (2003) have shown that the maximum impact of higher-order beliefs to diminish exponentially, i.e. the set $A_i^k[s^*, t_i]$ shrinks to a point as $k \to \infty$. Then, Proposition 2 shows that, in nice games, these stability conditions must imply that the game is dominance solvable.

In analysis of general equilibrium refinements, we will assume that the refinement is non-empty at least in finite games. Since the existence results are often for mixed-strategies, we now extend Proposition 1 to the mixed-strategy equilibria. Using interim formulation, we define a *mixed strategy* as any measurable function $\sigma_i : T_i \to \Delta(A_i)$. A mixed strategy profile σ^* is *Bayesian Nash equilibrium* iff $\text{supp}(\sigma_i^*(t_i)) \subseteq BR_i\left(\pi\left(\cdot|t_i, \sigma_{-i}^*\right)\right)$ for each i and t_i. We say that σ^* has full range iff for each a_i, there exists t_i with $\text{supp}(\sigma_i^*(t_i)) = \{a_i\}$. Clearly, under Assumption 1, every equilibrium on universal type space has full range. We also set

$$A_i^k[\sigma^*; t_i] = \left\{a_i | \text{supp}\left(\sigma_i^*\left(\tilde{t}_i\right)\right) = \{a_i\}, \quad \tilde{t}_i^m = t_i^m \quad \forall m \leq k\right\},$$

[16]We use the strict concavity assumption to make sure that a player's utility function for any fixed strategy profile of the others is always single-peaked in his own action. (Single-peakedness is not preserved in presence of uncertainty.)

the set of all actions that are played with probability 1 under σ^* by some type $\tilde{t}_i$ whose first k orders of beliefs are identical to those of t_i.

Proposition 3. *For any countable-action game, any (possibly mixed strategy) equilibrium σ^* with full range, any $k \leq \infty$, $i \in N$, and any t_i,*

$$W_i^k[t_i] \subseteq A_i^k[\sigma^*, t_i] \subseteq S_i^k[t_i].$$

That is, if σ^* has full range and we know only the first k orders of a player's beliefs, then for any $a_i \in W_i^k[t_i]$, we cannot rule out that a_i is played with probability 1 according to σ^*. The proof of this result simply focus on the types with pure actions and applies the proof of Proposition 1.

5. Robustness to higher-order beliefs

In this section we will formalize our notion of robustness to higher-order beliefs and bound the set of robust predictions for arbitrary refinements of equilibrium. Under Assumption 1, we will obtain a sharp characterization. We start by formally presenting some basic definitions.

Definition 2. A *solution concept* is any mapping Σ that maps each type space T to a set $\Sigma(T)$ of (mixed) strategy profiles σ with respect to T. An *equilibrium refinement* is any solution concept Σ such that for each T and $\sigma \in \Sigma(T)$, σ is a Bayesian Nash equilibrium of T.

We will frequently refer to solution concepts S^k, which yields all the strategy profiles with $s(t) \in S^k[t]$ for each $t \in T$ at each T, and W^k (similarly defined).

Definition 3. Given a solution concept Σ and a type space T, by a *prediction* of (Σ, T), we mean any formula Q with free variable $s : T \to A$ such that $Q(s)$ is true for each $s \in \text{supp}(\sigma)$ and each $\sigma \in \Sigma(T)$.

Although we allow solution concepts to be mixed, we are here focusing on the deterministic predictions, the predictions that remain true for each realization. Here a prediction can be about the behavior of a particular type. In an auction, for example, a prediction could be "the type with lowest valuation bids zero." A prediction can also be about a relation between the behavior of different types, e.g., "a player's bid is an increasing function of his valuation."

We first envision a researcher who subscribes to an equilibrium refinement Σ and can observe players' beliefs *precisely* up to an arbitrary but finite order k, where k is large. This is clearly an unrealistically generous assumption, but we will show that, despite this, the researcher cannot make very strong predictions.

Definition 4. Given a model T, a pair $(\tilde{T}, \tau)$ of a model $\tilde{T}$ and a mapping $\tau : T \to \tilde{T}$ is said to be a *k-perturbation* of T iff for each $t \in T$ and $\tilde{t} = \tau(t)$, we have $\tilde{t}_i^l = t_i^l$ for each $l \leq k$.

The definition of k-perturbation requires that whenever our researcher believes that a type profile t in T may describe the actual situation, he cannot rule out that the type profile $\tau(t)$ in $\tilde{T}$ describes the situation. That is, the researcher cannot reject the perturbation without rejecting the original model. A perturbation may result from relaxing the assumption that a certain fact is common knowledge, instead assuming that it is mutually known only up to kth order, and perhaps making some other assumption about the higher-order beliefs. Reflecting such a relaxation, the perturbed model will then have more type profiles. This is the case in the coordinated attack game of Section 2. As we discussed, type k (in $\tilde{T}$ or in $\check{T}$) agrees with the common knowledge type $t_i^{CK}(2)$ up to order k. Therefore, both $(\tilde{T}, \tau)$ and $\left(\check{T}, \tau\right)$ are k-perturbations of the complete information game T for any mapping τ whose values are both at least k.

Our robustness condition will require that the prediction remains valid for all perturbations defined as above. The motivation is clear. Imagine a researcher analyzing a model T, knowing that when he is asked to validate that his model applies to a particular situation, he will be able to observe only first k orders of beliefs. He knows that, if he can validate that T is consistent with the actual situation, $\tilde{T}$ will also be consistent with the actual situation. If a prediction of his model T does not remain true for a perturbation $\tilde{T}$, then he cannot justify that his prediction applies to a particular situation, even when he could validate his model. Therefore, the researcher would like to focus on the predictions of T that are robust to alternative specifications, such as $\tilde{T}$, given that he was going to apply his solution concept in those alternative specifications as well. We define robustness as follows.

Definition 5. A prediction Q of (Σ, T) is said to be *k-robust (to higher-order beliefs)* iff for each k-perturbation $(\tilde{T}, \tau)$, for each $\sigma \in \Sigma(\tilde{T})$, for each $s \in \text{supp}(\sigma)$, $Q(s \circ \tau)$ is true. Prediction Q is said to be *robust to higher-order beliefs* iff it is k-robust for some $k < \infty$.

That is, a prediction is said to be k-robust if it remains true in models where beliefs change at orders higher than k according to the perturbation mapping and we apply the same solution concept throughout. We could weaken our robustness requirement by requiring $Q(s \circ \tau)$ to be true only if s is the "unique" solution in the perturbed model, i.e., $\text{supp}(\sigma(\tau(t))) = \{s(t)\}$ at each $t \in T$. It will be clear that our results would remain valid under this substantially weaker requirement. Returning again to the coordinated attack example, the prediction of no attack for the complete information game $T = \left\{t^{CK}(2)\right\}$ is not robust under any equilibrium refinement Σ because for each k, $\left(\check{T}, \tau\right)$ with $\tau\left(t^{CK}(2)\right) > (k, k)$ is a k-perturbation of T, and for the unique member σ of $\Sigma\left(\check{T}\right)$, $\sigma\left(\tau\left(t^{CK}(2)\right)\right)$ assigns probability 1 to (Attack,Attack). Similarly, the prediction of attack is not robust.

Both of the non-robustness results in this example are special cases of the upcoming proposition. It states no refinement can make robust predictions that are any more powerful than the predictions that are generated by W^∞, the iterated elimination of actions that are never a strict best reply.

Proposition 4. *Let A be countable (or finite), T be any model, and Σ be an equilibrium refinement with $\sigma \in \Sigma(T^u)$ that has full range. Every k-robust prediction Q of (Σ, T) is also a prediction of $\left(W^k, T\right)$. In particular, if Q is robust to higher-order beliefs, then Q is a prediction of (W^∞, T). Conversely, if Q is a prediction of $\left(S^k, T\right)$, then Q is a k-robust prediction of (Σ, T).*

Proof. Take any k-robust prediction Q of (Σ, T) and any $s : T \to A$ with $s(t) \in W^k[t]$ for each $t \in T$. By Proposition 3, for each t, there exists $\tilde{t} = \tau(t)$ such that $\text{supp}(\sigma(\tau(t))) = \{s(t)\}$ and $\tilde{t}_i^l = t_i^l$ for each $l \leq k$ and each i. Taking (T^u, τ) as a k-perturbation of T, we then conclude that $Q(s)$ is true. Therefore, Q is a prediction of $\left(W^k, T\right)$. For the converse, take any prediction Q of $\left(S^k, T\right)$. Take also any k-perturbation $(\tilde{T},\tau)$ of T and any $s \in \text{supp}(\sigma)$ and $\sigma \in \Sigma(\tilde{T})$. For each $\tau(t)$, $s(\tau(t)) \in S^k[\tau(t)] = S^k[t]$, showing that $s \circ \tau \in S^k(T)$. Since Q is a prediction of (S^k, T), this shows that $Q(s \circ \tau)$ is true. Therefore, Q is a k-robust prediction of (Σ, T). $\square$

For nice games, using Proposition 2, we obtain a characterization:

Proposition 5. *For any nice game, any equilibrium refinement Σ with $\sigma \in \Sigma(T^u)$ that has full range, any model T, and any $k < \infty$, a prediction Q of (Σ, T) is k-robust if and only if Q is a prediction of (S^k, T).*

While the above propositions yield strong bounds for robust predictions, they leave number of issues unresolved. Firstly, they assume that the refinement is non-empty on universal type space. But most existing existence results for refinements assume finite models.[17] Secondly, our robustness condition above requires that the researcher does not restrict the set of models a priori. In particular, for simplicity, our proof uses the universal type space as one of possible models. Although we find this a reasonable condition, some may want to restrict the set of models by fiat. For example, it is customary in economics literature to assume that there is a common prior, and one may want to impose the common prior assumption on possible models, ignoring the models without a common prior all together. One may also want to focus on small models, such as finite type spaces, for perhaps to assure existence. Finally, our lower bound for countable-action games is in terms of W^∞, which may be small in certain games, weakening our results. We will next show that these potential issues are not crucial for our results. Indeed, requiring that the predictions remain valid if there is also very small misspecification in lower-order beliefs, we will obtain a characterization of robust predictions in terms of rationalizability, even when the set of models are restricted to be finite models with common prior.

Formally, a finite model T is said to have a common prior (with full support) if there exists a probability distribution $p \in \Delta(\Theta \times T)$ such that $\text{supp}(p) = \Theta' \times T$ for some $\Theta' \subseteq \Theta$ and $\kappa_{t_i} = p(\cdot|\Theta \times \{t_i\} \times T_{-i})$ for each $t_i \in T_i$. In the remainder

[17] Simon (2003) shows existence of equilibrium with non-measurable strategies for the union of countable type spaces, but he also shows that there may not exist an equilibrium with measurable strategies in universal type space for some payoff functions. We do assume that the strategies are measurable for ease of exposition, but that assumption does not play any role in our results.

of the section we will focus on the finite models with a common prior. We now imagine that a researcher can observe finitely many orders of beliefs with some small noise. We focus on the case that the noise is small in the sense that, for each $l \leq k$, the finite-order beliefs $\tilde{t}_i^l$ that the researcher finds possible converge to the observed beliefs t_i^l in the sense of "convergence in distribution", i.e., in weak topology. (Recall that t_i^l is a probability distribution.) To do this, we consider an arbitrary metric d on finite-order beliefs that metrizes the weak topology. The arbitrary metric d is taken to measure the distance according to the test the researcher employs. Given the observed or estimated beliefs, t_i^l, $l \leq k$, the researcher finds the set of beliefs $\tilde{t}_i^l$ with $d\left(\tilde{t}_i^l, t_i^l\right) \leq \epsilon$ for all $l \leq k$ possible (or cannot reject them at a particular level of confidence) for some $\epsilon > 0$, where ϵ is meant to measure the precision of the researcher's observations. Again, we focus on the limit $\epsilon \to 0$ and $k \to \infty$.

Definition 6. A model $(\tilde{T}, \tau)$ is said to be (ϵ, k)-*perturbation* of T iff (i) $\tilde{T}$ is finite and has a common prior, and (ii) $\tau : T \to \tilde{T}$ is such that for each t and $\tilde{t} = \tau(t)$, we have $d\left(\tilde{t}_i^l, t_i^l\right) \leq \epsilon$ for each $l \leq k$.

Definition 7. A prediction Q of (Σ, T) is said to be (ϵ, k)-*robust* iff for each (ϵ, k)-perturbation $(\tilde{T}, \tau)$ of T, for each $\sigma \in \Sigma(\tilde{T})$, for each $s \in \mathrm{supp}(\sigma)$, $Q(s \circ \tau)$ is true. Prediction Q is said to be *robust* iff it is (ϵ, k)-robust for some $\epsilon > 0$ and $k < \infty$.

On the one hand, we strengthen our robustness requirement by requiring ϵ to be positive but arbitrarily small, instead of setting it to zero as in the previous definition. On the other hand, we weaken our condition substantially by ignoring the perturbations that do not lead to a finite model with a common prior. In particular, the universal type space is no longer accepted as a possible perturbation. The next result, proven in Yildiz (2005, Proposition 3), will help us to characterize the robust predictions.

Lemma 2. *Assume that A is finite. Under Assumption 1, for any finite model T, any rationalizable strategy profile $s : T \to A$, for each $\epsilon > 0$ and $k < \infty$, there exist a (ϵ, k)-perturbation $(\tilde{T}, \tau)$ of T, where $\tilde{T}$ is finite and has a common prior, such that $S^\infty[\tau(t)] = \{s(t)\}$ for each $t \in T$.*

This result immediately yields a sharp characterization of the robust predictions.

Proposition 6. *Assume that A is finite. Under Assumption 1, for any equilibrium refinement Σ that is non-empty on finite models with a common prior, any finite model T, and any $k < \infty$, a prediction Q of (Σ, T) is robust if and only if Q is a prediction of (S^{∞}, T).*

Under the stronger richness assumption of Assumption 1, this result provides a characterization of robust predictions, which addresses all of the issues discussed above. Firstly, we only assume that Σ is non-empty on the finite models with a common prior. Since it is customary to prove such an existence result whenever a refinement is proposed, this assumption allows most equilibrium refinements. Secondly, our perturbation considers only the finite models with a common prior. Hence, the non-robustness implied by our characterization is not due to models that are usually assumed away by economists. Therefore, it is immune to the possible critique for earlier results based on models without a common prior, discussed in Section 2. Finally, we have a characterization in terms of usual rationalizability, muting the hope of obtaining positive robustness results in cases in which W^k is a small set.

We will now revisit the coordinated attack problem of Section 2 and illustrate how the methodology of global games can be extended to more general information structures and why the results will critically rely on the assumptions made on the information structure.

Example 2 (Izmalkov and Yildiz (2006))**.** In the coordinated attack problem of Section 2, assume that each player i observes a noisy signal $x_i = \theta + \varepsilon \eta_i$, where (η_1, η_2) is independently distributed from θ, with

$$\Pr_i \left(\eta_j > \eta_i | x_i\right) = q \qquad (\forall i \neq j) \tag{5.1}$$

for some constant $q \in (0,1)$, where $\Pr_i$ is the probability according to player i[18], and the support of θ contains an interval $[a, b]$ where $a < 0 < 5 < b$. That is, player i assigns probability q to the event that the other player is more optimistic ex-post. Carlsson and van Damme have shown that under the common-prior assumption,

[18] In this example, we do not have a common prior. We can construct a similar example with common prior, where the same unique actions survive iterated dominance, by using the Lipman construction as discussed above. In the new model, the signals will be two-dimensional, and the additive separability and independence assumptions of Carlsson and van Damme will be violated.

$q = 1/2$. Global games literature focuses on this case. Here, $q - 1/2$ measures the level of optimism of j according to i. Using the techniques and the mild assumptions of Carlsson and van Damme, one can easily check that, when $\varepsilon > 0$ is small, for each $x_i \neq 5(1-q)$, there is a unique rationalizable action $s_i^*(x_i)$, given by

$$s_i^*(x_i) = \begin{cases} \text{Attack} & \text{if } x_i > 5(1-q) \\ \text{No Attack} & \text{if } x_i < 5(1-q). \end{cases}$$

In the complete-information game, Attack is risk-dominant when $\theta > 5/2$, and No Attack is risk-dominant when $\theta < 5/2$. Under the common-prior assumption (i.e. when $q = 1/2$), players play according to risk dominance. In the general model, however, risk-dominance does not play any role. Given any x_i, we can make Attack uniquely rationalizable by choosing the level q of optimism sufficiently high, or make No Attack uniquely rationalizable by choosing the level q of optimism sufficiently low. Notice that, when ε is very small, the value of q has a very small impact on $x_i - x_j = \varepsilon(\eta_i - \eta_j)$ and on players' beliefs, and the players' beliefs converge to that of common knowledge at all orders as $\varepsilon \to 0$, but equilibrium behavior does not.

Taking θ as the vulnerability of economy, Morris and Shin (1998) have applied this idea to the currency-attack problem (where there is a continuum of players). They have illustrated that, attacks are closely related to the vulnerability of economy, as the likelihood of an attack is increasing with θ. They also noted that, in their model, investor sentiments do not play a role, which were given a prominent role in previous informal arguments based on multiple equilibria of the complete information game. In the above example, however, there is a precise measure concerning investor sentiments, namely $q - 1/2$, that determines belief in others' confidence in the economy. It plays an intuitive role similar to the vulnerability of the economy. Attack becomes more likely if the players become more optimistic or the economy becomes more vulnerable. That is, considering more general information structures, we can develop insightful models as in the global games literature, but focusing on their particular models which lead to risk-dominance is a very restrictive way of doing that.

Now, under the original assumptions of Carlsson and van Damme, despite the degenerate case of multiplicity in the complete-information game, there is an open set of nearby types that can only play risk-dominant equilibrium.[19] The global games literature would then argue that we ignore the complete-information case, even if that case is consistent with the researcher's observation of the situation. Nevertheless, when the complete-information case is consistent with his observation, there will also be an alternative open set of types that are consistent with his observation, leading to the opposite outcome of (Attack,Attack) as the only possibility—as in the above example. Then, the researcher will not be able to check whether the conditions assumed in the information structure of Carlsson and van Damme (1993) are satisfied even if he considers the complete-information to be degenerate. To put it differently, while the global games literature criticizes the multiplicity arguments for relying on simplifying common-knowlegde assumptions that we cannot check, our example shows that the same applies to their results. Their conclusions also critically depend on their simplifying assumptions on the noise structure.

More broadly, as Carlsson and van Damme (1993) and Kajii and Morris (1997) illustrate, by considering only some of the information structures that lead to a set of lower-order beliefs, one may be able to make sharper predictions. The analysis of such structure would be of great interest, but it is beyond the scope of this paper. Our paper shows that the resulting predictions will always critically depend on the assumptions built into these structures, and we cannot verify those assumptions and the conclusions with limited knowledge of the situation.

6. Without Full-Range Assumption

For ease of exposition, we have so far focused on equilibria with full range. Our full range assumption allowed us to consider large changes in higher-order beliefs. A researcher may be certain that it is common knowledge that the set of parameters are restricted to a small subset, or equivalently, the equilibrium considered may not vary much as the beliefs about the underlying uncertainty change. We will now present an extension of our main result to such cases. (For illustrations of proofs, see our earlier working paper.)

[19]When a type has unique rationalizable action, this action will be uniquely rationalizable for all types in an open neighborhood of this type (with respect to the product topology in the universal type space).

Local Rationalizability. For any $B_1 \times \cdots \times B_n \subset A$, define sets $S_i^k [B; t_i]$, $S_i^k [B; t_i]$, $i \in N$, $k \in \mathbb{N}$, $t_i \in T_i$, as before but set by setting

$$S_i^0 [B; t_i] = W_i^0 [B; t_i] = B_i.$$

Unlike before, the new sets can become larger as k increases. Hence, we define the limit sets by

$$\begin{aligned} S_i^\infty [B; t_i] &= \bigcap_{k=0}^{\infty} \bigcup_{m=k}^{\infty} S_i^m [B; t_i], \\ W_i^\infty [B; t_i] &= \bigcap_{k=0}^{\infty} \bigcup_{m=k}^{\infty} W_i^m [B; t_i]. \end{aligned}$$

Proposition 7. *For any equilibrium s^*, if the game has countable action spaces, then*

$$W_i^k [s^* (T^u); t_i] \subseteq A_i^k [s^*, t_i] \subseteq S_i^k [s^* (T^u); t_i] \quad (\forall i, k, t_i);$$

if the game is nice, then with notation of Proposition 2, for any $B \subseteq s^* (T)$,

$$S_i^k \left[B; \hat{t}_i\right] \subseteq A_i^k [s^*, t_i] = S_i^k \left[s^* (T^u); \hat{t}_i\right] \quad \left(\forall i, k, \hat{t}_i\right).$$

Replacing W^∞ and S^∞ with their local versions, this proposition establishes that the bounds for A^∞ provided by Propositions 4 and 5 remains valid even if one does not assume that the equilibrium has full range. Then, using this result instead of Propositions 1 and 2, we can obtain similar bounds to those of Propositions 4 and 5 without assuming full range assumption, but instead replacing W^∞ and S^∞ with their local versions for some set of actions which are known to be played by some types.

The last statement in the proposition implies that, for nice games, even the slight changes in very higher-order beliefs will have substantial impact on equilibrium behavior, unless the game is locally dominance-solvable. There are important games in which a slight failure of common knowledge assumption in very high orders leads to substantially different outcomes—as we show next.

7. Applications

In this section, we will provide three applications of our results. First, we will show that in Cournot oligopoly with sufficiently firms, even a slight relaxation of the common-knowledge assumption will preclude us from making any prediction beyond the elementary fact that no firm will produce more than the monopoly outcome. Second, we will show that an equilibrium (or rationalizable) strategy is continuous with respect to the product topology at a given type *if and only if* that type has a unique rationalizable action. Third, we will show that, under our perturbation, there is a robust equilibrium if and only if the game is dominance solvable.

7.1. **Cournot Oligopoly.** In a Cournot oligopoly with sufficiently many firms, any production level that is less than or equal to the monopoly production is rationalizable (Bernheim (1984), Basu (1992)).[20] Then, our Proposition 2 implies that a researcher cannot rule out any such output level as the *equilibrium* output for a firm no matter how many orders of beliefs he specifies, assuming the set of payoff parameters is rich. Using Proposition 7, we will now show that the richness assumption is not needed for this conclusion. Even a slight doubt about the payoffs in very high orders will lead a researcher to fail to rule out any outcome that is less than the monopoly outcome as a firm's equilibrium output. More broadly, this establishes that our non-robustness results apply to some very important games in economics, even without our richness assumptions.

Consider n firms with identical constant marginal cost $c > 0$. Simultaneously, each firm i produces q_i at cost $q_i c$ and sell its output at price $P(Q;\theta)$ where $Q = \sum_i q_i$ is the total supply. For some fixed $\bar{\theta}$, we assume that Θ is a closed interval with $\bar{\theta} \in \Theta \neq \{\bar{\theta}\}$. We also assume that $P(0;\bar{\theta}) > 0$, $P(\cdot;\bar{\theta})$ is strictly decreasing when it is positive, and $\lim_{Q\to\infty} P(Q;\bar{\theta}) = 0$. Therefore, there exists a unique $\hat{Q}$ such that $P(\hat{Q};\bar{\theta}) = c$. We assume that, on $[0,\hat{Q}]$, $P(\cdot;\bar{\theta})$ is continuously twice-differentiable and $P' + QP'' < 0$.

It is well known that, under the assumptions of the model, (i) the profit function, $u(q,Q;\bar{\theta}) = q(P(q+Q) - c)$, is strictly concave in own output q; (ii) the unique best response $q^*(Q_{-i})$ to others' aggregate production Q_{-i} is strictly decreasing on $[0,\hat{Q}]$ with slope bounded away from 0 (i.e., $\partial q^*/\partial Q_{-i} \leq \lambda$ for some $\lambda < 0$); (iii) equilibrium outcome at $t^{CK}(\bar{\theta})$, $s^*(t^{CK}(\bar{\theta}))$, is unique and symmetric (Okuguchi

[20]Borgers and Janssen (1995) show that if we replicate both consumers and the firms in such a way that the cobweb dynamics is stable for the resulting demand and supply curves, then the Cournot oligopoly will be dominance-solvable. In that case, by Proposition 2, equilibrium outcomes will not be sensitive to higher-order beliefs.

and Suzumura (1971)). We will also assume that θ is a payoff-relevant parameter in the following sense: $q^*(Q_{-i};\theta)$ is a continuous and strictly increasing function of θ at $\left(Q_{-i},\bar{\theta}\right)$ where $Q_{-i}=(n-1)\,s_j^*\left(t^{CK}\left(\bar{\theta}\right)\right)$.

Lemma 3. *In the Cournot oligopoly above, there exists $\bar{n}<\infty$ such that for any $n>\bar{n}$ and any $B=\left[s_1^*\left(t_1^{CK}\left(\bar{\theta}\right)\right)-\epsilon, s_1^*\left(t_1^{CK}\left(\bar{\theta}\right)\right)+\epsilon\right]^n\subset A$ with $\epsilon>0$, we have*

$$S_i^\infty\left[B;t^{CK}\left(\bar{\theta}\right)\right]=\left[0,q^M\right]\quad(\forall i\in N),$$

where q^M is the monopoly output under $P\left(\cdot;\bar{\theta}\right)$ and $s^*\left(t^{CK}\left(\bar{\theta}\right)\right)$ is the unique equilibrium of the complete information game $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$.

This is a straightforward extension of a result by Basu (1992) for rationalizability to local rationalizability. The proof is in the appendix. Together with Proposition 7, this lemma yields the following.

Proposition 8. *In the Cournot oligopoly above, let $\Theta=\left[\bar{\theta}-\varepsilon,\bar{\theta}+\varepsilon\right]$ for arbitrarily small $\varepsilon>0$. Then, for any equilibrium s^* on the universal type space,*

$$A_i^\infty\left[s^*,t_i^{CK}\left(\bar{\theta}\right)\right]=\left[0,q^M\right]\quad(\forall i\in N),$$

where q^M is the monopoly output under $P\left(\cdot;\bar{\theta}\right)$.

Proof. Since we can put a large upper bound on q, by (i) above, we have a nice game. By the hypothesis, there exists $B\subset s^*(T)$ as in Lemma 3. Hence, Lemma 3 and Proposition 7 imply

$$\left[0,q^M\right]=S_i^\infty\left[B;t_i^{CK}\left(\bar{\theta}\right)\right]\subseteq A_i^\infty\left[s^*,t_i^{CK}\left(\bar{\theta}\right)\right]\subseteq\left[0,q^M\right],$$

yielding the desired equality. □

Our proposition suggests that, with sufficiently many firms, any equilibrium prediction that is not implied by strict dominance will be invalid whenever we slightly deviate from the idealized complete information model. To see this, consider two researchers. One is confident that it is common knowledge that $\theta=\bar{\theta}$. The other is slightly skeptical: he is only willing to concede that it is common knowledge that $\left|\theta-\bar{\theta}\right|\leq\varepsilon$ and agrees with the kth-order mutual knowledge of $\theta=\bar{\theta}$. He is an arbitrarily generous skeptic; he is willing to concede the above for arbitrarily small $\varepsilon>0$ and arbitrarily large finite k. Our proposition states that the skeptic nonetheless cannot rule out any output level that is not strictly dominated.

7.2. **Continuity of equilibrium.** It is well-known that some equilibria of some games are discontinuous in the universal type space under the product topology. For example, if we combine the e-mail game of Rubinstein (1989) with the complete information case, there will be one continuous equilibrium, prescribing No Attack everywhere, and one discontinuous equilibrium, prescribing Attack only in the complete information case. After all, best response of a single player may be discontinuous when A is finite. Similarly, a Nash equilibrium may be discontinuous function of complete information games, while the Nash equilibrium correspondence is generically lower-hemicontinuous on the space of complete-information games. Now we will show that, in nice and finite-action games every equilibrium strategy will be discontinuous at every type with multiple rationalizable actions.

We consider an arbitrary metric d_A on A. (When A is finite, d_A will be equivalent to the discrete metric.) A sequence $(a^m)_{m\in\mathbb{N}}$ is said to *converge* to some $a \in A$ iff for each $\epsilon > 0$, there exists k such that $d_A(a^m, a) < \epsilon$ for each $m > k$. When A is finite this simplifies to the requirement that $a^m = a$ for $m > k$.

Definition 8. A strategy profile s is said to be *continuous (with respect to product topology) at t* iff for each sequence $\left(\tilde{t}\,[m]\right)_{m\in\mathbb{N}}$ of type profiles

$$\left[\tilde{t}^k\,[m] \to t^k \quad \forall k\right] \Rightarrow \left[s\left(\tilde{t}\,[m]\right) \to s(t)\right].$$

Proposition 2 implies that if an equilibrium s^* with full range is continuous at t, then $S^\infty[t] = \{s^*(t)\}$, yielding the following discontinuity result.

Proposition 9. *For any nice game, every equilibrium $s^* : T^u \to A$ with full range is discontinuous at each type profile t for which there are more than one rationalizable action profiles. In particular, if a nice game possesses an equilibrium s^* that is continuous and has a full range, then the game is dominance solvable.*

Proof. Take any equilibrium s^* with full range and any t with $a \in S^\infty[t]$ such that $s^*(t) \neq a$. Then, for each k, there exists $\tilde{t}\,[k]$ such that $\tilde{t}^m\,[k] = t^m$ for each $m \leq k$ and $s^*\left(\tilde{t}\,[k]\right) = a \neq s^*(t)$. Clearly, $s^*\left(\tilde{t}\,[k]\right)$ does not converge to $s^*(t)$. But, by definition, for each m and each $k > m$, $\tilde{t}^m\,[k] = t^m$, and hence $\tilde{t}^m\,[k] \to t^m$ as $k \to \infty$. Therefore, s^* is discontinuous at t. $\square$

Under a stability condition, Nyarko (1996) shows that equilibrium is continuous on universal type space. Our result then shows that his condition implies that the game is dominance-solvable. Considering a class of simple dominance-solvable games that satisfies the stability condition of Nyarko (1996), Morris (2002) shows, among other things, that continuity of equilibrium is not uniform over these games. That is, even in dominance-solvable games, we may need to pay special attention to how we specify the higher-order beliefs.

One can obtain a discontinuity result for countable-action games at each type with $|W^{\infty}[t]| > 1$, similar to the previous proposition. For mixed-strategy equilibria, one can obtain similar discontinuity results under substantially weaker continuity requirements (see our working paper). We will now obtain a characterization for all rationalizable strategies in finite-action games.

Proposition 10. *Assume that A is finite. Under Assumption 1, a rationalizable strategy profile $s : T^u \to A$ is continuous at a finite type profile $\hat{t}$ if and only if $\hat{t}$ has a unique rationalizable action, i.e., $\left|S^{\infty}\left[\hat{t}\right]\right| = 1$. This characterization remains intact if the domain of s is $T^{CPA} = \{t | t \in T,\ T$ is finite and has a common prior$\}$.*

Proof. The "if" part follows from the fact that S^{∞} is upper-semicontinuous (Dekel, Fudenberg, and Morris (2004)). To prove the "only if" part, take any rationalizable strategy profile s and any $a \in S^{\infty}\left[\hat{t}\right]$ with $s\left(\hat{t}\right) \neq a$. By Lemma 2, there exists a sequence of types $t[m] \to \hat{t}$ with $s(t[m]) \in S^{\infty}[t[m]] = \{a\}$. Since $s(t[m]) = a$ for each $t[m]$, $s(t[m])$ does not converge to $s\left(\hat{t}\right)$. The last statement follows from the fact that $t[m] \in T^{CPA}$ in Lemma 2. $\square$

Proposition 10 establishes that, at each finite type, either there is a unique rationalizable action and all rationalizable strategies are continuous (in fact, locally constant) at that type, or there are multiple rationalizable actions and all rationalizable strategies are discontinuous at that type. As discussed in Section 2, this raises a serious difficulty in usefulness of continuity arguments to refine equilibrium (or rationalizability). Whenever there is a possibility of refining rationalizability, all of the rationalizable strategies are discontinuous at the relevant type, and imposing continuity leads to the empty set. Note that generically there exists a unique rationalizable outcome, and hence all rationalizable strategies are continuous at generic types, where there is no need for a refinement.

7.3. **Robustness of equilibrium.** Kajii and Morris (1997) introduced an ex ante notion of robustness of equilibrium to incomplete information. We will now investigate how the results will change if we consider an analogous interim notion of robustness for equilibrium, or drop the common-prior assumption in their formulation. Following Kajii and Morris (1997), we will assume that the set of payoff profiles is not restricted, so that for each payoff profile $v : A \to \mathbb{R}^N$, there exists θ such that $u(\theta, \cdot) = v$. With our notation, for pure strategy equilibria, their robustness can be defined as follows.

Definition 9 (Kajii and Morris (1997))**.** An equilibrium $\hat{a} \in A$ of a complete-information game $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$ is said to be *robust to incomplete information* iff for every $\delta > 0$, there exists $\varepsilon > 0$ such that for each finite T with common prior p such that $p\left(\left\{(\theta, t) \,|u_i(\theta, \cdot) = u_i\left(\bar{\theta}, \cdot\right), \kappa_{t_i}\left(u_i(\theta, \cdot) = u_i\left(\bar{\theta}, \cdot\right)\right) = 1 \forall i\right\}\right) \geq 1 - \varepsilon$, there exists an equilibrium σ^* of T with $p\left(\operatorname{supp}\left(\sigma^*(t)\right) = \{\hat{a}\}\right) \geq \delta$.

That is, if the common prior of T puts high probability on the event that the payoffs are as described in the complete-information game and everybody knows his payoffs, then T will have an equilibrium in which most of the types will play according to the original equilibrium. Here, there is no restriction on players' interim beliefs, other than what the restriction on the prior implies. This definition envisions a researcher who believes that his complete information model is the true model with high probability according to the prior of the correct model. As we discussed before, we are interested in the robustness in the interim stage when we do not have any information about the ex ante stage, which the researcher constructs to model the interim beliefs. Accordingly, we define interim robustness, again for pure strategy equilibria of complete information games, as follows.

Definition 10. An equilibrium $\hat{a} \in A$ of a complete-information game $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$ is said to be *interim robust to incomplete information* iff there exists $\varepsilon > 0$ and $k < \infty$ such that for each (ε, k)-perturbation (T, τ) of $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$, there exists an equilibrium σ^* of T with $\operatorname{supp}\left(\sigma^*\left(\tau\left(t^{CK}\left(\bar{\theta}\right)\right)\right)\right) = \{\hat{a}\}$, where T is finite and has a common prior.

In this definition, we keep their requirement that the perturbed game has a common prior, but we do not require the perturbed game to assign a high probability to the original model. Instead, we require only that the perturbed model has a type that has similar beliefs in the interim stage. Also, instead of requiring that most of the types play according to the original equilibrium, we only require the type with similar beliefs to do so.

Kajii and Morris (1997) showed that, although a unique Nash equilibrium need not be robust to incomplete information, if there is a unique correlated equilibrium, then it will be robust. Moreover, any $(p_1, \ldots, p_n)$-dominant equilibrium with $p_1 + \cdots + p_n < 1$, such as a risk-dominant equilibrium, is robust. These are stringent sufficient conditions, but they yield existence of robust equilibria in many cases, such as generic two-player, two-action games. In contrast, we will now show that existence of interim robust equilibrium is equivalent to dominance-solvability. Therefore, when we do not impose that the prior of the correct model assigns high probability to the actual case, there will be significantly fewer robust equilibria, even if we impose stringent conditions on the beliefs of the perturbed type.

Proposition 11. *Let A be finite and let $\bar{\theta}$ have generic payoffs. Under Assumption 1, $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$ has an interim robust equilibrium $\hat{a}$ if and only if there is a unique rationalizable action profile for $t^{CK}\left(\bar{\theta}\right)$.*

Proof. The "if" part follows from the fact that S^∞ is upper-semicontinuous, so that whenever $S^\infty\left[t^{CK}\left(\bar{\theta}\right)\right] = \{\hat{a}\}$, we have $S^\infty[t] = \{\hat{a}\}$ in a neighborhood of $t^{CK}\left(\bar{\theta}\right)$. Then, when ε is small and k large, for every (ε, k)-perturbation (T, τ), $S^\infty\left[\tau\left(t^{CK}\left(\bar{\theta}\right)\right)\right] = \{\hat{a}\}$. Letting σ^* be any equilibrium of T, we conclude that $\mathrm{supp}\left(\sigma^*\left(\tau\left(t^{CK}\left(\bar{\theta}\right)\right)\right)\right) = \{\hat{a}\}$. To prove the "only if" part, assume that $\left|S^\infty\left[t^{CK}\left(\bar{\theta}\right)\right]\right| > 1$, and take any equilibrium $\hat{a}$ of $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$. There exists $a \in S^\infty\left[t^{CK}\left(\bar{\theta}\right)\right]$ with $a \neq \hat{a}$. Then, by Lemma 2, for each $\varepsilon > 0$ and $k < \infty$, there exists (ε, k)-perturbation (T, τ) of $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$ such that $S^\infty\left[\tau\left(t^{CK}\left(\bar{\theta}\right)\right)\right] = \{a\}$. Then, for any equilibrium σ of T, $\sigma\left(\tau\left(t^{CK}\left(\bar{\theta}\right)\right)\right)$ must assign zero probability to $\hat{a}$. Therefore, $\hat{a}$ is not interim robust. □

What happens if we drop the common-prior assumption in Kajii-Morris definition of robustness, keeping their (prior, rather than interim) notion of perturbation? To answer this question, let us modify Definition 9.

Definition 11. An equilibrium $\hat{a} \in A$ of a complete-information game $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$ is said to be *weakly robust to incomplete information without CPA* iff there exists

$\varepsilon > 0$ such that for each finite T, defined by $\kappa_{t_i} \equiv p_i(\cdot|t_i)$ for priors $(p_1,\dots,p_n)$ and such that $p_i\left(\left\{(\theta,t)\,|u_j(\theta,\cdot) = u_j\left(\bar{\theta},\cdot\right), \kappa_{t_j}\left(u_j(\theta,\cdot) = u_j\left(\bar{\theta},\cdot\right)\right) = 1\forall j\right\}\right) \geq 1-\varepsilon$ for each i, there exist an equilibrium σ^* of T and $t \in T$ such that σ^* assigns positive probability to $\hat{a}$ at t, and t is assigned a positive probability by some p_i.

Above, we drop the common-prior assumption in the Kajii-Morris definition, making robustness more stringent. We also weaken the robustness condition significantly by simply requiring that $\hat{a}$ is played in equilibrium with positive probability according to some player, as opposed to $\hat{a}$ being the pure equilibrium outcome with high probability. We show that the impact of dropping the common-prior assumption is so strong that now robustness roughly implies that we have a dominant-strategy equilibrium. This is because without a common prior, the restrictions on prior beliefs have no implications for interim beliefs beyond second order.

Proposition 12. *For $N = \{1,2\}$, let $\hat{a} = (\hat{a}_1,\hat{a}_2)$ be a Nash equilibrium of $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$ where neither of $\hat{a}_i$ is a dominant strategy in $\left\{t^{CK}\left(\bar{\theta}\right)\right\}$. Then, $\hat{a}$ is not weakly robust to incomplete information without CPA.*

Proof. We will define, for arbitrary ε, a model T as in Definition 11 which will not possess any equilibrium with the desired properties. For each i, since $\hat{a}_i$ is not a dominant strategy, there exist $\bar{a}_i$ and $\tilde{a}_{-i}$ such that $u_i(\bar{a}_i,\tilde{a}_{-i}) > u_i(\hat{a}_i,\tilde{a}_{-i})$. Let θ_i be a state at which $\tilde{a}_i$ is a strictly dominant action and the payoffs of the other player j are as in $\bar{\theta}$. Define T with $T_i = \{t_i,t_i',t_i''\}$, $i \in N$, with prior beliefs[21] (p_1,p_2) given in the following tables for each triple in $\Theta \times T_1 \times T_2$

$\bar{\theta}$	t_2	t_2'	t_2''
t_1	0,0	$1-\varepsilon-\varepsilon^2$,$\varepsilon^2$	0,0
t_1'	ε^2,$1-\varepsilon-\varepsilon^2$	0,0	0,0
t_1''	0,0	0,0	0,0

θ_1	t_2	t_2'	t_2''
t_1	0,0	0,0	0,0
t_1'	0,0	0,0	0,0
t_1''	0,0	$\frac{\varepsilon}{2}$,$\frac{\varepsilon}{2}$	0,0

θ_2	t_2	t_2'	t_2''
t_1	0,0	0,0	0,0
t_1'	0,0	0,0	$\frac{\varepsilon}{2}$,$\frac{\varepsilon}{2}$
t_1''	0,0	0,0	0,0

Now, define $F = \left\{\left(\bar{\theta},t_1,t_2'\right),\left(\bar{\theta},t_1',t_2\right)\right\}$. On F players' payoffs are given as in $\bar{\theta}$, and each player knows his own payoffs, i.e., both types t_i and t_i' put zero probability on $\theta = \theta_i$. Since $p_1(F) = p_2(F) = 1-\varepsilon$, T satisfies the condition in Definition 11. However, we will show that there does not exist any equilibrium σ^*

[21]Notice that p_1 and p_2 have common support, and hence our result would not change if we imposed the common-support assumption.

where $\hat{a}$ is played with positive probability anywhere on the common support of p_1 and p_2. Take any equilibrium σ^*. First, each type t_i'' puts probability 1 on $\theta = \theta_i$. Hence, by definition of θ_i, type t_i'' must play $\tilde{a}_i$ with probability 1. Now, type t_i' puts probability $2\varepsilon/(2\varepsilon+1)$ on $(\bar{\theta}, t_{-i})$ and probability $1/(2\varepsilon+1)$ on (θ_{-i}, t''_{-i}). When ε is small, type t_i' puts nearly probability 1 on type t''_{-i}, who plays $\tilde{a}_{-i}$. Then, under equilibrium beliefs, $\bar{a}_i$ yields higher expected payoff than $\hat{a}_i$, and hence t_i' puts zero probability on $\hat{a}_i$. Everywhere in the common support of p_1 and p_2, one of the types is either of the form t_i' or t_i'', which put zero probability on $\hat{a}_i$, and hence σ^* puts zero probability on $\hat{a}$. □

Oyama and Tercieux (2005) extend Proposition 12 to n players. They show that if an equilibrium is weakly robust to incomplete information without CPA, then $n-1$ players play dominant strategies. As mentioned above, the extreme lack of robustness in these results stems from the fact that without CPA, restrictions on priors do not put any restriction on beliefs at second order and higher. Then, Proposition 1 suggests that robustness of an equilibrium would require that W^1 is singleton, leading to Proposition 12. Oyama and Tercieux (2005) also consider a weaker robustness notion for equilibrium without CPA, combining our approach of restricting higher-order beliefs with robustness of Kajii and Morris (1997). They consider perturbations where each player's prior puts high probability on the event that the first k orders of beliefs are as specified, for arbitrarily large k. Consistent with Proposition 1, they show that there is no robust equilibrium in this very weak sense whenever $|W^\infty| > 1$.

8. Conclusion

Most predictions in economics are based on some equilibrium refinement, applied to specific models. In this paper, we recognize that, when we need to check whether such a prediction applies to a particular situation, we could only know finitely many orders of players' interim beliefs. There are many models compatible with this partial information. We ask which predictions remain valid under our limited information, in the sense that we could make that prediction independent of which compatible model is chosen, given that we will apply the same refinement at each model. We show that predictions remain valid only if they are true for all strategies that survive iterated elimination of actions that are never a strict best reply. For generic payoffs, we find a characterization: a prediction remains valid if and only if it is true for all rationalizable strategies. Therefore, equilibrium refinements will be useful in generating extra predictions only when we have information about the entire infinite hierarchy of beliefs. We prove this by establishing that every equilibrium refinement has to be highly sensitive to specification of higher-order beliefs when there are multiple rationalizable actions. This sensitivity is what drives our result. We also show that our result remains intact when we restrict our attention to small (i.e. finite) type spaces, or impose the common-prior assumption.

Appendix A. Omitted Proofs

Lemma 4 (Moulin (1984) and Battigalli (2003))**.** *For any nice game, for any i, t_i, k, and any $a_i \in S_i^k [t_i]$, there exists a pure strategy $\hat{s}_{-i}$ such that $\hat{s}_{-i}(t_{-i}) \in S_{-i}^{k-1}[t_{-i}]$ for each t_{-i} and*

$$BR_i(\pi(\cdot|t_i, \hat{s}_{-i})) = \{a_i\}.$$

Proof of Proposition 2. For any $a_i \in S_i^k[t_i] = W_i^k[t_i]$, by Lemma 4, there exists $\hat{s}_{-i}$ such that $\hat{s}_{-i}(\cdot) \in S_{-i}^{k-1}[\cdot] = S_{-i}^{k-1}[\cdot]$ and a_i is a strict best reply against $\pi(\cdot|t_i, \hat{s}_{-i})$. Since κ_{t_i} has countable support, $P(\cdot|t_i, \hat{s}_{-i})$, the probability distribution induced by κ_{t_i} and $\hat{s}_{-i}$ on $\Theta \times L \times A_{-i}$, has a countable support:

$$\mathrm{supp} P(\cdot|t_i, \hat{s}_{-i}) = \{(\theta, l, \hat{s}_{-i}(\theta, l, h)) \,|\, (\theta, l, h) \in \mathrm{supp}\kappa_{t_i}\}.$$

Hence our proof of Proposition 1 applies. That is, there exists $\tilde{t}_i \in T_i$ (not necessarily in T_i) such that $s_i^*(\tilde{t}_i) = a_i$ and $\tilde{t}_i^m = t_i^m$ for each $m \leq k$, yielding the equality above. $\square$

Proof of Lemma 3. Let $\bar{n}$ be any integer greater than $1 + 1/|\lambda|$, where λ is as in (ii). Take any $n > \bar{n}$. By (iii), $B = [\underline{q}^0, \bar{q}^0]^n$ for some $\underline{q}^0, \bar{q}^0$ with $\underline{q}^0 < \bar{q}^0$. By (ii), for any $k > 0$, $S^k\left[B; t^{CK}(\bar{\theta})\right] = [\underline{q}^k, \bar{q}^k]^n$, where

$$\bar{q}^k = q^*\left((n-1)\,\underline{q}^{k-1}\right) \text{ and } \underline{q}^k = q^*\left((n-1)\,\bar{q}^{k-1}\right).$$

Define $\underline{Q}^k \equiv (n-1)\,\underline{q}^k$, $\bar{Q}^k \equiv (n-1)\,\bar{q}^k$, and $Q^* = (n-1)\,q^*$, so that

$$\bar{Q}^k = Q^*\left(\underline{Q}^{k-1}\right) \text{ and } \underline{Q}^k = Q^*\left(\bar{Q}^{k-1}\right).$$

Since $(n-1)\,\lambda < 1$, the slope of Q^* is strictly less than -1. Hence $\underline{Q}^k$ decreases with k and becomes 0 at some finite $\bar{k}$, and $\bar{Q}^k$ increases with k and takes value $Q^*(0) = (n-1)\,q^M$ at $\bar{k}+1$. That is, $S^k\left[B; t^{CK}(\bar{\theta})\right] = \left[0, q^M\right]^n$ for each $k > \bar{k}$. Therefore, $S^\infty\left[B; t^{CK}(\bar{\theta})\right] = \left[0, q^M\right]^n$. $\square$

References

[1] Aumann, Robert (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1-18.

[2] Basu, K. (1992): "A characterization of the class of rationalizable equilibria of oligopoly games," *Economics Letters* 40, 187-191.

[3] Battigalli, P. (2003): "Rationalizability in infinite, dynamic games with complete information," *Research in Economics*, 57, 1-38.

[4] Battigalli, P. and M. Siniscalchi (2002): "Strong Belief and Forward-Induction Reasoning," *Journal of Economic Theory* 106, 356-391.

[5] Battigalli, P. and M. Siniscalchi (2003) "Rationalization and Incomplete Information," *Advances in Theoretical Economics* 3-1, Article 3.

[6] Bergemann, D. and S. Morris (2003): "Robust Mechanism Design," mimeo.

[7] Bernheim, D. (1984): "Rationalizable strategic behavior," *Econometrica*, 52, 1007-1028.

[8] Börgers T. (1994): "Weak Dominance and Approximate Common Knowledge," *Journal of Economic Theory*, 64, 265-276.

[9] Börgers T. and M. Janssen (1995): "On the dominance solvability of large Cournot games," *Games and Economic Behavior* 8, 297-321.

[10] Brandenburger, A. and E. Dekel (1987): "Rationalizability and Correlated Equilibria," *Econometrica*, 55, 1391-1402.

[11] Brandenburger, A. and E. Dekel (1993): "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189-198.

[12] Carlsson, H. and E. van Damme (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61, 989-1018.

[13] Cremer, J. and. R. McLean (1988): "Full extraction of the Surplus in Bayesian and Dominant Strategy Auctions", *Econometrica*, 56, 1247-1257.

[14] Dekel, E. and D. Fudenberg (1990): "Rational Behavior with Payoff Uncertainty," *Journal of Economic Theory*, 52, 243-267.

[15] Dekel, E. D. Fudenberg, S. Morris (2003): "Interim Rationalizability," mimeo.

[16] Dekel, E. D. Fudenberg, S. Morris (2004): "Dense Types," mimeo.

[17] Dekel, Eddie and Faruk Gul (1997): "Rationality and Knowledge in Game Theory," in Kreps, D. and K. Wallis (eds): *Advances in Economics and Econometrics: Theory and Applications*, Cambridge University Press, UK.

[18] Ely, J. and M. Peski (2004): "Hierarchies of beliefs and interim rationalizability," mimeo.

[19] Feinberg, Y. (2002): "Subjective Reasoning in Dynamic Games," forthcoming in *Games and Economic Behavior*.

[20] Feinberg, Y. and A. Skrzypacz (2002): "Uncertainty about uncertainty and delay in bargaining," forthcoming in *Econometrica*.

[21] Harsanyi, J. (1967): "Games with Incomplete Information played by Bayesian Players. Part I: the Basic Model," *Management Science* 14, 159-182.

[22] Heifetz, A. and Z. Neeman (2003): "On the Generic Infeasibility of Full Surplus Extraction in Mechanism Design," mimeo.

[23] Kajii, A. and S. Morris (1997): "The Robustness of Equilibria to Incomplete Information," *Econometrica*, 65, 1283-1309.

[24] Kajii, A. and S. Morris (1998): "Payoff Continuity in Incomplete Information Games," *Journal of Economic Theory*, 82, 267-276.

[25] Lipman, B. (2003): "Finite Order Implications of Common Priors," *Econometrica*, 71, 1255-1267.

[26] Mertens and Zamir (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 10, 619-632.

[27] Monderer, D. and D. Samet (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170-190.

[28] Morris, S. (2002): "Typical Types," mimeo.

[29] Morris, S. and H. Shin (1998): "Unique Equilibrium in a Model of Self-Fulfilling Attacks," *American Economic Review* 88, 587-597

[30] Morris, S. and Ui (2002): "Generalized Potentials and Robust Sets of Equilibria," mimeo.

[31] Moulin, H. (1984): "Dominance solvability and Cournot stability," *Mathematical Social Sciences* 7, 83-102.

[32] Neeman, Z. (2004) "The Relevance of Private Information in Mechanism Design," *Journal of Economic Theory* 117, 55-77.

[33] Nyarko, Y. (1996): "Convergence in Economic Models with Bayesian Hierarchies of Beliefs," *Journal of Economic Theory*, 74, 266-296.

[34] Okuguchi, K. and Suzumura, K. (1971): "Uniqueness of the Cournot Oligopoly Equilibrium," *Economic Studies Quarterly*, 22, 81-83.

[35] Oyama, D. and O. Tercieux (2005): "Strategic Implications of (Non-)Common Priors II: Robustness of Equilibria," mimeo (University of Tokyo and PSE).

[36] Pearce, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029-1050.

[37] Rubinstein, A. (1989): "The Electronic Mail Game: Strategic Behavior Under 'Almost Common Knowledge'," *The American Economic Review*, Vol. 79, No. 3, 385-391.

[38] Simon, R. (2003): "Games of Incomplete Information, Ergodic Theory, and Measurability of Equilibria," *Israeli Journal of Mathematics* 138, 73-92.

[39] Tan, T. and S. Werlang (1988): "The Bayesian foundations of solution concepts of games," *Journal of Economic Theory* 45, 370-391

[40] Townsend, R. (1983): "Forecasting the Forecasts of Others," *Journal of Political Economy*, 9-4, 546-88.

[41] Ui, T. (2001). "Robust Equilibria of Potential Games," *Econometrica* 69, 1373-1380.

[42] Weinstein, J. and M. Yildiz (2003): "Impact of Higher-Order Uncertainty," *MIT Department of Economics Working Paper* 03-14.

[43] Wilson, R. (1987): "Game-Theoretic Analyses of Trading Processes," in: Truman Bewley (ed.) *Advances in Economic Theory: Fifth World Congress*, Cambridge UK: Cambridge University Press, 33-70.

[44] Yildiz (2005): "Generic Uniqueness of Rationalizable Actions," mimeo.

MIT

URL: `http://econ-www.mit.edu/faculty/myildiz/index.htm`