



Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Sensitivity of equilibrium behavior to higher-order beliefs in nice games[☆]Jonathan Weinstein^{a,*}, Muhamet Yildiz^b^a Northwestern University, Evanston, IL, United States^b MIT, Cambridge, MA, United States

ARTICLE INFO

Article history:

Received 14 April 2009

Available online xxxx

JEL classification:

C72

C73

Keywords:

Higher-order beliefs

Incomplete information

Robustness

Sensitivity

Universal type space

ABSTRACT

We analyze “nice” games (where action spaces are compact intervals, utilities continuous and strictly concave in own action), which are used frequently in classical economic models. Without making any “richness” assumption, we characterize the sensitivity of any given Bayesian Nash equilibrium to higher-order beliefs. That is, for each type, we characterize the set of actions that can be played in equilibrium by some type whose lower-order beliefs are all as in the original type. We show that this set is given by a local version of interim correlated rationalizability. This allows us to characterize the robust predictions of a given model under arbitrary common knowledge restrictions. We apply our framework to a Cournot game with many players. There we show that we can never robustly rule out any production level below the monopoly production of each firm.

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

Most economic applications, even those that model incomplete information, fix a specific type space. They thus make common-knowledge assumptions that are difficult to verify in the modeling stage. Unfortunately, as in the well-known e-mail game example of Rubinstein (1989), the equilibrium behavior may be highly sensitive to these assumptions, so the researcher may not be able to know whether his predictions are valid. Indeed, for finite action games, we showed in Weinstein and Yildiz (2007a) [hereafter WY] that whenever there are multiple rationalizable actions and the space of basic uncertainty is rich, rationalizable strategies are highly sensitive to common-knowledge assumptions. In fact, by relaxing these assumptions suitably one can make any rationalizable action the unique solution. In this paper, changing our solution concept to Bayesian Nash equilibria, we extend this result in two important directions.

First, this paper relaxes the richness assumption in WY. There, we assume that the set of underlying payoff parameters is rich enough so that each action could be dominant at some parameter value. In other words, there are no common-knowledge restrictions on payoffs. An application, however, may impose a natural structure on payoffs, and the researcher may be willing to assume that this structure is common knowledge. For example, it may be natural to assume that the bidders in an auction care only about whether they win the object and how much they pay in case they win. Such a bidder would be indifferent towards how he loses the object, and in that case submitting a low bid cannot be a strictly dominant

[☆] This paper is based on an earlier working paper, Weinstein and Yildiz (2004). We thank Daron Acemoglu, Pierpaolo Battigalli, Tilman Börgers, Glenn Ellison, Drew Fudenberg, Stephen Morris, an anonymous associate editor, two anonymous referees, and the seminar participants at various universities and conferences. We thank Institute for Advanced Study for their generous support and hospitality. Yildiz also thanks Princeton University and Cowles Foundation of Yale University for their generous support and hospitality.

* Corresponding author.

E-mail address: j-weinstein@kellogg.northwestern.edu (J. Weinstein).

action for him. If the researcher is willing to assume that this is indeed common knowledge, then the richness assumption of WY would fail. In this paper, we show that we can characterize the sensitivity of equilibrium strategies without the full richness assumption. The characterization now depends on a local version of interim correlated rationalizability, rather than the usual interim correlated rationalizability defined in Dekel et al. (2007).¹

Second, we generalize from finite-action games to nice games (Moulin, 1984), which are commonly used in classical economic models. In these games, the action spaces are compact intervals, and the utility functions are continuous and strictly concave in own action, as in Cournot competition and differentiated Bertrand competition. The assumption of finite-action games with finite types is used in WY to ensure measurability of a certain constructed mapping. There is no way to completely solve the measurability problem for general infinite games, but here we are able to solve the measurability problem by showing that rationalizability in nice games has a special structure. The keys are that in nice games, any rationalizable action is a best reply to a *deterministic* theory of the other players' actions, and that at each step in the elimination process, the set of remaining actions is an interval. A further advantage of nice games is the uniqueness of best replies. This avoids our having to allow small perturbations to lower-order beliefs; instead we can assume the parameter is mutually known up to arbitrary order, with no reference to a topology on each order of beliefs. The use of a fixed equilibrium also plays a role, by telling us that a player has a consistent theory of the actions taken by each type of the other players.

Since we do not make any richness assumption on the set of payoff functions, our result allows us to analyze the robustness of equilibrium predictions in complete information games under weaker robustness concepts, such as that given by the uniform topology. Specifically, suppose that instead of assuming common knowledge of payoffs, we assume that the exact payoffs are mutually known up to an arbitrary finite order and that it is common knowledge that the payoffs are in an arbitrarily small neighborhood of the actual payoffs. Then, our result states that for an equilibrium prediction to remain valid under these slightly weaker assumptions, it must be true for all locally interim correlated rationalizable strategies.

In some important games this leads to disturbing conclusions. As an example, we consider a Cournot oligopoly with linear cost function and sufficiently many firms. We can show that in such a game any production level that is less than or equal to the monopoly production is locally rationalizable. Suppose we weaken the complete information assumption just slightly, by assuming instead that the payoffs are mutually known up to an arbitrarily high finite order and it is common knowledge that the payoffs are within an arbitrarily small neighborhood of the original payoffs. Then, our theorem tells us that even a fixed equilibrium now yields no sharper prediction than the trivial one given by mere individual rationality, that firms' productions do not exceed the monopoly level.

In the next section, we lay out our model. In Section 3, we introduce our notion of sensitivity of equilibrium strategies and present our general result. In Section 4, we study the robustness of equilibrium predictions in complete information games to the mild perturbations as in the Cournot example above. In Section 5, we present our application on Cournot oligopoly. Finally, in Section 6, we discuss the literature and the role of our modeling assumptions in more detail. In particular, we present two extensions, one to multidimensional action spaces and one to infinite type spaces. Some of the proofs are relegated to Appendix A.

2. Basic definitions

We consider n -player *nice games* with a possibly unknown payoff-relevant parameter $\theta \in \Theta^*$, where Θ^* is a compact metric space, and with a finite set $N = \{1, 2, \dots, n\}$ of players. In a nice game, the action space of each player i is $A_i = [0, 1]$; the space of action profiles is $A = [0, 1]^n$,² and the utility function $u_i : \Theta^* \times A \rightarrow \mathbb{R}$ of player i is continuous in the action profile $a = (a_i, a_{-i}) \in A$ and strictly concave³ in own action $a_i \in A_i$. We fix the players, action space and utility function and 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. Formally, by a (finite) *model*, we mean a finite set $\Theta \times T_1 \times \dots \times T_n$ associated with beliefs $\kappa_{t_i} \in \Delta(\Theta \times T_{-i})$ for each $t_i \in T_i$, where $\Theta \subseteq \Theta^*$.⁴ By associating each type t_i with a belief κ_{t_i} , we encode all the relevant information about the model into types. We designate t_i as a generic type and $\Theta \times T$ as a model that contains t_i .

Remark 1. In our formulation, in any model it is common knowledge that the payoff functions are in $\{u(\theta, \cdot) \mid \theta \in \Theta^*\}$. Since Θ^* is arbitrary, this allows arbitrary common knowledge restrictions on payoff functions.

¹ Local rationalizability starts with a subset of actions for each player and applies the best-response function iteratively, instead of starting with the set of all possible actions. The word local here does not refer to local best replies or nearby types.

² **Notation:** Given any list $X_1, \dots, X_n$ of sets, we write $X = X_1 \times \dots \times X_n$ with typical element x , $X_{-i} = \prod_{j \neq i} X_j$ with typical element x_{-i} , and $(x_i, x_{-i}) = (x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n)$. Likewise, for any family of functions $f_j : X_j \rightarrow Y_j$, we define $f_{-i} : X_{-i} \rightarrow Y_{-i}$ by $f_{-i}(x_{-i}) = (f_j(x_j))_{j \neq i}$. Given any metric space (X, d) , we write $\Delta(X)$ for the space of probability distributions on X , endowed with Borel σ -algebra and the weak topology.

³ We use the strict concavity assumption to make sure that for any belief over the other players' actions, a player's utility function is always single-peaked in his own action. (Mere single-peakedness is not strong enough, because it is not preserved in the presence of uncertainty.)

⁴ It is standard to define a type space as a pair (T, κ) . For convenience, we suppress the belief map κ in our notation of type spaces. This will cause no confusion, as we will never use the same labels for types in two different spaces, so we can safely associate types themselves with beliefs rather than considering them as merely labels.

Given any type t_i in a type space $\Theta \times T$, we can compute the first-order belief $h_i^1(t_i) \in \Delta(\Theta^*)$ of t_i (about θ), second-order belief $h_i^2(t_i) \in \Delta(\Theta^* \times \Delta(\Theta^*)^n)$ of t_i (about θ and the first-order beliefs), etc., using the joint distribution of the types and θ . Note that $h_i^1(t_i)$ is computed from κ_{t_i} , $h_i^2(t_i)$ is computed from κ_{t_i} and from those $\kappa_{t_{-i}}$ with $\kappa_{t_{-i}}(t_{-i}) > 0$, and so on. Using the mapping $h_i : t_i \mapsto (h_i^1(t_i), h_i^2(t_i), \dots)$, we can embed all such models in the universal type space (Mertens and Zamir, 1985; Brandenburger and Dekel, 1993). We will be interested in the subset T^u of the universal type space consisting of hierarchies that can arise in finite models. That is, we consider $T^u = T_1^u \times \dots \times T_n^u$ where

$$T_i^u = \{h_i(t_i) \mid t_i \in T_i \text{ for some finite model } \Theta \times T\}.$$

A strategy of a player i with respect to T_i is any function $s_i : T_i \rightarrow A_i$. Given any type t_i and any profile s_{-i} of strategies, we write $\pi(\cdot | t_i, s_{-i}) \in \Delta(\Theta \times A_{-i})$ for the joint distribution of the underlying uncertainty and the other players' actions induced by t_i and s_{-i} . We define $\pi(\cdot | t_i, s_{-i})$ similarly for functions $s_{-i} : \Theta \times T_{-i} \rightarrow A_{-i}$. For each $i \in N$ and for each belief $\pi \in \Delta(\Theta \times A_{-i})$, we write $BR_i(\pi)$ for the unique action $a_i \in A_i$ that maximizes the expected value of $u_i(\theta, a_i, a_{-i})$ under the probability distribution π . A strategy profile $s^* = (s_1^*, s_2^*, \dots)$ is a *Bayesian Nash equilibrium* if and only if at each t_i ,

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

Note that under our assumptions, there exists a Bayesian Nash equilibrium s^* on T^u (see Yildiz, 2009).

We will consider singleton selections from Bayesian Nash equilibria of models, picking a Bayesian Nash equilibrium for each model such that when we put all these models together, the resulting strategy profile is a Bayesian Nash equilibrium of the larger game. This can be thought of as a consistency requirement of a theory of selection for various games. More precisely, we will fix a Bayesian Nash equilibrium $s^* : T^u \rightarrow A$ in T^u and pick the Bayesian Nash equilibrium $s|_T^*$ with

$$s|_T^*(t) = s^*(h(t)) \quad (\forall t \in T)$$

as the solution in type space T . (Notice that $s|_T^*$ is a Bayesian Nash equilibrium of T .) Multiple equilibria are introduced to our analysis trivially, by considering sets of equilibria s^* on T^u , which does not affect our analysis.

Our formulation also restricts the equilibrium action to depend only on the hierarchy of beliefs. That is, if there are two types t_i and t'_i in possibly two different models with identical belief hierarchies (i.e. $h_i(t_i) = h_i(t'_i)$), then the equilibrium actions are the same for t_i and t'_i . In particular, in a model with redundant types, all types with identical belief hierarchies play the same action, ruling out the extra equilibria introduced by the redundant types. This is the only restriction imposed on the solution of individual models: Yildiz (2009) shows that given any family of models $\Theta^\alpha \times T^\alpha$ with equilibria $s^\alpha : T^\alpha \rightarrow A$ such that the types with identical belief hierarchies play the same action, there exists an equilibrium $s^* : T^u \rightarrow A$ such that $s^*(t) = s^\alpha(t)$ for all $t \in T^\alpha$ and for all α .

The next result establishes that $s^*(T^u)$ is a product of convex intervals:

Lemma 1. For any equilibrium $s^* : T^u \rightarrow A$ and any $i \in N$, $s_i^*(T_i^u)$ is convex.

As we detail in Appendix A, this very useful result follows from the facts that the (single-valued) best-response function is continuous with respect to beliefs and that T^u is a convex set when types are represented by their beliefs. This is despite the fact that s^* is highly discontinuous and T^u is a large, complicated type space. The result also applies to the standard universal type space.

2.1. Local interim rationalizability

We will show that the sensitivity of equilibrium strategies is characterized by a local version of (interim correlated) rationalizability (for the original version, see Dekel et al. (2007), Battigalli (2003), Battigalli and Siniscalchi (2003)). Interim correlated rationalizability allows correlations not only within players' strategies but also between their strategies and θ . For any set $B = B_1 \times \dots \times B_n \subset A$ and any i and t_i , we set

$$S_i^0[B, t_i] = B_i$$

and define sets $S_i^k[B, t_i]$ for $k > 0$ iteratively by

$$S_i^k[B, t_i] = \{BR_i(\text{marg}_{\Theta \times A_{-i}} \pi) \mid \pi \in \Delta(\Theta \times T_{-i} \times A_{-i}), \text{marg}_{\Theta \times T_{-i}} \pi = \kappa_{t_i}, \pi(a_{-i} \in S_{-i}^{k-1}[B, t_{-i}]) = 1\}.$$

The set $S_i^k[B, t_i]$ consists of best replies to beliefs that assign positive probability only to the actions that are in $S_{-i}^{k-1}[B, \cdot]$. As in Dekel et al. (2007), $S_i^k[B, t_i]$ only depends on $h_i^k(t_i)$, not the particular type space or any higher-order beliefs. That is,

$$S_i^k[B, t_i] = S_i^k[B, \tilde{t}_i] \quad \text{whenever } h_i^m(t_i) = h_i^m(\tilde{t}_i) \text{ for all } m \leq k. \quad (2.1)$$

We define the limit set, which we call the set of *locally interim correlated rationalizable actions with respect to B*, by

$$S_i^\infty[B, t_i] = \bigcap_{k=0}^{\infty} \bigcup_{m=k}^{\infty} S_i^m[B, t_i]. \quad (2.2)$$

Note that unlike in the typical elimination process, the sequence of sets $S_i^k[B, t_i]$ may become larger or be incomparable to one another. Therefore, to make sure the rationalizable set $S_i^\infty[B, t_i]$ is well-defined, we have used in (2.2) the limsup, or the set of actions contained in $S_i^k[B, t_i]$ for infinitely many k . In the following section we will see that in the case of primary interest, when B is the range of an equilibrium, the sets indeed decrease and we have an ordinary elimination process. Nevertheless, the general definition is useful: since $S_i^k[B, t_i]$ and hence $S_i^\infty[B, t_i]$ are monotone with respect to set inclusion in the argument B , when we know $B \subset s^*(T^u)$ we can use $S_i^\infty[B, t_i]$ as a lower bound in Proposition 1. This is the motivation for using the limsup in (2.2) rather than the liminf or intersection; it gives the largest lower bound in the sequel. Finally, note that when $B = A$, we have simply the usual elimination process for interim correlated rationalizability, and can write $S_i^k[t_i]$ for $S_i^k[A, t_i]$.

The following result is an extension of earlier results by Moulin (1984) and Battigalli (2003) to local interim correlated rationalizability.

Lemma 2. For any convex $B = B_1 \times \cdots \times B_n$, any i, t_i, k , and any $a_i \in S_i^k[B, t_i]$,

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

for some $\hat{s}_{-i} : \Theta \times T_{-i} \rightarrow A_{-i}$ with $\hat{s}_{-i}(\theta, t_{-i}) \in S_{-i}^{k-1}[B, t_{-i}]$ for all (θ, t_{-i}) .

That is, in a nice game, every rationalizable action is a best reply to a *deterministic* theory about how the other players' actions are related to their types *and* the underlying parameter. Here, the action of another player j may vary with θ or a third player's type because we allow all possible correlations.

3. Sensitivity to higher-order beliefs

In this section, we will introduce a straightforward measure of sensitivity of a strategy to higher-order beliefs and present our general result, which gives a characterization of sensitivity in terms of local interim correlated rationalizability.

Fix any strategy $s_i^* : T_i^u \rightarrow A_i$ on T_i^u and any type t_i of a player i . According to strategy s_i^* , type t_i will play $s_i^*(h_i(t_i))$. Now imagine a researcher who only knows the first k orders of beliefs of player i and knows that he plays s_i^* . All the researcher can conclude from this information is that i plays one of the actions in⁵

$$A_i^k[s_i^*, t_i] \equiv \{s_i^*(h_i(\tilde{t}_i)) \mid h_i(\tilde{t}_i) \in T_i^u, h_i^m(\tilde{t}_i) = h_i^m(t_i) \forall m \leq k\}.$$

That is, an action is in $A_i^k[s_i^*, t_i]$ if and only if it is played according to s_i^* by a type $\tilde{t}_i$ that comes from a finite model and whose first k order beliefs are as in t_i . Therefore, $A_i^k[s_i^*, t_i]$ measures precisely how sensitive the strategy s_i^* is to the specification of beliefs at orders higher than k when the first k orders of beliefs are as specified by t_i . 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

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

If the researcher knew only that the strategy of i is in a given set S_i , rather than knowing what his strategy is, then he could conclude from his information only that i will play an action in

$$A_i^k[S_i, t_i] = \bigcup_{s_i^* \in S_i} A_i^k[s_i^*, t_i].$$

The main result of this paper characterizes the sets $A_i^k[s_i^*, t_i]$ by local rationalizability:

Proposition 1. For any equilibrium s^* and any (i, k, t_i) ,

$$A_i^k[s_i^*, t_i] = S_i^k[s^*(T^u), t_i].$$

In particular, when $s^*(T^u) = A$, $A_i^k[s_i^*, t_i] = S_i^k[t_i]$. Also, for any $B \subseteq s^*(T^u)$, $S_i^\infty[B, \hat{t}_i] \subseteq \bigcup_{m \geq k} S_i^m[B, \hat{t}_i] \subseteq A_i^k[s_i^*, t_i]$.

⁵ Note that it would suffice to require $h_i^k(\tilde{t}_i) = h_i^k(t_i)$, which by coherence entails agreement in all lower-order beliefs. We find it more intuitive to refer to agreement at orders 1 through k .

Proposition 1 tells us a way of determining how sensitive an arbitrary equilibrium s^* is to the specifications of beliefs at orders higher than k : Consider the set of *all* actions that are played by some type according to s^* , without requiring any connection to the beliefs at hand. Apply the best response operator to this set k times, allowing all possible correlations. The resulting set is precisely the set of actions that could be played by types whose first k orders of beliefs are as specified at the beginning. When the parameter space is rich enough so that all actions are played by some types, this set is simply the set of actions that survive k th-order elimination of strictly dominated actions in the interim stage. When we allow k to be arbitrarily high, this set is simply the set of all (locally) interim correlated rationalizable actions. It is immediate from their definition that the sets $A_i^k[s_i^*, t_i]$ are decreasing in k , i.e. $A_i^k[s_i^*, t_i] \supseteq A_i^{k+1}[s_i^*, t_i]$. The equality in the proposition then implies that the sets $S_i^k[s^*(T^u), t_i]$ are also decreasing in k . That is, when we start with the range $s^*(T^u)$ of an equilibrium, the above process is, in fact, an elimination process:

Corollary 1. For any equilibrium s^* and any (i, k, t_i) , $S_i^k[s^*(T^u), t_i] \supseteq S_i^{k+1}[s^*(T^u), t_i]$.

Sometimes, it may be difficult to know the set of actions played by arbitrary types according to s^* , but we may still know the behavior of certain types, e.g., the common knowledge types. In that case, we can still use Proposition 1 to find a lower bound: consider the set of actions that are known to be played by some type and apply the best response correspondence k times. In that case, $S_i^k[B, t_i]$ may not be decreasing in k , and one can find a better lower bound by iterating the procedure further. Since $S_i^m[B, t_i] \subseteq A_i^m[s_i^*, t_i] \subseteq A_i^k[s_i^*, t_i]$ for each $m \geq k$, we have $\bigcup_{m \geq k} S_i^m[B, t_i] \subseteq A_i^k[s_i^*, t_i]$.

A comparison of this result with that of WY is useful. In WY, we consider a finite action game and assume that the parameter space is so rich that every action becomes dominant at some parameter value. Then, we show that for each $a_i \in S_i^k[t_i]$ and each rationalizable strategy s_i , we can perturb first k order beliefs arbitrarily slightly and change the higher-order beliefs to obtain a type $\tilde{t}_i$ such that $s_i(\tilde{t}_i) = a_i$. Here, we consider nice games instead of finite-action games. At the expense of focusing on Bayesian Nash equilibria, rather than arbitrary rationalizable strategies, we strengthen the result in two ways. First, we do not make any richness assumption, allowing arbitrary common knowledge restriction on payoffs. Instead, we give a general characterization, $A_i^k[s_i^*, t_i] = S_i^k[s^*(T^u), t_i]$, that depends on the range of equilibrium on T^u . Second, since the best reply is always unique, we do not need to perturb the lower-order beliefs at all, and hence our result does not refer directly to any topology on beliefs.

We will now give the proof for $k = 1$. Our general proof, which is in Appendix A, uses the same arguments inductively. The inclusion $A_i^1[s_i^*, t_i] \subseteq S_i^1[s^*(T^u), t_i]$ follows from the definitions and (2.1). Indeed, for any $a_i \in A_i^1[s_i^*, t_i]$, we have $a_i = s_i^*(\tilde{t}_i)$ for some $\tilde{t}_i$ with $h_i^1(\tilde{t}_i) = h_i^1(t_i)$, implying also $a_i = BR_i(\pi(\cdot | \tilde{t}_i, s_{-i}^*))$. Then, $a_i \in S_i^1[s^*(T^u), \tilde{t}_i] = S_i^1[s^*(T^u), t_i]$, where the last equality is by (2.1).

To show the inclusion $S_i^1[s^*(T^u), t_i] \subseteq A_i^1[s_i^*, t_i]$, take any type t_i , from a finite type space $\Theta \times T$, and any $a_i \in S_i^1[s^*(T^u), t_i]$. We need to construct a new type $\tilde{t}_i$, from a finite type space $\tilde{\Theta} \times \tilde{T}$, such that

- (1) $h_i^1(\tilde{t}_i) = h_i^1(t_i)$, i.e., $\sum_{\tilde{t}_{-i} \in \tilde{T}_{-i}} \kappa_{\tilde{t}_i}(\theta, \tilde{t}_{-i}) = \sum_{t_{-i} \in T_{-i}} \kappa_{t_i}(\theta, t_{-i})$ for each $\theta \in \Theta$, and
- (2) $S_i^1(\tilde{t}_i) = a_i$, i.e., $a_i = BR_i(\pi(\cdot | \tilde{t}_i, s_{-i}^*))$.

By Lemma 1, $s^*(T^u)$ is a product of convex sets. Hence, by Lemma 2, $a_i = BR_i(\pi(\cdot | t_i, s_{-i}))$ for some function $s_{-i} : \Theta \times T_{-i} \rightarrow A_{-i}$ with $s_{-i}(\theta, t_{-i}) \in s_{-i}^*(T_{-i}^u)$. By definition, for each $a_{-i} \in s_{-i}(\Theta \times T_{-i})$, which is contained in $s_{-i}^*(T_{-i}^u)$, there exists $t_{-i}[a_{-i}] \in T_{-i}^u$ such that $s_{-i}^*(t_{-i}[a_{-i}]) = a_{-i}$. We will define our type $\tilde{t}_i$ by beliefs

$$\kappa_{\tilde{t}_i}(\theta, t_{-i}[a_{-i}]) = \pi(\theta, a_{-i} | t_i, s_{-i}) \quad (\forall \theta \in \Theta, a_{-i} \in s_{-i}(\Theta \times T_{-i})). \quad (3.1)$$

That is, we assign the probability of an action under $\pi(\cdot | t_i, s_{-i})$ to a type who plays that action in equilibrium, while we keep the probabilities of θ intact. It is then straightforward to check that the two conditions above are satisfied. First,

$$\sum_{t_{-i}[a_{-i}]} \kappa_{\tilde{t}_i}(\theta, t_{-i}[a_{-i}]) = \sum_{a_{-i}} \pi(\theta, a_{-i} | t_i, s_{-i}) = \sum_{t_{-i} \in T_{-i}} \kappa_{t_i}(\theta, t_{-i}).$$

Here, the first equality is by (3.1). To see the second equality, note that both expressions are equal to the probability type t_i assigns on θ . The second condition is satisfied (i.e. $a_i = BR_i(\pi(\cdot | \tilde{t}_i, s_{-i}^*))$) because, for each (θ, a_{-i}) ,

$$\pi(\theta, a_{-i} | \tilde{t}_i, s_{-i}^*) = \sum_{s_{-i}^*(t_{-i}) = a_{-i}} \kappa_{\tilde{t}_i}(\theta, t_{-i}) = \kappa_{\tilde{t}_i}(\theta, t_{-i}[a_{-i}]) = \pi(\theta, a_{-i} | t_i, s_{-i})$$

and $a_i = BR_i(\pi(\cdot | t_i, s_{-i}))$. (It is also clear that $\tilde{t}_i$ comes from a finite space.⁶)

⁶ For each a_{-i} , we have $t_{-i}[a_{-i}] \in T_{-i}^{a_{-i}}$ for some finite model $\Theta^{a_{-i}} \times T^{a_{-i}}$. We define the finite model $\tilde{\Theta} \times \tilde{T}$ by $\tilde{\Theta} = (\bigcup_{a_{-i}} \Theta^{a_{-i}}) \cup \Theta$, $\tilde{T}_i = (\bigcup_{a_{-i}} T_i^{a_{-i}}) \cup \{t_i\}$, and $\tilde{T}_j = \bigcup_{a_{-i}} T_j^{a_{-i}}$ for all other players j . The belief of $\tilde{t}_i$ is defined as above, and the beliefs of all the other types are kept as in the original model each comes from.

It is crucial for the proof and the result that s^* is an equilibrium. Since s^* is an equilibrium, each type plays a best response to the same strategy profile s_{-i}^* of the other players. This puts a strong restriction on the actions played by different types. Indeed, we conclude that $\tilde{t}_i$ plays a_i by assuming that both t_i and $\tilde{t}_i$ play a best reply to s_{-i}^* . On the other hand, in a rationalizable strategy, each type's action may be a best response to different rationalizable strategies of the other players. In that case, we could not make such assumptions, and the result need not be true.

Lemma 2 also plays a crucial role in the proof. In order for the belief $\kappa_{\tilde{t}_i}$ to be well-defined, the mapping $a_{-i} \mapsto t_{-i}[a_{-i}]$ needs to be measurable on the set of actions a_{-i} type t_i assigns positive probability when he plays a_i as a best reply. Lemma 2 allows us to focus on the “degenerate belief” for which that set, which is contained in $s_{-i}(\Theta \times T_{-i})$, is finite, and therefore the mapping $a_{-i} \mapsto t_{-i}[a_{-i}]$ is trivially measurable on that domain. Since we have uncountably many actions, without Lemma 2, we would need to consider beliefs that put positive probability on an uncountable set of actions. On such a domain, the mapping $a_{-i} \mapsto t_{-i}[a_{-i}]$ need not be measurable.

4. Minimally robust predictions of equilibrium

In this section, we study the minimally robust equilibrium predictions of complete-information models. For any given complete-information model, instead of assuming that payoffs are common knowledge, we assume that the payoffs are mutually known up to a finite order k , for arbitrarily high k , and that it is common knowledge that the payoffs are within an arbitrarily small neighborhood. We show that if a prediction is robust to such a mild relaxation, then it must hold for all locally rationalizable outcomes near the solution.

Take $\Theta^* \subset \mathbb{R}$ and fix any $\bar{\theta} \in \text{int}(\Theta^*)$,⁷ so that all values in $\Theta_\varepsilon^* \equiv [\bar{\theta} - \varepsilon, \bar{\theta} + \varepsilon]$ are possible for some $\bar{\varepsilon} > 0$. Consider the complete information model $T^{u,0} \equiv \{t^{CK}(\bar{\theta})\}$ in which it is common knowledge that $\theta = \bar{\theta}$. Towards relaxing this common knowledge restriction slightly, for each $\varepsilon \in (0, \bar{\varepsilon}]$, take $\Theta_\varepsilon^* \equiv [\bar{\theta} - \varepsilon, \bar{\theta} + \varepsilon]$ as the parameter space. Consider the space $T^{u,\varepsilon}$ of all belief hierarchies of types that come from finite type spaces in which it is common knowledge that $\theta \in \Theta_\varepsilon^*$. Fix any Bayesian Nash equilibrium $s^* : T^{u,\bar{\varepsilon}} \rightarrow A$. Note that since $T^{u,\bar{\varepsilon}}$ contains all type spaces $T^{u,\varepsilon}$ with $\varepsilon \in [0, \bar{\varepsilon}]$ as belief closed subspaces, the restriction of s^* to each of these spaces induces a Bayesian Nash equilibrium on that space. In the complete information model $T^{u,0}$, s^* yields $s^*(t^{CK}(\bar{\theta}))$ as the unique outcome. If one relaxes the common knowledge restriction, by instead assuming that it is common knowledge that θ is in the ε -neighborhood Θ_ε^* of $\bar{\theta}$ and that players have mutual certainty of $\theta = \bar{\theta}$ up to order k ,⁸ then the set of possible outcomes according to s^* is

$$A^{k,\varepsilon}[s^*, \bar{\theta}] \equiv \{s^*(h(\tilde{t})) \mid h(\tilde{t}) \in T^{u,\varepsilon}, h^m(\tilde{t}) = h^m(t^{CK}(\bar{\theta})) \forall m \leq k\}.$$

Definition 1. The *minimally-robust prediction* of s^* at $\bar{\theta}$ is

$$A^{\infty,0}[s^*, \bar{\theta}] \equiv \bigcap_{k < \infty, \varepsilon \in (0, \bar{\varepsilon})} A^{k,\varepsilon}[s^*, \bar{\theta}].$$

Note that if $a \in A^{\infty,0}[s^*, \bar{\theta}]$, we cannot rule out that a is the outcome of s^* even if we know that the players have k th-order mutual certainty of $\theta = \bar{\theta}$ and it is common knowledge that θ is in ε -neighborhood of $\bar{\theta}$, regardless of the sizes of k and ε . Conversely, if $a \notin A^{\infty,0}[s^*, \bar{\theta}]$, then there exist k and $\varepsilon > 0$ such that we can rule out a as the outcome of s^* whenever we know that the players have k th-order mutual certainty of $\theta = \bar{\theta}$ and it is common knowledge that θ is in ε -neighborhood of $\bar{\theta}$. Note also that we consider only the predictions of the form “the outcome is in set B ”. Robustness of other predictions can be deduced from the robustness of such basic predictions (see WY). In general, for a prediction to be minimally robust, it must hold for all outcomes in $A^{\infty,0}[s^*, \bar{\theta}]$.

Towards finding a lower bound for the minimally-robust prediction, for any $a \in A$ and any $\varepsilon > 0$, we write $B(a, \varepsilon) = ([a_1 - \varepsilon, a_1 + \varepsilon] \times \cdots \times [a_n - \varepsilon, a_n + \varepsilon]) \cap A$ for the ε -neighborhood of a .

Definition 2. For any $a \in A$ and $\bar{\theta}$, the *locally rationalizable set* at $(a, \bar{\theta})$ is

$$S^\infty[a, \bar{\theta}] \equiv \bigcap_{\varepsilon > 0} S^\infty[B(a, \varepsilon), t^{CK}(\bar{\theta})].$$

Note that in order to compute the locally rationalizable set, one does not need to consider the payoff or information perturbations. In the complete information model, one slightly perturbs the equilibrium outcome and applies best response operator iteratively. The difference between various notions of rationalizability disappears because there is no payoff uncertainty. Finally, we assume that θ is payoff-relevant around the equilibrium value of the complete information game:

⁷ $\mathbb{R}$ denotes the set of real numbers, and for any $X \subset \mathbb{R}$, $\text{int}(X)$ denotes the interior of X .

⁸ Here, we use the term *certainty* instead of *knowledge* to emphasize that the truth axiom is not assumed, i.e., a player may be certain of something that happens to be false. For example, the *common knowledge* assumption that $\theta = \bar{\theta}$ in model $T^{u,0} = \{t^{CK}(\bar{\theta})\}$ is weakened to *common certainty* when we embed $t^{CK}(\bar{\theta})$ into $T^{u,\varepsilon}$ with $\varepsilon > 0$, as $T^{u,\varepsilon}$ contains a type t_i that assigns positive probability on $(\theta, t_i^{CK}(\bar{\theta}))$ for some $\theta \neq \bar{\theta}$.

Assumption 1. There exists a neighborhood of $\bar{\theta}$ on which, for each $i \in N$, $f_i(\theta) \equiv \arg \max_{a_i} u_i(\theta, a_i, s_{-i}^*(t_{-i}^{CK}(\bar{\theta})))$ is continuous and either strictly increasing or strictly decreasing in θ .

Assumption 1 ensures that the minimally robust prediction of s^* cannot be sharper than local rationalizability:

Proposition 2. Under Assumption 1, the minimally robust prediction of equilibrium s^* at $\bar{\theta}$ contains the locally interim rationalizable set at $(\bar{\theta}, s^*(t_{-i}^{CK}(\bar{\theta})))$:

$$A^{\infty,0}[s^*, \bar{\theta}] \supseteq S^\infty[s^*(t_{-i}^{CK}(\bar{\theta})), \bar{\theta}].$$

Proof. By Assumption 1, for any $\varepsilon > 0$, there exists $\varepsilon' > 0$ such that for any $a \in B(s^*(t_{-i}^{CK}(\bar{\theta})), \varepsilon')$ and any $i \in N$, there exists $\theta' \in \Theta_\varepsilon^*$ such that a_i is the best reply to $s_{-i}^*(t_{-i}^{CK}(\bar{\theta}))$ under θ' . Then, the type $t_i' \in T_i^{u,\varepsilon}$ that puts probability 1 on $(\theta', t_{-i}^{CK}(\bar{\theta}))$ plays a_i according to s^* . Hence,

$$B(s^*(t_{-i}^{CK}(\bar{\theta})), \varepsilon') \subseteq s^*(T^{u,\varepsilon}).$$

Thus, by Proposition 1, $S^\infty[B(s^*(t_{-i}^{CK}(\bar{\theta})), \varepsilon'), t_{-i}^{CK}(\bar{\theta})] \subseteq A^{\infty,\varepsilon}[s^*, \bar{\theta}]$, showing that

$$S^\infty[s^*(t_{-i}^{CK}(\bar{\theta})), t_{-i}^{CK}(\bar{\theta})] \subseteq A^{\infty,\varepsilon}[s^*, \bar{\theta}].$$

Since ε is arbitrary, it follows that

$$S^\infty[s^*(t_{-i}^{CK}(\bar{\theta})), t_{-i}^{CK}(\bar{\theta})] \subseteq \bigcap_{\varepsilon > 0} A^{\infty,\varepsilon}[s^*, \bar{\theta}] = A^{\infty,0}[s^*, \bar{\theta}]. \quad \square$$

Proposition 2 establishes that in order for a prediction to be robust, it must hold for all locally rationalizable actions near equilibrium. In other words, even under the strong common knowledge restrictions, one cannot make any sharper prediction than local rationalizability around equilibrium. The converse is not necessarily true, as the next example illustrates.

Example 1. Take $N = \{1, 2\}$, $u_i(\theta, a_1, a_2) = \theta a_1 a_2 - a_i^3/3$, $\Theta^* = [0, 2]$, and $\bar{\theta} = 1$. For every $k > 0$ and every $\varepsilon \in (0, 1)$, consider a finite type space $\{\bar{\theta}, \bar{\theta} + \varepsilon\} \times T^{k,\varepsilon}$ in which $\theta = \bar{\theta}$ can be mutually known up to order k and not for any higher order (e.g. consider a version of the e-mail game in which the players send messages until it becomes k th-order mutual knowledge that $\theta = \bar{\theta}$). Note that $(1, 1)$ is an equilibrium of $\{t_{-i}^{CK}(\bar{\theta})\}$, and each $\{\bar{\theta}, \bar{\theta} + \varepsilon\} \times T^{k,\varepsilon}$ has an equilibrium in which each type plays 0. Since $t_{-i}^{CK}(\bar{\theta})$ is not contained in any $T^{k,\varepsilon}$, by the existence result in Yildiz (2009), there then exists an equilibrium $s^*: T^u \rightarrow A$ such that $s^*(t_{-i}^{CK}(\bar{\theta})) = (1, 1)$ and $s^*(t) = (0, 0)$ for all $t \in T^{k,\varepsilon}$ and for all (k, ε) . One can easily check that the complete-information game is locally dominance solvable around the equilibrium: $S^\infty[s^*(t_{-i}^{CK}(\bar{\theta})), t_{-i}^{CK}(\bar{\theta})] = \{(1, 1)\}$. But, by construction, $(0, 0) \in A^{\infty,0}[s^*, \bar{\theta}]$. Hence,

$$A^{\infty,0}[s^*, \bar{\theta}] \not\subseteq S^\infty[s^*(t_{-i}^{CK}(\bar{\theta})), t_{-i}^{CK}(\bar{\theta})].$$

Therefore, a prediction may not be minimally robust even if it holds for all locally rationalizable actions. In this example, s^* is discontinuous at $t_{-i}^{CK}(\bar{\theta})$, while there exists a continuous equilibrium s^{**} on $T^{u,\varepsilon}$, in which all types play 0. It is tempting to seek for a characterization by assuming continuity at $t_{-i}^{CK}(\bar{\theta})$. Unfortunately, equilibrium must be discontinuous when there are multiple locally rationalizable actions (by Proposition 2) or when there are multiple rationalizable actions and a rich set of parameters (by the result of WY). This example also shows that Assumption 1 is not superfluous in Proposition 2. To see this, take equilibrium s^{**} . Assumption 1 does not hold, as $f_i(\theta) = \arg \max_{a_i} u_i(\theta, a_i, s_{-i}^{**}(t_{-i}^{CK}(\bar{\theta}))) = 0$ at each θ . In contrast to Proposition 2, we have $A_i^{\infty,0}[s_i^{**}, \bar{\theta}] = \{0\}$ while $S_i^\infty[s_i^{**}(t_{-i}^{CK}(\bar{\theta})), t_{-i}^{CK}(\bar{\theta})] = [0, 1)$.

A number of papers, such as Moulin (1984) and Guesnerie (2002), have analyzed the relationship between local dominance-solvability and Cournot stability of equilibrium, i.e. stability under best-reply dynamics. Assuming common knowledge of payoffs, they show that local dominance solvability is sufficient (but not always necessary) for stability under small perturbations to equilibrium actions. Here, we assume that players do not deviate from the equilibrium strategies, but instead examine a very different kind of robustness, to small perturbations of the interim beliefs. In contrast to the prior results, we show that local dominance solvability is a necessary but not a sufficient condition for our notion of robustness.

5. Application: Cournot oligopoly

In this section, we will show that in Cournot oligopoly with sufficiently many firms, the minimally robust prediction of Bayesian Nash equilibrium is that no firm produced more than its monopoly outcome, a trivial implication of profit maximization. That is, 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.

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). We will show that this result extends to local interim correlated rationalizability when the equilibrium of the complete information game is in the interior of A . Then, 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, even if he assumes that it is common knowledge that the payoffs are in an arbitrarily small neighborhood of the true value and subscribes to strong refinements of equilibrium that yield unique solutions. 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.

On the other hand, Börgers 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 1, equilibrium outcomes will not be sensitive to higher-order beliefs.

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 $\bar{\theta} \in \text{int}(\Theta^*)$. We also assume that $P(0; \bar{\theta}) > 0$, $P(\cdot; \bar{\theta})$ is strictly decreasing when it is positive, and $\lim_{Q \rightarrow \infty} P(Q; \bar{\theta}) = 0$. Therefore, there exists a unique $\bar{Q}$ such that $P(\bar{Q}; \bar{\theta}) = c$. We assume that, on $[0, \bar{Q}]$, $P(\cdot; \bar{\theta})$ is continuously twice-differentiable and $P' + Q P'' < 0$.

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 $[0, \bar{Q}]$ 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 and Suzumura, 1971). We further impose Assumption 1: $q^*(Q_{-i}; \theta)$ is a continuous and strictly increasing function of θ at $(Q_{-i}, \bar{\theta})$ where $Q_{-i} = (n-1)s^*(t^{CK}(\bar{\theta}))$.

Lemma 3. *In the Cournot oligopoly above, there exists $\bar{n} < \infty$ such that for any $n > \bar{n}$ and any $\varepsilon > 0$,*

$$S^\infty[B(s^*(t^{CK}(\bar{\theta})), \varepsilon), t^{CK}(\bar{\theta})] = [0, q^M]^n,$$

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

This is a straightforward extension of a result by Basu (1992) for rationalizability to local rationalizability. The proof is in Appendix A. Together with Proposition 2, this lemma yields the following.

Proposition 3. *In the Cournot oligopoly above, there exists $\bar{n} < \infty$ such that for any $n > \bar{n}$ and any Bayesian Nash equilibrium s^* , the minimally robust prediction of s^* is that no firm produces more than its monopoly output q^M under $P(\cdot; \bar{\theta})$:*

$$A^{\infty,0}[s^*, \bar{\theta}] = [0, q^M]^n.$$

Proof. Since we can put a large upper bound on q , by (i) above, we have a nice game. Moreover, by Lemma 3, $S^\infty[B(a, \varepsilon), t^{CK}(\bar{\theta})] = [0, q^M]^n$ for all $\varepsilon > 0$, showing that $S^\infty[s^*(t^{CK}(\bar{\theta})), \bar{\theta}] = [0, q^M]^n$. Hence, by Proposition 2, $A^{\infty,0}[s^*, \bar{\theta}] \supseteq S^\infty[s^*(t^{CK}(\bar{\theta})), \bar{\theta}] = [0, q^M]^n$. Conversely, by Proposition 1 and the definition of $A^{\infty,0}$, $A^{\infty,0}[s^*, \bar{\theta}] \subseteq S^1[t^{CK}(\bar{\theta})] \subseteq [0, q^M]^n$. $\square$

Proposition 3 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 $|\theta - \bar{\theta}| \leq \varepsilon$ and agrees with the k th-order mutual certainty 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 . Proposition 3 states that the skeptic nonetheless cannot rule out any output level that is implied by simple profit maximization.

6. Extensions and concluding remarks

We have so far made several assumptions, such as unidimensional action spaces and finite type spaces. In this section we will describe how our results can be extended beyond these assumptions and comment on our modeling choices and the literature.

6.1. Multidimensional action spaces

We have so far confined ourselves to nice games, in which action spaces are unidimensional intervals and the utility functions are continuous and strictly concave in own action. These assumptions are made in order to ensure two properties:

- (1) the rationalizable actions are best replies to degenerate beliefs, as in point rationalizability of Bernheim (1984), and
- (2) there is always a unique best reply.

Our characterization is valid whenever these two properties hold. All of the properties of nice games are needed to imply the first property. The second property, though, holds even with multidimensional action spaces. More broadly, with unique best replies, the sensitivity of equilibrium strategies is characterized by local point rationalizability. In order to state this formally, we define *local interim correlated point rationalizability*, denoted by $\dot{S}^\infty$, for incomplete information games as follows. For any set $B = B_1 \times \cdots \times B_n \subset A$ and any i and t_i , we set

$$\dot{S}_i^0[B, t_i] = B_i$$

and define sets $\dot{S}_i^k[B, t_i]$ for $k > 0$ iteratively by

$$\dot{S}_i^k[B, t_i] = \{BR_i(\pi(\cdot|t_i, \hat{s}_{-i})) \mid \hat{s}_{-i} : \Theta \times T_{-i} \rightarrow A_{-i}, \hat{s}_{-i}(\theta, t_{-i}) \in S_{-i}^{k-1}[B, t_{-i}] \forall (\theta, t_{-i})\}.$$

We define the set of locally interim correlated point rationalizable actions with respect to B by

$$\dot{S}_i^\infty[B, t_i] = \bigcap_{k=0}^{\infty} \bigcup_{m=k}^{\infty} \dot{S}_i^m[B, t_i].$$

The following characterization holds as long as the best reply is always unique.

Proposition 4. Assume that for each finite model $\Theta \times T$, for each $i \in N$, and for each $\hat{s}_{-i} : \Theta \times T_{-i} \rightarrow A_{-i}$, there exists a unique best response $BR_i(\pi(\cdot|t_i, \hat{s}_{-i}))$. Then, for any equilibrium $s^* : T^u \rightarrow A$ and any (i, k, t_i) ,

$$A_i^k[s_i^*, t_i] = \dot{S}_i^k[s^*(T^u), t_i].$$

That is, if one fixes the first k orders of beliefs according to t_i and varies the higher-order beliefs using types from finite type spaces, then s_i^* traces $\dot{S}_i^k[s^*(T^u), t_i]$ without going outside of it. This characterization of sensitivity of equilibrium via local interim correlated point rationalizability holds as long as there is always a unique best reply, e.g., when the action spaces are compact and convex sets and the utility functions are continuous and strictly concave in own action. In particular, it holds in generalized nice games with multidimensional action spaces.

Our proof for Proposition 1 first obtains $S_i^k[s^*(T^u), t_i] = \dot{S}_i^k[s^*(T^u), t_i]$ for all (i, k, t_i) (by Lemmas 1 and 2) and then proves the characterization for $\dot{S}_i^k[s^*(T^u), t_i]$. The latter part of the proof uses only the fact that the constructed type has a unique best reply to s_{-i}^* . Since the uniqueness of the best reply is already assumed in Proposition 4, our proof of Proposition 1 also proves Proposition 4.

6.2. Infinite type spaces

We have confined ourselves to finite type spaces, only because we can ensure existence of equilibrium in the space T^u of finite types. Our main result extends to the equilibria of larger type spaces when they exist. Indeed, in our working paper (Weinstein and Yildiz, 2004) we have proved Proposition 1 for all countable type spaces, by simply taking T^u as the belief hierarchies of types that come from countable type spaces. More importantly, Proposition 1 implies a lower bound on the sensitivity of equilibria of larger type spaces already, extending our main contribution to those spaces:

Corollary 2. For any type space $\hat{T} \supseteq T^u$, for any equilibrium $s^* : \hat{T} \rightarrow A$ on $\hat{T}$, for any $i \in N$, $t_i \in T^u$, and $k \geq 0$,

$$\hat{A}_i^k[s_i^*, t_i] \supseteq S_i^k[s^*(T^u), t_i]$$

where $\hat{A}_i^k[s_i^*, t_i] \equiv \{s_i^*(h_i(\tilde{t}_i)) \mid h_i(\tilde{t}_i) \in \hat{T}_i, h_i^m(\tilde{t}_i) = h_i^m(t_i) \forall m \leq k\}$.

Proof. Since $\hat{T} \supseteq T^u$, $\hat{A}_i^k[s_i^*, t_i] \supseteq A_i^k[s_i^*, t_i] = S_i^k[s^*(T^u), t_i]$, where the last equality is by Proposition 1 and the fact that the restriction of s^* to T^u is an equilibrium on T^u . $\square$

If one enlarges the type space by including more types, equilibrium only becomes more sensitive to higher-order beliefs because now there are more ways to vary the higher-order beliefs. Hence, our lower bound for the sensitivity of equilibrium, which is our main contribution, remains valid. Note, however, that as we enlarge the type space, the equilibrium may cease to exist. In that case, our result becomes vacuous.

None of the results mentioned here addresses the sensitivity of equilibrium strategies at the types that come from uncountable type spaces, although the corollary allows them in the space. The only reason for this is that the measurability problem in our proof cannot be avoided when uncountable types are included in our construction without further structure. Apart from this technical problem with our proof, there is no reason to suspect that the high sensitivity at finite and countable types would disappear at uncountable types.

6.3. Consistency and extension property of equilibria

In a genuine case of incomplete information we have no ex ante stage, and can only work with the players' hierarchies of beliefs. Since Harsanyi (1967), these beliefs are modeled by a type in a type space. There can be multiple types from various type spaces that model the same situation. In that case, it is natural to require that the solution is the same for all such types, for otherwise the solution would be dependent on the way we model the situation. This is the consistency restriction used in this paper.

Nevertheless, Friedenberg and Meier (2008) have recently shown that a type space T may have an equilibrium that cannot be extended to a larger type space T' , even if T does not have any "redundant" types, which have identical belief hierarchies. As we mentioned earlier, this cannot happen within T^u : any equilibrium of a finite type space with no redundant types can be extended to T^u (Yildiz, 2009). When it does happen in a larger type space, consistent equilibrium selections may induce a refinement on the solutions of individual games, in a way that may not be anticipated by the researcher. For example above, one would necessarily exclude the non-extendable equilibrium of T by fixing an equilibrium s^* on T' and considering only the solutions induced by s^* on subspaces. In order to analyze the sensitivity of the latter equilibria to higher-order beliefs, one needs to consider a more liberal restriction than consistency on the equilibrium selection.

6.4. Further literature review

In Weinstein and Yildiz (2007b), considering action spaces that come from a compact metric space, and assuming a global stability condition, we showed that higher-order beliefs have exponentially decreasing impact on every Bayesian Nash equilibrium in the universal type space. In the present context, this implies that the equilibrium strategies are not sensitive to higher-order beliefs in nice games that satisfy the global stability condition, which are dominance solvable. In this paper, we observe more broadly that the sensitivity of equilibrium strategies is bounded above by local rationalizability, a fact that immediately follows from (2.1), which has been originally observed by Dekel et al. (2007) for rationalizability. The contribution of this paper is the lower bound: every equilibrium strategy has to be so sensitive to higher-order beliefs that it traces all locally rationalizable actions. For example, while Weinstein and Yildiz (2007b) emphasize the decreasing impact of higher-order beliefs in the dominance-solvable game of Cournot duopoly, here we emphasize that, in Cournot oligopoly, the equilibrium strategies are so sensitive to higher-order beliefs that one cannot make any non-trivial minimally robust prediction.

Recently, several papers took complementary approaches to weakening the richness assumption in WY. First, note that fixing a non-trivial dynamic game tree contradicts the richness assumption of WY. Chen (2008) shows, however, that the conclusion of WY remains intact under a weaker richness assumption that is satisfied by dynamic game trees with unrestricted payoff functions. Weinstein and Yildiz (2009) extends this result further to the games that are continuous at infinity, allowing uncountable action spaces in normal form. Without imposing a richness assumption, Penta (2008a) proves that the conclusion of WY holds for the (rationalizable) actions that can be traced back to dominance regions through successive best responses. Penta (2008b) analyzes the sensitivity to higher-order beliefs under common-knowledge restrictions on payoffs in a similar formulation to that of Battigalli and Siniscalchi (2003). He proves an analogous result to WY by using "interim sequential rationalizability".

Appendix A. Proofs

A.1. Proof of Lemma 1

Take any $a_i, a'_i \in s_i^*(T_i^u)$ and any $a''_i \in A_i$ with $a_i > a''_i > a'_i$. By definition, there exist finite belief-closed subspaces $\Theta \times T$ and $\Theta' \times T'$ of $\Theta^* \times T^u$ with types $t_i \in T_i$ and $t'_i \in T'_i$ such that $s_i^*(t_i) = a_i$ and $s_i^*(t'_i) = a'_i$. Now, for every $\alpha \in [0, 1]$, define

$$\beta_i(\alpha) = BR_i(\alpha\pi(\cdot|t_i, s_{-i}^*) + (1 - \alpha)\pi(\cdot|t'_i, s_{-i}^*)). \quad (A.1)$$

By the Maximum Theorem, β_i is continuous, and since s^* is an equilibrium, $\beta_i(0) = BR_i(\pi(\cdot|t'_i, s_{-i}^*)) = s_i^*(t'_i) = a'_i$ and $\beta_i(1) = a_i$. Hence, by the Intermediate-Value Theorem, there exists $\alpha^* \in (0, 1)$ such that $\beta_i(\alpha^*) = a''_i$. Now, consider the type $t_i^* \in T_i^u$ with

$$\kappa_{t_i^*} = \alpha^* \kappa_{t_i} + (1 - \alpha^*) \kappa_{t'_i}. \quad (A.2)$$

(Note that t_i^* is in the finite belief-closed subspace $\tilde{\Theta} \times \tilde{T}$ where $\tilde{\Theta} = \Theta \cup \Theta'$, $\tilde{T}_j = T_j \cup T'_j$ for any $j \neq i$, and $\tilde{T}_i = T_i \cup T'_i \cup \{t_i^*\}$.) Clearly,

$$s_i^*(t_i^*) = BR_i(\pi(\cdot|t_i^*, s_{-i}^*)) = BR_i(\alpha^* \pi(\cdot|t_i, s_{-i}^*) + (1 - \alpha^*) \pi(\cdot|t'_i, s_{-i}^*)) = \beta_i(\alpha^*) = a''_i,$$

where the first equality is due to the fact that s^* is an equilibrium, the second equality is by (A.2), the third is by definition of β_i , and the last is by definition of α^* .

A.2. Proof of Lemma 2

We first prove a preliminary technical result:

Lemma 4. Suppose $U(x_1, x_2) : [0, 1] \times C \rightarrow \mathbb{R}$ is continuous and for each fixed x_2 is strictly concave in x_1 , where $C \subset \mathbb{R}^m$ is a product of finite closed intervals. Then for any distribution π on C with $\pi(R) = 1$ for some product of (open, half-open, or closed) intervals R , there exists $x_2^* \in R$ with

$$\operatorname{argmax}_{x_1} \int_{x_2} U(x_1, x_2) d\pi = \operatorname{argmax}_{x_1} U(x_1, x_2^*).$$

Proof. Let $x_1^* = \operatorname{argmax}_{x_1} \int_{x_2} U(x_1, x_2) d\pi$. Define $g(x_2) = \operatorname{argmax}_{x_1} U(x_1, x_2)$. Strict concavity implies that g is singleton-valued. By the Maximum Theorem, its graph is closed, hence it is continuous. Assume without loss of generality that C is the closure of R . Then, since C is compact and connected, the image of C under g is a closed interval, say $[a, b]$. If $x_1^* < a$, then $U(x_1^*, x_2) < U(a, x_2)$ for every x_2 so

$$\int_{x_2} U(x_1^*, x_2) d\pi < \int_{x_2} U(a, x_2) d\pi,$$

a contradiction. By a similar argument for $x_1^* > b$, we conclude that $x_1^* \in [a, b]$, so there always exists $x_2^* \in C$ with the desired property. It remains (by symmetry) only to contradict the possibility that

$$a = x_1^* = g(x_2^*) < g(x_2) \quad (\forall x_2 \in R),$$

where x_2^* is on the boundary of R .

Assume that this is so. Since C is compact and connected, the range of U is compact and connected, so assume without loss that it is $[0, 1]$. Strict concavity tells us that U has a well-defined, finite, strictly decreasing right-derivative with respect to x_1 on $(0, 1] \times C$, which we denote U' . Note furthermore that for each $x_1 > a$ and each x_2 ,

$$U(x_1, x_2) - U(a, x_2) > (x_1 - a)U'(x_1, x_2),$$

implying

$$\int_{x_2} U(x_1, x_2) d\pi - \int_{x_2} U(a, x_2) d\pi > (x_1 - a) \int_{x_2} U'(x_1, x_2) d\pi.$$

Hence, if we can find $x_1 > a$ for which $\int_{x_2} U'(x_1, x_2) d\pi > 0$ we will have succeeded by contradicting optimality of a . Indeed, we will show that it has positive $\liminf$ as $x_1 \rightarrow a$. We will decompose this integral into positive and negative parts. Concavity tells us that $U'(x_1, x_2) > 0 \Leftrightarrow x_1 < g(x_2)$. Furthermore, there is a uniform lower bound on the derivative:

$$U'(x_1, x_2) > (U(1, x_2) - U(x_1, x_2)) / (1 - x_1) > -1 / (1 - x_1).$$

We thus have

$$\begin{aligned} \int_{x_2} U'(x_1, x_2) d\pi &= \int_{x_2} [U'(x_1, x_2)]^+ d\pi + \int_{x_2} [U'(x_1, x_2)]^- d\pi \\ &\geq \int_{x_2 \in g^{-1}((x_1, 1])} U'(x_1, x_2) d\pi - \pi(g^{-1}((a, x_1])) / (1 - x_1). \end{aligned}$$

The negative part goes to 0 as $x_1 \rightarrow a$ by monotone convergence of probabilities. To show that the positive part does not go to zero, observe that $\pi \circ g^{-1}$ must assign positive mass to some interval $(c, 1]$ with $c > a$. For each $x_2 \in g^{-1}((c, 1])$ we have $U(a, x_2) < U(c, x_2)$. This means that for some $\delta > 0$ we have $\pi(\{x_2 : U(a, x_2) < U(c, x_2) - \delta\}) > 0$. But on this event, concavity gives $\liminf_{x_1 \rightarrow a} U' > \delta / (c - a) > 0$, bounding the positive integral away from 0 as desired. $\square$

One can easily show that every $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ with $\operatorname{marg}_{\Theta \times T_{-i}} \pi = \kappa_{t_i}$ and $\pi(a_{-i} \in S_{-i}^{k-1}[B, t_{-i}]) = 1$ is induced by type t_i and a mixed belief $\sigma_{-i} \in \Delta(\hat{S}_{-i}^{k-1}[B])$ where $\hat{S}_{-i}^{k-1}[B] = \prod_{(\theta, t_{-i})} S_j^{k-1}[B, t_j]$ is the set of all functions $\hat{s}_{-i} : \Theta \times T_{-i} \rightarrow A_{-i}$ with $\hat{s}_{-i}(\theta, t_{-i}) \in S_{-i}^{k-1}[B, t_{-i}]$ for each (θ, t_{-i}) . We are ready to prove the following result, which immediately implies Lemma 2.

Lemma 5. For any nice game, any convex $B = B_1 \times \dots \times B_n$, and for any i, t_i, k , the following are true.

- (1) $S_i^k[B, t_i]$ is convex.
 (2) For each $a_i^k \in S_i^k[B, t_i]$, there exists $\hat{s}_{-i} \in \hat{S}_{-i}^{k-1}[B]$ such that

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

Proof. We will use induction on k . For $k=0$, part (1) is true by definition. Assume that part (1) is true for some $k-1$, i.e., $S_j^{k-1}[B, t_j]$ is an interval in $A_j = [0, 1]$ for each j . This implies that $\hat{S}_{-i}^{k-1}[B]$ is connected. Moreover, by the Maximum Theorem, $\beta_i(\cdot; t_i)$ that maps each $s_{-i} \in \hat{S}_{-i}^{k-1}[B]$ to $BR_i(\pi(\cdot|t_i, s_{-i}))$ has a closed-graph and hence is continuous. (It is a function.) Since $\hat{S}_{-i}^{k-1}[B]$ is connected, this implies that $\beta_i(\hat{S}_{-i}^{k-1}[B]; t_i)$ is connected, and hence it is convex as it is unidimensional. We **claim** that $\beta_i(\hat{S}_{-i}^{k-1}[B]; t_i) = S_i^k[B, t_i]$. This readily proves part (1). Part (2) follows from the definition of $\beta_i(\hat{S}_{-i}^{k-1}[B]; t_i)$.

Towards proving our **claim**, let C be the set of functions $s_{-i} : \Theta \times T_{-i} \rightarrow A_{-i}$, so C is a finite product of intervals $[0, 1]$. Define $U : [0, 1] \times C \rightarrow \mathbb{R}$ by setting $U(a_i, s_{-i}) = \int u_i(\theta, a_i, s_{-i}(\theta, t_{-i})) d\kappa_{t_i}$. Since u_i is strictly concave in a_i and continuous in a , U is strictly concave in a_i and continuous. For every $a_i \in S_i^k[B, t_i]$, there exists a belief $\sigma_{-i} \in \Delta(\hat{S}_{-i}^{k-1}[B])$ such that $a_i = \arg \max_{a_i'} \int U(a_i', s_{-i}) d\sigma_{-i}(s_{-i})$. Since $\hat{S}_{-i}^{k-1}[B] \subseteq C$ is a product of intervals by the inductive hypothesis, Lemma 4 then implies that there exists $\hat{s}_{-i}$ such that $a_i = \arg \max_{a_i'} U(a_i', \hat{s}_{-i}) = BR_i(\pi(\cdot|t_i, \hat{s}_{-i}))$. Therefore, $\beta_i(\hat{S}_{-i}^{k-1}[B]; t_i) = S_i^k[B, t_i]$. $\square$

A.3. Proof of Proposition 1

We proceed inductively on k , showing first $S_i^k[s^*(T^u), t_i] \subseteq A_i^k[s^*, t_i]$. For $k=0$, both sides are equal to $s^*(T^u)$. For any given k and any player i , write each $h_{-i}(t_{-i})$ as $h_{-i}(t_{-i}) = (\lambda, \eta)$ where $\lambda = (h_{-i}^1(t_{-i}), h_{-i}^2(t_{-i}), \dots, h_{-i}^{k-1}(t_{-i}))$ and $\eta = (h_{-i}^k(t_{-i}), h_{-i}^{k+1}(t_{-i}), \dots)$ are the lower and higher-order beliefs, respectively. Let $L = \{\lambda \mid \exists \eta : (\lambda, \eta) \in T_{-i}^u\}$. The induction hypothesis is that

$$S_{-i}^{k-1}[s^*(T^u), \lambda] \equiv \bigcup_{\eta'} S_{-i}^{k-1}[s^*(T^u), (\lambda, \eta')] \subseteq A_{-i}^{k-1}[s^*, (\lambda, \eta)] \quad (\forall (\lambda, \eta) \in T_{-i}^u).$$

Fix any type t_i and any $a_i \in S_i^k[s^*(T^u), t_i]$. 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 $a_i \in A_i^k[s^*, t_i]$. Now, by Lemmas 1 and 2, $a_i = BR_i(\pi(\cdot|t_i, \hat{s}_{-i}))$ for some $\hat{s}_{-i} : \Theta^* \times T_{-i}^* \rightarrow A_{-i}$ with $\hat{s}_{-i}(\theta, t_{-i}) \in S_{-i}^{k-1}[s^*(T^u), t_{-i}]$. By the induction hypothesis, for each a_{-i} in the image of $\hat{s}_{-i}$, $a_{-i} \in A_{-i}^{k-1}[s^*, (\lambda, \eta)]$ for some η . Hence, there exists a mapping $\mu : \text{supp}(\text{marg}_{\Theta^* \times L \times A_{-i}} \pi(\cdot|t_i, \hat{s}_{-i})) \rightarrow \Theta^* \times T_{-i}^u$,

$$\mu : (\theta, \lambda, a_{-i}) \mapsto (\theta, \lambda, \tilde{\eta}(a_{-i}, \theta, \lambda)), \quad (\text{A.3})$$

such that

$$s_{-i}^*(\lambda, \tilde{\eta}(a_{-i}, \theta, \lambda)) = a_{-i}. \quad (\text{A.4})$$

We define $\tilde{t}_i$ by

$$\kappa_{\tilde{t}_i} \equiv (\text{marg}_{\Theta^* \times L \times A_{-i}} \pi(\cdot|t_i, \hat{s}_{-i})) \circ \mu^{-1},$$

the probability distribution induced on $\Theta^* \times T_{-i}^u$ by the mapping μ and the probability distribution π . Notice that, since $h_i^k(t_i)$ has finite support and $\hat{s}_{-i}$ is in pure strategies, the set $\text{supp}(\text{marg}_{\Theta^* \times L \times A_{-i}} \pi(\cdot|t_i, \hat{s}_{-i}))$ is finite, in which case μ is trivially measurable. Hence $\kappa_{\tilde{t}_i}$ is well-defined. By a well-known isomorphism of Mertens and Zamir (1985), $\kappa_{\tilde{t}_i}$ is the belief of a type $\tilde{t}_i$, such that

$$h_i^k(\tilde{t}_i) = \text{marg}_{\Theta^* \times L} \kappa_{\tilde{t}_i}. \quad (\text{A.5})$$

Since $\text{supp}(\kappa_{\tilde{t}_i})$ is finite, $h_i(\tilde{t}_i) \in T_i^u$ (as in footnote 6). By construction of μ , the first k orders of beliefs (about (θ, λ)) are identical under t_i and $\tilde{t}_i$:

$$\begin{aligned} \text{marg}_{\Theta^* \times L} \kappa_{\tilde{t}_i} &= \text{marg}_{\Theta^* \times L} [(\text{marg}_{\Theta^* \times L \times A_{-i}} \pi) \circ \mu^{-1}] = \text{marg}_{\Theta^* \times L} (\text{marg}_{\Theta^* \times L \times A_{-i}} \pi) \\ &= \text{marg}_{\Theta^* \times L} \pi = \text{marg}_{\Theta^* \times L} (\text{marg}_{\Theta^* \times T_{-i}^u} \pi) = \text{marg}_{\Theta^* \times L} \kappa_{t_i}, \end{aligned}$$

where the second equality is by (A.3), which implies that μ leaves θ and λ intact. Therefore, $h_i^k(\tilde{t}_i) = h_i^k(t_i)$, which also implies that $h_i^m(\tilde{t}_i) = h_i^m(t_i)$ for each $m \leq k$. Towards showing that $s_i^*(\tilde{t}_i) = a_i$, let $\tilde{\pi} = \kappa_{\tilde{t}_i} \circ \gamma^{-1} \in \Delta(\Theta^* \times T_{-i}^u \times A_{-i})$ be the equilibrium belief of type $\tilde{t}_i$, where $\gamma : (\theta, \lambda, \eta) \mapsto (\theta, \lambda, s_{-i}^*(\lambda, \eta))$. By construction,

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

[By (A.4) and the definition of γ , $\text{proj}_{\Theta^* \times L \times A_{-i}} \circ \gamma \circ \mu$ is the identity mapping, yielding the last equality.] Therefore,

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

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

$$s_i^*(\tilde{t}_i) \in BR_i(\pi(\cdot | \tilde{t}_i, s_{-i}^*)) = BR_i(\text{marg}_{\Theta^* \times A_{-i}} \pi) = \{a_i\}. \quad (\text{A.6})$$

To see the inclusion $A_i^k[s^*, t_i] \subseteq S_i^k[s^*(T^u), t_i]$, observe that for any $\tilde{t}_i$ with $h_i^m(\tilde{t}_i) = h_i^m(t_i)$ for each $m \leq k$, we have

$$s_i^*(\tilde{t}_i) \in S_i^k[s^*(T^u), \tilde{t}_i] = S_i^k[s^*(T^u), t_i],$$

where the last equality is by (2.1).

A.4. 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 $k > 0$, $S^k[B; t^{CK}(\bar{\theta})] = [\underline{q}^k, \bar{q}^k]^n$, where

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

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

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

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, $S^k[B; t^{CK}(\bar{\theta})] = [0, q^M]^n$ for each $k > \bar{k}$. Therefore, $S^\infty[B; t^{CK}(\bar{\theta})] = [0, q^M]^n$.

References

- Basu, K., 1992. A characterization of the class of rationalizable equilibria of oligopoly games. *Econ. Letters* 40, 187–191.
- Battigalli, P., 2003. Rationalizability in infinite, dynamic games with complete information. *Res. Econ.* 57, 1–38.
- Battigalli, P., Siniscalchi, M., 2003. Rationalization and incomplete information. *Advanc. Theoretical Econ.* 3 (1). Article 3.
- Bernheim, D., 1984. Rationalizable strategic behavior. *Econometrica* 52, 1007–1028.
- Börger, T., Janssen, M., 1995. On the dominance solvability of large Cournot games. *Games Econ. Behav.* 8, 297–321.
- Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *J. Econ. Theory* 59, 189–198.
- Chen, Y., 2008. A structure theorem for rationalizability in dynamic games. Northwestern University Working Paper.
- Dekel, E., et al., 2007. Interim correlated rationalizability. *Theoretical Econ.* 2, 15–40.
- Friedenberg, A., Meier, M., 2008. Context of the game. Working Paper.
- Guesnerie, R., 2002. Anchoring economic predictions in common knowledge. *Econometrica* 70, 439–480.
- Harsanyi, J., 1967. Games with incomplete information played by Bayesian players. Part I: The basic model. *Manage. Sci.* 14, 159–182.
- Mertens, J., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. *Int. J. Game Theory* 10, 619–632.
- Moulin, H., 1984. Dominance solvability and Cournot stability. *Math. Soc. Sci.* 7, 83–102.
- Okuguchi, K., Suzumura, K., 1971. Uniqueness of the Cournot oligopoly equilibrium. *Econ. Stud. Quart.* 22, 81–83.
- Penta, A., 2008a. On the structure of rationalizability on arbitrary spaces of uncertainty. PIER Working Paper 09-21.
- Penta, A., 2008b. Higher-order beliefs in dynamic environments. Working Paper.
- Rubinstein, A., 1989. The electronic mail game: strategic behavior under ‘almost common knowledge’. *Amer. Econ. Rev.* 79, 385–391.
- Weinstein, J., Yildiz, M., 2004. Finite-order implications of any equilibrium, MIT Working Paper.
- Weinstein, J., Yildiz, M., 2007a. A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica* 75, 365–400.
- Weinstein, J., Yildiz, M., 2007b. Impact of higher-order uncertainty. *Games Econ. Behav.* 60, 200–212.
- Weinstein, J., Yildiz, M., 2009. A structure theorem for rationalizability in infinite-horizon games. Working Paper.
- Yildiz, M., 2009. Invariance to representation of information. MIT Working Paper.