FINITE-ORDER IMPLICATIONS OF ANY EQUILIBRIUM

JONATHAN WEINSTEIN AND MUHAMET YILDIZ

Abstract. Present economic theories make a common-knowledge assumption that implies that the first or the second-order beliefs determine all higher-order beliefs. We analyze the role of such closing assumptions at finite orders by instead allowing higher orders to vary arbitrarily. Assuming that the space of underlying uncertainty is sufficiently rich, we show that the resulting set of possible outcomes, under an arbitrary fixed equilibrium, must include all outcomes that survive iterated elimination of strategies that are never a strict best reply. For many games, this implies that, unless the game is dominance-solvable, every equilibrium will be highly sensitive to higher-order beliefs, and thus economic theories based on such equilibria may be misleading. Moreover, every equilibrium is discontinuous at each type for which two or more actions survive our elimination process.

Key words: higher-order uncertainty, rationalizability, incomplete information, equilibrium.

JEL Numbers: C72, C73.

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

We thank Jeffrey Ely, Stephen Morris, and the seminar participants at Berkeley, Stanford, and Yale for helpful comments.

> "Game theory . . . is deficient to the extent it assumes other features to be common knowledge, such as one player's probability assessment about another's preferences or information. I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of common knowledge assumption will the theory approximate reality." Wilson (1987)

1. Introduction

Most economic theories are based on equilibrium analysis of models that are closed after specifying the first and second-order beliefs, i.e., the beliefs about underlying uncertainty and the beliefs about other players' beliefs about underlying uncertainty. These models assume that the specified belief structure is common knowledge, i.e., conditional on the first and the second-order beliefs, all of the players' higher-order beliefs are common knowledge. Since these assumptions may easily fail in the in the actual incomplete-information situation modeled, these theories may be misleading when the impact of higher-order beliefs on equilibrium behavior is large. There are examples that suggest that this impact might indeed be large in some situations (see Rubinstein (1989), Feinberg and Skrzypacz (2002), and also Milgrom and Weber (1985)). To overcome this fundamental deficiency, one may want to close the model at higher orders, specifying more orders of beliefs, and hence weakening the common knowledge assumption. As Wilson (1987), one might hope that by specifying more and more orders of beliefs, the theory would "approximate reality." We demonstrate that this is not the case. We show that regardless of how many orders of beliefs are specified, the closing assumption is required to gain any predictive power beyond that of iterated elimination of strategies that are never strict best reply.

Consider a situation where players have incomplete information about some payoff-relevant parameter. Each player has a probability distribution about the parameter, which represents his first-order beliefs, a probability distribution about other players' first-order beliefs, which represents his second-order beliefs, and so on. Imagine a researcher who has computed an equilibrium of this game, where a type of a player is an infinite hierarchy of his beliefs,[1] and would like to make a prediction about the action of a player i according to this equilibrium. Fix a type t_i of player i as his actual type, and write $A_i^1(t_i)$ for the set of all actions that are played by some alternative type of i whose first-order beliefs agree with t_i. This set is the set of actions that the researcher cannot rule out if he only knows the first-order beliefs and assumes that the player plays according to the equilibrium. Similarly, write $A_i^k(t_i)$ for the set of actions that the researcher cannot rule out by looking at the first k orders of beliefs. Write $A_i^\infty(t_i)$ for the limit of these (decreasing) sets as k approaches infinity, i.e., the set of all actions that cannot be ruled out by the researcher by looking at (arbitrarily many) finite orders of beliefs. This definition can be put another way. Consider two researchers who agree on the equilibrium played. One researcher is certain that player i is of type t_i. The other (slightly suspicious) researcher is willing to agree with this assessment for the first k orders of beliefs but does not have any further assumption. The set $A_i^k(t_i)$ is precisely the set of actions that will not be ruled out by the second researcher.

In a model that is closed at order k, all higher-order beliefs are determined by the first k orders of beliefs and the assumption that is made when the model is closed. We wish to emphasize the sensitivity of the model's predictions to the closing assumption. In a given equilibrium, the model predicts a unique action for each possible set of beliefs at orders 1 through k, namely the equilibrium

[1] For technical reasons, we assume that beliefs at each finite order have countable support.

action for the complete type implied by this set of beliefs and the closing assumption. But in the general model, every other action in $A_i^k(t_i)$ is played by a type whose first k orders of beliefs will be exactly as this type (but will fail the closing assumption.) Therefore, we cannot rule out any action in $A_i^k(t_i)$ without resorting to the closing assumption.

Our main result gives a lower bound for $A_i^k(t_i)$. We assume that the space of underlying uncertainty is rich enough so that our fixed equilibrium has full range, i.e., every action is played by some type. For countable-action games, we show that $A_i^k(t_i)$ includes all actions which survive the first k iterations of eliminating all actions which are never a *strict* best reply under t_i. In particular, $A_i^\infty(t_i)$ includes all actions that survive iterated elimination of actions that cannot be a strict best reply. On the other hand, $A_i^k(t_i)$ is a subset of actions that survive the first k iterations of eliminating strictly dominated actions, and hence $A_i^\infty(t_i)$ is a subset of rationalizable actions. When there are no ties, these elimination procedures lead to the same outcome, and therefore $A_i^\infty(t_i)$ is precisely equal to the set of rationalizable outcomes. We extend this characterization to *nice games*, where the action spaces are one-dimensional compact intervals and the the utility functions are strictly concave in own action and continuous—as in many classical economic models.

To illustrate the main argument in the proof of the lower bound, we now explain why $A_i^1(t_i)$ includes all actions that survive the first round of elimination process. Let $\tilde{t}_i$ vary over the set of types that agree with t_i at first order (i.e., concerning the underlying parameter) but may have any beliefs at higher orders (i.e., concerning the other players' type profile.) Our full range assumption implies that there are types $\tilde{t}_i$ with any beliefs whatsoever about other players' equilibrium action profile. Given any action a_i of i that is a strict best reply to his fixed belief about the parameter and some belief about the other players' actions, there is a type $\tilde{t}_i$ who has these beliefs in equilibrium, and therefore must play the strict best reply, a_i, in equilibrium. This argument will be formalized as part of an inductive proof of the main result.

For general games, Nash equilibrium has weak epistemic foundations (Aumann and Brandenburger (1995)) in comparison to iterative admissibility (Brandenburger and Keisler (2000)) and rationalizability (Bernheim (1985), Pearce (1985)). Yet, in application, researchers frequently use equilibrium analysis and further focus on a particular equilibrium, invoking refinement arguments and such, so that they can make predictions. Our result shows that the predictive power—beyond that of our elimination process—obtained in this way comes from the assumption that is (implicitly) made when the model is closed. That is, for any such prediction, there are types that are ruled out by the closing assumption and that behave inconsistently with the prediction in the focused equilibrium. Therefore, our result suggests that the closing assumption deserves a close scrutiny, and needs to be justified at least as much as the explicit assumptions of the model. It would be highly desirable to investigate whether the types that behave inconsistently with the prediction of the closed model can be excluded by a weaker set of assumptions.

From an evolutionary point of view, when there are privately observed signals, if a myopic adjustment process converges, its limit is a Nash equilibrium of the incomplete-information game with these signals as private information. Our result suggests then that the limit behavior of such a process may depend on the elusive signals about the distribution of other players' signals, higher-order signals about these signals, and so on. Incidentally, our result has counterparts in sophisticated and Bayesian learning models: the learning of sophisticated agents leads to equilibrium if and only if the game is dominance solvable (Milgrom and

Roberts (1991)), and in a specific model, any sequence of rationalizable action profiles can be a sample path in Bayesian learning (Nyarko (1996)).

Our result also points to a close link between higher-order reasoning and the equilibrium impact of higher-order uncertainty. When there are no ties, assuming kth-order mutual knowledge of payoffs and that a fixed equilibrium (with full range) is played is equivalent to assuming kth-order mutual knowledge of rationality and common knowledge of payoffs.[2] This implies that when the equilibrium impact of high-order uncertainty is large, the impact of high-order failures of rationality is also large. In that case, predictions may be unreliable without a very accurate knowledge of players' reasoning capacity.

It is well known that some Nash equilibria may be discontinuous in product topology and with respect to higher-order uncertainty, as in the electronic-mail game of Rubinstein (1989). There is an interest in understanding how severe this discontinuity is. Monderer and Samet (1989,1997) and Kajii and Morris (1998) have analyzed the weakest topologies that make the equilibrium continuous over all games (see also Milgrom and Weber (1985) for a continuity result.) These topologies are quite strong, but since they focus on the worst case games, such as the electronic-mail game, it is not clear whether the equilibria used in applications will be highly sensitive to higher-order uncertainty. Our result implies (and we formally establish) that, if the space of underlying uncertainty is sufficiently rich, *every equilibrium* is discontinuous (with respect to product topology and higher-order beliefs) for *every game* at *every type* for which two or more actions survive our elimination process.

[2]The relationship between assumptions about rationality and payoff uncertainty is not straightforward; A_i^∞ may differ from both rationalizability and iterative admissibility.

As a precedent to our main result, Brandenburger and Dekel (1987) show that every rationalizable outcome is the outcome of a subjective correlated equilibrium (see Section 3 for a discussion.) Battigalli and Siniscalchi (2003) extended this result to dynamic games and investigated the implications of the restrictions on first-order beliefs and common strong belief in sequential rationality on the rationalizable outcomes, which coincide with all equilibrium outcomes. It seems that, using their methodology (and that of Battigalli (2003)), one can obtain sharp predictions in sequential games using relatively mild assumptions. Note that in sequential games, our lower bound is usually weak, and assumptions about sequential rationality yield strong predictions (Battigalli and Siniscalchi (2002) and Feinberg (2002)).

Our next section contains the basic definitions and preliminary results. Section 3 is the heart of the paper. There we develop our main notions and prove our main theorem. Our main theorem is extended to the nice games as a characterization in Section 4, and to mixed strategies and to the spaces of uncertainty that are not necessarily rich in Section 5. In Section 6, we present our discontinuity results and discuss their methodological implications for global games and robustness of equilibria. Section 7 contains a very negative result about Cournot oligopoly as an application. Section 8 concludes. Some of the proofs are relegated to the appendix.

2. Basic Definitions and Preliminary Results

Notation 1. 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, suppressing the fixed σ-algebra on

Y which at least contains all open sets and singletons; we use the product σ-algebra in product spaces. The support of π is denoted by suppπ.

We consider a game with finite set of players $N = \{1, 2, \ldots, n\}$. The source of underlying uncertainty is a payoff-relevant parameter $\theta \in \Theta$ where (Θ, d) is a compact, complete and separable metric space, with d a metric on set Θ. Each player i has action space A_i and utility function $u_i : \Theta \times A \to \mathbb{R}$, where $A = \prod_i A_i$. We endow the game with the universal type space of Brandenburger and Dekel (1993), a variant of an earlier construction by Mertens and Zamir (1985), with an additional assumption that the players' beliefs at each finite order have countable (or finite) support.[3] Types are defined using the auxiliary sequence $\{X_k\}$ of sets defined 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 $\hat{\Delta}\left(X_{k-1}\right)$ is the set of probability distributions on X_{k-1} that have countable (or finite) support. We endow each X_k with the weak topology and the σ-algebra generated by this topology. A player i's first order beliefs (about the underlying uncertainty θ) are represented by a probability distribution t_i^1 on X_0, second order beliefs (about all players' first order beliefs and the underlying uncertainty) are represented by a probability distribution t_i^2 on X_1, etc. Therefore, a *type* t_i of a player i is a member of $\prod_{k=1}^{\infty} \hat{\Delta}\left(X_{k-1}\right)$. Since a player's kth-order beliefs contain information about his lower-order beliefs, we need the usual coherence requirements. We write $T = \prod_{i \in N} T_i$ for 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. We will use the variables $t_i, \tilde{t}_i \in T_i$ as generic types of any player i and $t, \tilde{t} \in T$ as generic type profiles. For every $t_i \in T_i$, there exists

[3]This assumption is made to 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.

a probability distribution κ_{t_i} on $\Theta \times T_{-i}$ such that

$$t_i^k = \delta_{t_i^{k-1}} \times \mathrm{marg}_{\Theta \times [\Delta(X_{k-2})]^{N\setminus\{i\}}} \kappa_{t_i}, \qquad (\forall k) \tag{2.1}$$

and $t_i^1 = \mathrm{marg}_{\Theta}\kappa_{t_i}$, where $\delta_{t_i^{k-1}}$ is the probability measure that puts probability 1 on the set $\left\{t_i^{k-1}\right\}$ and marg denotes the marginal distribution. Conversely, given any distribution κ_{t_i} on $\Theta \times T_{-i}$, we can define $t_i \in T_i$ via (2.1), as long as $\mathrm{marg}_{\Theta \times [\Delta(X_{k-2})]^{N\setminus\{i\}}} \kappa_{t_i}$ is always countable.

A *strategy* of a player 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 $\pi(\cdot|t_i, s_{-i}) \in \Delta(\Theta \times A_{-i})$ for the joint distribution of the underlying uncertainty and the other players' actions induced by t_i and s_{-i}; $\pi(\cdot|t_i, \sigma_{-i})$ is similarly defined for correlated mixed strategy profile σ_{-i}. For each $i \in N$ and for each belief $\pi \in \Delta(\Theta \times A_{-i})$, we write $BR_i(\pi)$ for the set of actions $a_i \in A_i$ that maximize the expected value of $u_i(\theta, a_i, a_{-i})$ under the probability distribution π. A strategy profile $s^* = (s_1^*, s_2^*, \ldots)$ is a *Bayesian Nash equilibrium* iff at each t_i,

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

An equilibrium s^* is said to have *full range* iff

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

The following assumption implies that every equilibrium s^* has full range.

Assumption 1 (Richness of Θ)**.** *Given any $i \in N$, any $\mu \in \Delta(A_{-i})$, and any a_i, there exists a probability distribution ν on Θ with countable support and such that*

$$BR_i(\nu \times \mu) = \{a_i\}.$$

Lemma 1. *Under Assumption 1, every equilibrium s^* has full range.*

Proof. The proofs that are omitted in the text are in the appendix. □

Elimination Processes. We will use interim notions and allow correlations not only within players' strategies but also between their strategies and the underlying uncertainty θ. Such correlated rationalizability is introduced by Battigalli (2003), Battigalli and Siniscalchi (2003) and Dekel, Fudenberg, and Morris (2003). 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. Write M_i for the set of all measurable functions from $\Theta \times T_i$ to A_i. Towards defining rationalizability, define sets $S_i^k[t_i]$, $i \in N$, $t_i \in T_i$, $k = 0, 1, \ldots$, iteratively as follows. Set $S_i^0[t_i] = A_i$. For each $k > 0$, let $\hat{S}_{-i}^{k-1} \subset M_{-i}$ be the set of all measurable functions $f : \Theta \times T_{-i} \to A_{-i}$ such that $f(\theta, t_{-i}) \in S_{-i}^{k-1}[t_{-i}]$ for each t_{-i}. Let also Σ_{-i}^{k-1} be the set of all probability distributions on $\hat{S}_{-i}^{k-1}$. Note that Σ_{-i}^{k-1} is the set of all possible beliefs of player i on other players' allowable actions that are not eliminated in the first $k-1$ rounds. Write

$$S_i^k[t_i] = \bigcup_{\sigma_{-i} \in \Sigma_{-i}^{k-1}} BR_i(\pi(\cdot|t_i, \sigma_{-i}))$$

for the set of all all actions a_i of i that are best reply against some of his beliefs in Σ_{-i}^{k-1}. The set of all rationalizable actions for player i (with type t_i) is

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

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

$$W_i^k[t_i] = \left\{ a_i | BR_i(\pi(\cdot|t_i, \sigma_{-i})) = \{a_i\} \text{ for some } \sigma_{-i} \in \Delta\left(\hat{W}_{-i}^{k-1}\right) \right\},$$

where $\hat{W}_{-i}^{k-1} \subset M_{-i}$ is the set of all functions $f : \Theta \times T_{-i} \to A_{-i}$ such that $f(\theta, t_{-i}) \in W_{-i}^{k-1}[t_{-i}]$ for each t_{-i}. 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 admissibility (i.e., iterative elimination of weakly dominated strategies).[4] In some games, iterative admissibility may yield strong predictions. For example, in finite perfect information games it leads to backwards induction outcomes. Nevertheless, in generic normal-form games (as defined in Definition 1 below), all these concepts are equivalent and usually have weak predictive power.

3. Equilibrium predictions with finite-order beliefs

We are interested in how robust equilibrium is against the failure of assumptions made at high orders, such as the failure of the common knowledge assumption at high orders. We now formalize our notion of robustness.

Let us fix an equilibrium s^* 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 kth-order 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

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

[4]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

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

We are now ready to prove our main result for countable-action games, i.e., games where each player i has a countable or finite action space A_i.

Proposition 1. *For any countable-action game, 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,

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

Proof. The inclusion $A_i^k [s^*, t_i] \subseteq S_i^k [t_i]$ is established by Proposition 8 in the Appendix for general games and all equilibria. We will now show that $W_i^k [t_i] \subseteq A_i^k [s^*, t_i]$. For $k = 0$, the statement is given by 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 $L = \{l | \exists h : (l, h) \in T_{-i}\}$. The induction hypothesis is that

$$W_{-i}^{k-1} [l] \equiv \bigcup_{h'} W_{-i}^{k-1} [(l, h')] \subseteq A_{-i}^{k-1} [s^*, (l, h)] \quad (\forall (l, h) \in T_{-i}).$$

Fix any type t_i and any $a_i \in W_i^k [t_i]$. We will construct a type $\tilde{t}_i$ such that $s_i^* \left(\tilde{t}_i\right) = a_i$ 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 definition, for some $\sigma_{-i} \in \Delta(\hat{W}_{-i}^{k-1})$, a_i is the unique best reply for type t_i if t_i assigns probability distribution σ_{-i} on the other players' strategies, i.e., $BR_i (\pi (\cdot | t_i, \sigma_{-i})) = \{a_i\}$. Let $P (\cdot | t_i, \sigma_{-i})$ be the probability distribution on $\Theta \times L \times A_{-i}$ induced by κ_{t_i} and σ_{-i}. By the induction hypothesis, for each $(\theta, l, a_{-i}) \in \mathrm{supp} P(\cdot|t_i, \sigma_{-i})$, $a_{-i} \in W_{-i}^{k-1}[l] \subseteq A_{-i}^{k-1}[s^*, (l,h)]$ for some h. Hence, there exists a mapping $\mu : \mathrm{supp} P(\cdot|t_i, \sigma_{-i}) \to \Theta \times T_{-i}$,

$$\mu : (\theta, l, a_{-i}) \mapsto \left(\theta, l, \tilde{h}(a_{-i}, \theta, l)\right), \tag{3.1}$$

such that

$$s_{-i}^*\left(l, \tilde{h}(a_{-i}, \theta, l)\right) = a_{-i}. \tag{3.2}$$

We define $\tilde{t}_i$ by

$$\kappa_{\tilde{t}_i} \equiv P(\cdot|t_i, \sigma_{-i}) \circ \mu^{-1},$$

the probability distribution induced on $\Theta \times T_{-i}$ by the mapping μ and the probability distribution $P(\cdot|t_i, \sigma_{-i})$. Notice that, since t_i^k has countable support and the action spaces are countable, the set $\mathrm{supp} P(\cdot|t_i, \sigma_{-i})$ is countable, in which case μ is trivially measurable. Hence $\kappa_{\tilde{t}_i}$ is well-defined. By construction of μ, the first k orders of beliefs (about (θ, l)) are identical under t_i and $\tilde{t}_i$:

$$\mathrm{marg}_{\Theta \times L} \kappa_{\tilde{t}_i} = \mathrm{marg}_{\Theta \times L} P(\cdot|t_i, \sigma_{-i}) \circ \mu^{-1} = \mathrm{marg}_{\Theta \times L} P(\cdot|t_i, \sigma_{-i}) = \mathrm{marg}_{\Theta \times L} \kappa_{t_i},$$

where the second inequality is by (3.1) and the last equality is by definition of $P(\cdot|t_i, \sigma_{-i})$. Moreover, using the mapping $\gamma : (\theta, l, h) \mapsto \left(\theta, l, s_{-i}^*(l,h)\right)$, we can check that the distribution induced by $\kappa_{\tilde{t}_i}$ and s_{-i}^* on $\Theta \times L \times A_{-i}$ is

$$P\left(\cdot|\tilde{t}_i, s_{-i}^*\right) \equiv \kappa_{\tilde{t}_i} \circ \gamma^{-1} = P(\cdot|t_i, \sigma_{-i}) \circ \mu^{-1} \circ \gamma^{-1} = P(\cdot|t_i, \sigma_{-i}),$$

where the last equality is due to the fact that μ is the inverse of the restriction of γ to $\mathrm{supp}\kappa_{\tilde{t}_i}$. Therefore,

$$\pi\left(\cdot|\tilde{t}_i, s_{-i}^*\right) = \mathrm{marg}_{\Theta \times L} P\left(\cdot|\tilde{t}_i, s_{-i}^*\right) = \mathrm{marg}_{\Theta \times L} P(\cdot|t_i, \sigma_{-i}) = \pi(\cdot|t_i, \sigma_{-i}).$$

That is, the equilibrium beliefs of $\tilde{t}_i$ about $\Theta \times A_{-i}$ are identical to the beliefs of t_i about $\Theta \times A_{-i}$ when t_i assigns probability distribution σ_{-i} on the other players' strategies. Since a_i is the only best reply to these beliefs, $\tilde{t}_i$ must play a_i in equilibrium:

$$s_i^*\left(\tilde{t}_i\right) \in BR_i\left(\pi\left(\cdot|\tilde{t}_i, s_{-i}^*\right)\right) = BR_i\left(\pi\left(\cdot|t_i, \sigma_{-i}\right)\right) = \{a_i\}\,. \tag{3.3}$$

$\square$

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. In the next section, we present another class of games in which μ is measurable; μ may not be measurable in general.

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 are as specified. Then, we cannot rule out that a_i will be played in equilibrium s^*. Put it differently, if we have a model that is closed at order k and if an action survives k rounds of iterated elimination of strategies that cannot be a best reply for a type within the model, then we cannot rule out action a_i as an equilibrium action for that type without invoking the closing assumption. Hence, any prediction that does not follow from the first k steps of this elimination process comes from the closing assumptions, rather than the assumptions that lead to a specific equilibrium s^*. This suggests that, contrary to the current practice in economics, a researcher needs to justify his closing assumption at least as much as the other assumptions, such as the rationale for equilibrium selection.

The part $A_i^k[s^*, t_i] \subseteq S_i^k[t_i]$ is proved in the Appendix for general games and for all equilibria (see Proposition 8). Although the proof for the general case is somewhat involved, it is straightforward for the complete-information case.

Notation 2. For any $\bar{\theta} \in \Theta$, write $t_i^{CK}(\bar{\theta})$ for the type of a player i who is certain that it is common knowledge that $\theta = \bar{\theta}$.

For the types of the form $t_i^{CK}(\bar{\theta})$, the proof is as follows. Notice that $A_i^k\left[s^*, t_i^{CK}(\bar{\theta})\right] \subseteq S_i^k\left[t_i^{CK}(\bar{\theta})\right]$ means that, if a player i has kth-order knowledge of $\theta = \bar{\theta}$, then his equilibrium action (according to s^*) will survive kth round of iterated elimination of strictly dominated strategies under the restriction that $\Theta = \{\bar{\theta}\}$. Towards an induction, assume that $A_j^{k-1}\left[s^*, t_j^{CK}(\bar{\theta})\right] \subseteq S_j^k\left[t_j^{CK}(\bar{\theta})\right]$ for each j. That is, if a player j has $k-1$st-order knowledge of $\theta = \bar{\theta}$, j will play an action in S_j^{k-1} (under $\Theta = \{\bar{\theta}\}$.) But any type of player i who has kth-order knowledge of $\theta = \bar{\theta}$ is certain that every other player j has $k-1$st-order knowledge of $\theta = \bar{\theta}$. Hence, in equilibrium, he is certain that j plays an action in S_j^{k-1}. Since his equilibrium action is a best response to such a belief (with support contained in S_{-i}^{k-1}), it must survive the kth round of elimination.

Brandenburger and Dekel (1987) show that, given any rationalizable outcome of a game, by adding payoff irrelevant types, we can construct a type space with an equilibrium that yields the original rationalizable outcome.[5] Since we need to choose a different equilibrium for each rationalizable outcome, there is a large multiplicity of equilibria, and thus equilibrium as a solution concept does not have more predictive power than rationalizability has. One may want to ignore this multiplicity by focusing on a particular equilibrium or relying on one of the many refinements developed in the last few decades—in order to cope with the

[5]See Bergemann and Morris (2003) for an important application of this adea.

usual multiplicity problem. Our result shows that, once we abandon artificial restrictions on type space, this unpredictability reappears as high variability of any fixed equilibrium with respect to very high-order beliefs, albeit with somewhat smaller scope due to the stronger elimination process. There is an intuitive relation between these two results. Since type spaces in general can be embedded in the universal type space, given any fixed set of beliefs at finite orders and a fixed equilibrium behavior in a fixed type space, we can envision a coherent type whose lower-order beliefs are as the former but whose higher-order beliefs assigns high probability to the latter. That type will give a best reply to the equilibrium in the latter. Hence, multiplicity of equilibria in various type spaces tends to yield high variability of equilibria with respect to the higher-order beliefs in universal type space. This intuition, however, does not yield a proof. This is because the distinctions among the types in Brandenburger and Dekel are all payoff irrelevant, and thus their type spaces are *not* contained in the universal type space in general. More importantly, the equilibrium behavior in various type spaces need not be different within a fixed equilibrium.

Our next example shows that either of the inclusions in Proposition 1 may be strict. Hence, (i) some rationalizable strategies may not be in A_i^∞, showing the distinction between the results of Brandenburger and Dekel and ours, and (ii) A_i^∞ may include some weakly dominated strategies, distinguishing our result from the characterization of Brandenburger and Keisler (2000).

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 $A_i^k\left[s^*; t_i^{CK}(\theta_1)\right] = \{a^1\}$.

Proposition 1 yields a characterization whenever the payoffs are generic in the following (standard) sense.

Definition 1. 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) $\sum_{a_i} \alpha(a_i) u_i(\theta, a_i, a_{-i}) = \sum_{a_i} \alpha(a_i) u_i\left(\theta, a_i, a_{-i}'\right) = 0$.

When the payoffs are generic at $\bar{\theta}$ and it is common knowledge that $\theta = \bar{\theta}$, then 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. In that case, 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,*

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

That is, in generic finite-action games, a researcher's predictions based on finite orders of players' beliefs and equilibrium will be equivalent to the predictions that follow from rationalizability. This characterization will be generalized to the following widely-used class of games.

4. Nice games

We will now consider a 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 the firm, etc. We will show that $A_i^k[s^*, t_i] = S_i^k[t_i]$ for each k whenever equilibrium s^* has full range.

Definition 2. 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 in a_i.

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.) We use the continuity assumption to make sure that a player's strategy best response is continuous with respect to the other players' strategies. For the complete information types, our results in this section will be true under the weaker condition that $u_i(\theta, \cdot, a_{-i})$ is single-peaked with a maximand that is continuous in a_{-i}. Now, since our players have always unique best reply, our elimination processes will be equivalent, yielding the functional equation

$$W = S. \tag{4.1}$$

Moreover, our next lemma ensures that, despite our uncountable action spaces, we only need to consider countably many actions for types with countable supports, allowing us to circumvent the measurability issue discussed in Remark 1.

Lemma 2. *For any nice game, for any i, t_i, k, and any $a_i \in S_i^k[t_i]$, there exists $\hat{s}_{-i} \in \hat{S}_{-i}^{k-1}$ such that*

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

Together with (4.1), Lemma 2 gives us our main result for this section.

Proposition 2. *For any nice game, let s^* be any equilibrium with full range. Let also $\hat{\Theta} \times \hat{T}$ be a countable subset of $\Theta \times T$ such that for each $\hat{t}_i \in \hat{T}_i$, supp$\kappa_{\hat{t}_i} \subseteq \hat{\Theta} \times \hat{T}_{-i}$. Then, for any $k \in \mathbb{N}$, $i \in N$, and $\hat{t}_i \in \hat{T}_i$,*

$$S_i^k \left[\hat{t}_i\right] = A_i^k \left[s^*, \hat{t}_i\right];$$

in particular,

$$S_i^\infty \left[\hat{t}_i\right] = A_i^\infty \left[s^*, \hat{t}_i\right].$$

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

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

Hence our proof of Proposition 1 applies. That is, there exists $\tilde{t}_i \in T_i$ (not necessarily in $\hat{T}_i$) such that $s_i^*\left(\tilde{t}_i\right) = a_i$ and $\tilde{t}_i^m = \hat{t}_i^m$ for each $m \leq k$. □

5. Extensions

For ease of exposition, we have so far focused on pure strategy equilibria with full range. In this section, we will extend our results for mixed strategy equilibria and beyond the full-range assumption.

5.1. Mixed Strategies.

Since all equilibria in nice games are in pure strategies, we will focus on the countable-action games. Using interim formulation, we define a *mixed strategy* as any measurable function $\sigma_i : T_i \to \Delta\left(A_i\right)$. A mixed strategy profile σ^* is *Bayesian Nash equilibrium* iff supp$\sigma_i^*\left(t_i\right) \subseteq BR_i\left(\pi\left(\cdot|t_i, \sigma_{-i}^*\right)\right)$ for each i and t_i. Writing $T_i^{\sigma^*} = \left\{t_i | \left|\text{supp}\sigma_i^*\left(t_i\right)\right| = 1\right\}$ for the set of types who play pure strategies, we define a mapping $s_i^{\sigma^*} : T_i^{\sigma^*} \to A_i$

by $\text{supp}\sigma_i^*(t_i) = \left\{s_i^{\sigma^*}(t_i)\right\}$. We then use this "pure part of" σ^* to extend our previous definitions and results to mixed strategies. We say that σ^* has full range iff $s^{\sigma^*}\left(T^{\sigma^*}\right) = A$ and set

$$A_i^k[\sigma^*; t_i] = \left\{s_i^{\sigma^*}\left(\tilde{t}_i\right) | \tilde{t}_i \in T_i^{\sigma^*}, \tilde{t}_i^m = t_i^m \forall m \leq k\right\},$$

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. Clearly, every equilibrium has full range under Assumption 1, i.e., when Θ is sufficiently rich.

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,*

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

Proof. In the proof of Proposition 1, insert s^{σ^*} everywhere s^* appears, and restrict the range of μ and the domain of γ to $\Theta \times T_{-i}^{\sigma^*}$. Notice that, by (3.3), $\tilde{t}_i \in T_i^{\sigma^*}$. □

That is, if σ^* has full range (e.g., if Θ is sufficiently rich) 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^*. For $k \geq 1$, the full-range assumption can replaced by the weaker assumption that $A_i \subseteq \cup_{t_i} \text{supp}\sigma_i^*(t_i)$.

5.2. **Without full range.** Our full range assumption allowed us to consider large changes. 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 extension of our main result to such cases.

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

$$\begin{aligned} S_i^0 [B; t_i] &= B_i, \\ S_i^k [B; t_i] &= \bigcup_{\sigma_{-i} \in \Delta\left(\hat{S}_{-i}^{k-1}[B]\right)} BR_i \left(\pi \left(\cdot | t_i, \sigma_{-i}\right)\right), \end{aligned}$$

where $\hat{S}_{-i}^{k-1} [B] \subset M_{-i}$ is the set of all measurable functions $f : \Theta \times T_{-i} \to A_{-i}$ such that $f(\theta, t_{-i}) \in S_{-i}^{k-1} [B; t_{-i}]$ for each t_{-i}. Notice that this is the same procedure as iterated strict dominance, except that the initial set is restricted to a subset. Unlike iterated strict dominance, these sets can become larger as k increases. Hence we define the set of *locally rationalizable* strategies by

$$S_i^\infty [B; t_i] = \bigcap_{k=0}^{\infty} \bigcup_{m=k}^{\infty} S_i^m [B; t_i].$$

Notice that the set $S^\infty [B; t_i]$ may be much larger than B. We define local version of W^∞, similarly, by setting $W_i^0 [B; t_i] = B_i$,

$$W_i^k [B; t_i] = \left\{ a_i \in A_i | BR_i \left(\theta^i, \sigma_{-i}\right) = \{a_i\} \text{ for some } \sigma_{-i} \in \Delta \left(\hat{W}_{-i}^{k-1} [B]\right) \right\},$$

and $W_i^\infty [B; t_i] = \bigcap_{k=0}^\infty \bigcup_{m=k}^\infty W_i^m [B; t_i]$. Notice that we consider all actions in our process, which is no longer an elimination process.

Proposition 4. *For any equilibrium s^*, if the game has countable action spaces, then*

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

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

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

The last statement 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 in Section 7.

6. Continuity of Equilibrium

It is well known that equilibrium may be discontinuous with respect to the product topology. In this section we will introduce our notion of continuity with respect to higher-order beliefs, which appears much weaker than continuity with respect to the product topology. We will show that even this weaker continuity property is violated in every equilibrium on a very large set.

6.1. **Equilibrium in pure strategies.** We consider an arbitrary metric d on A. 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 $d(a^m, a) < \epsilon$ for each $m > k$. Given any subset $B \subseteq A$, we write $D(B) = \sup\{d(a,b)\,|a, b \in B\}$ for the diameter of B. As usual, we write $A^k[s^*, t] = \prod_i A_i^k[s^*, t_i]$, $S^k[s^*, t] = \prod_i S_i^k[t_i]$, etc.

Definition 3. An equilibrium 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

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

An equilibrium s^* is said to be *continuous with respect to higher-order beliefs at* t iff for each $\epsilon > 0$, there exists k such that for each $\tilde{t}$,

$$\left[\tilde{t}^m = t^m \quad \forall m \le k\right] \Rightarrow d\left(s^*\left(\tilde{t}\right), s^*(t)\right) < \epsilon.$$

The latter continuity concept is uniform continuity with respect to the product topology (on the type space) of discrete topologies on each order of beliefs. Of course, continuity with respect to discrete topology is much weaker than other topologies. The next lemma presents some basic facts.

Lemma 3. *For any equilibrium s^* and for any t, the following are true.*

(1) *If s^* is continuous at t, then s^* is continuous with respect to higher-order beliefs at t.*
(2) *s^* is continuous with respect to higher-order beliefs at t iff $D\left(A^k\left[s^*, t\right]\right) \to 0$ as $k \to \infty$.*
(3) *If s^* is continuous with respect to higher-order beliefs at t, then $A^\infty\left[s^*, t\right] = \{s^*(t)\}$.*

For nice games, Lemma 3.3 and Proposition 2 imply that if an equilibrium s^* with full range is continuous with respect to higher-order beliefs at t, then $S^\infty[t] = \{s^*(t)\}$, yielding the following discontinuity result. (One can obtain a similar result for countable-action games by replacing $S^\infty[t] = \{s^*(t)\}$ with $|W^\infty[t]| \leq 1$.)

Proposition 5. *For any nice game, every equilibrium s^* with full range is discontinuous with respect to the higher-order beliefs (and the product topology) 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 with respect to higher-order beliefs or with respect to the product topology, then the game is dominance solvable.*

6.2. **Mixed-strategy equilibria in finite-action games.** Endow the space of mixed action profiles, $\Delta(A)$, with an arbitrary metric d. Extend the definitions of continuity with respect to product topology and higher-order beliefs to mixed-strategy equilibria σ^* by replacing s^* with σ^* in Definition 3.

Definition 4. A mixed-strategy equilibrium σ^* is said to be *weakly continuous with respect to higher-order beliefs at t* iff there exists k such that for each $\tilde{t}$,

$$\left[\tilde{t}^m = t^m \quad \forall m \leq k\right] \Rightarrow \operatorname{supp}\left(\sigma^*\left(\tilde{t}\right)\right) \cap \operatorname{supp}\left(\sigma^*(t)\right) \neq \emptyset.$$

Now, continuity in product topology implies continuity with respect to higher-order beliefs. For finite-action games, the latter in turn implies weak continuity with respect to higher-order beliefs, as we show in the Appendix (see Lemma 6). There we also show that strong and weak continuity of σ^* at t with respect to higher-order beliefs imply that $|A^\infty[\sigma^*, t]| \leq 1$ and $A^\infty[\sigma^*, t] \subseteq \mathrm{supp}(\sigma^*(t))$, respectively. This yields the following result.

Proposition 6. *For any finite-action game and any equilibrium σ^* with full range, $|W^\infty[t]| \leq 1$whenever (i) σ^* is continuous with respect to higher-order beliefs at t, or (ii) σ^* is weakly continuous with respect to higher-order beliefs at t and $\sigma^*(t)$ is pure.*

Proof. Either of the conditions (i) and (ii) implies that $|A^\infty[\sigma^*, t]| \leq 1$ (see Lemma 6). Hence, by Proposition 3, $|W^\infty[t]| \leq |A^\infty[\sigma^*, t]| \leq 1$. □

That is, for sufficiently rich Θ, every equilibrium is discontinuous with respect to higher-order beliefs (and the product topology) at each type profile for which two or more action profiles survive iterated elimination of strategies that cannot be a strict best reply. These type profiles include the generic instances of complete-information without dominance solvability. At each such type profile, even the weakest continuity property fails if the equilibrium actions are pure.

Example 2. Consider the coordinated attack game with payoff matrix

	Attack	No Attack
Attack	1,1	-2,0
No Attack	0,-2	0,0

where there are two pure strategy equilibria: the efficient equilibrium (Attack, Attack) and the risk-dominant equilibrium (No Attack, No Attack). Since each action is a strict best reply, no action is eliminated in our elimination process. Therefore, Lemma 1 and Proposition 6 imply that when we embed the coordinated-attack game in a rich type space as a type profile, every equilibrium must be discontinuous with respect to higher-order beliefs at that type profile.

Rubinstein's (1989) electronic-mail game presents a type space in which any equilibrium that selects the efficient equilibrium in the coordinated-attack game must be discontinuous with respect to higher-order uncertainty. In that example there is also a continuous equilibrium, which selects the risk-dominant action profile for each type profile. Our example shows that the latter continuous equilibrium is an artifact of the small type space utilized, and in fact in a rich type space, no equilibrium could have been robust against higher-order beliefs, and thus every equilibrium theory would have been sensitive to the assumptions about higher-order uncertainty, strengthening Rubinstein's position.

Following Carlsson and Van Damme (1993), global games literature investigate the equilibria in nearby type profiles that are generated by a model that is closed at the first order. At these type profiles, the game is dominant solvable, and the resulting equilibrium action profile converges to the risk-dominant equilibrium as these type profiles approach the coordinated-attack game. In this way, they select the risk-dominant equilibrium. Our result uncovers a difficulty in this methodology: every equilibrium must be discontinuous at the limiting type profile, and an equilibrium selection argument based on continuity is problematic—as there is another path that we could have taken the limit in which we would have selected the other equilibrium. This is despite the fact that the equilibrium outcome is robust against higher-order beliefs in these nearby type profiles themselves (by Proposition 8 in the Appendix.)

On a positive note, Kajii and Morris (1997) show that the risk-dominant equilibrium is robust to incomplete information under common prior assumption.

That is, if the common prior puts sufficiently high weight on the original game, then the incomplete information game has an equilibrium in which the risk-dominant equilibrium is played with high probability according to the common prior. Similar positive results are obtained by others, such as Ui (2001), Morris and Ui (2003). This suggests that, when there is a common prior, it may put low probability on the paths that converge to other equilibria.

7. Application: Cournot Oligopoly

In a linear Cournot duopoly, the game is dominant-solvable, and hence Proposition 8 (in the Appendix) implies that higher-order beliefs have negligible impact on equilibrium. (This has also been shown by Weinstein and Yildiz (2003) and is also implied by a result of Nyarko (1996).) On the other hand, in a linear Cournot duopoly with three or more firms, any production level that is less than or equal to the monopoly production is rationalizable, and hence Proposition 2 implies that a researcher cannot rule out any such output level for a firm no matter how many orders of beliefs he specifies. We will now show a more disturbing fact. Focusing complete-information types, $t^{CK}(\theta)$, for fairly general oligopoly models we will show that when there are sufficiently many firms, any such outcome will be in $S_i^{\infty}\left[B; t^{CK}(\theta)\right]$ for every neighborhood B of $s^*\left(t^{CK}(\theta)\right)$. Therefore, by Proposition 4, even a slight doubt about the model in very high orders will lead a researcher to fail to rule out any outcome that is less than monopoly outcome as a firm's equilibrium output.

General Cournot Model. 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 is positive, and

$\lim_{Q\to\infty} P\left(Q;\bar{\theta}\right) = 0$. Therefore, there exists a unique $\hat{Q}$ such that

$$P\left(\hat{Q};\bar{\theta}\right) = c.$$

(In order to have a nice game, we can impose an upper bound for q, larger than $\hat{Q}$, without affecting the equilibria.) We assume that, on $[0,\hat{Q}]$, $P\left(\cdot;\bar{\theta}\right)$ is continuously twice-differentiable and

$$P' + QP'' < 0.$$

It is well known that, under the assumptions of the model, (i) the profit function, $u\left(q,Q;\bar{\theta}\right) = q\left(P\left(q+Q\right)-c\right)$, is strictly concave in own output q; (ii) the unique best response $q^*\left(Q_{-i}\right)$ 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 $t^{CK}\left(\bar{\theta}\right)$, $s^*\left(t^{CK}\left(\bar{\theta}\right)\right)$, is unique and symmetric (Okuguchi and Suzumura (1971)).

Lemma 4. *In the general Cournot model, for any equilibrium s^*, 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 $P\left(\cdot;\bar{\theta}\right)$.

Proof. Let $\bar{n}$ be any integer greater than $1 + 1/\left|\lambda\right|$, 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 $k > 0$, $S^k\left[B;t^{CK}\left(\bar{\theta}\right)\right] = [\underline{q}^k, \bar{q}^k]^n$, where

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

Define $\underline{Q}^k \equiv \left(n-1\right)\underline{q}^k$, $\bar{Q}^k \equiv \left(n-1\right)\bar{q}^k$, and $Q^* = \left(n-1\right)q^*$, 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, $S^k\left[B; t^{CK}\left(\bar{\theta}\right)\right] = \left[0, q^M\right]^n$ for each $k > \bar{k}$. Therefore, $S^\infty\left[B; t^{CK}\left(\bar{\theta}\right)\right] = \left[0, q^M\right]^n$. □

Together with Proposition 4, this lemma yields the following.

Proposition 7. *In the general Cournot model, assume that* $\Theta = \left[\bar{\theta} - \varepsilon, \bar{\theta} + \varepsilon\right]$ *for arbitrarily small* $\varepsilon > 0$, *and that the best-response function* $q^*(Q_{-i};\theta)$ *is a continuous and strictly increasing function of* θ *at* $\left(Q_{-i}, \bar{\theta}\right)$ *where* $Q_{-i} = (n-1)s_j^*\left(t^{CK}\left(\bar{\theta}\right)\right)$ *is the others' aggregate output in equilibrium. Then,*

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

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

Proof. By (i) above, we have a nice game. By the hypothesis, there exists $B \subset s^*(T)$ as in Lemma 4. Hence, Lemma 4 and Proposition 4 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. □

In Proposition 7, the assumption that $q^*(Q_{-i};\theta)$ is responsive to θ guaranties that θ is a payoff-relevant parameter. 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 the confident researcher and his slightly skeptical friend in the Introduction. The former is confident that it is common knowledge that $\theta = \bar{\theta}$, while the latter 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.

8. Conclusion

It is a common practice in economics to close the model after only specifying the first or second order beliefs, using a (mostly implicit) common knowledge assumption. In this paper, we have investigated the role of this assumption in predictions according to an arbitrary fixed equilibrium. Finding strong lower and upper bounds for the variations with respect to this assumption, we have shown that it is this casually made common knowledge assumption that drives any prediction that we could not have made already by iteratively eliminating strategies that can never be strict best reply. In games like Cournot oligopoly, this implies that no interesting conclusions could have been reached without making a precise common knowledge assumption. In some other games, such as sequential games, our lower bounds are weak, and one may plausibly make sharp predictions using much weaker assumptions. Therefore, it is essential for assuring the reliability of theories to pay special care to closing assumption and justify it at least as much as the other assumptions.

When there are two or more actions that survive our elimination process, there is an inherent unpredictability which cannot be avoided without making an assumption on infinitely many orders of beliefs, as all of these actions are played with probability 1 by some types whose finite-order beliefs agree for arbitrarily high orders. In that case, equilibrium is necessarily discontinuous with respect to higher-order beliefs and in product topology. Moreover, when there are no ties, there is a one-to-one relationship between this sensitivity to higher-order beliefs and sensitivity to higher-order assumptions about players' rationality. It then becomes very difficult in analyzing these situations to justify the common knowledge of rationality as a good approximating assumption.

Appendix A. Proofs and further results

Proof of Lemma 1. Take any i and any a_i. Take any $\gamma \in \Delta(T_{-i})$ with countable support, and let $\mu = \gamma \circ (s^*_{-i})^{-1} \in \Delta(A_{-i})$. Let ν be as in Assumption 1. Define t_i as the type such that $\kappa_{t_i} = \nu \times \gamma$. Notice that $\pi(\cdot|t_i, s^*_{-i}) = \kappa_{t_i} \circ \beta^{-1} = (\nu \times \gamma) \circ \beta^{-1} = \nu \times \left(\gamma \circ (s^*_{-i})^{-1}\right) = \nu \times \mu$. Hence, $s^*(t_i) = BR_i(\pi(\cdot|t_i, s^*_{-i})) = BR_i(\nu \times \mu) = a_i$. □

Proposition 8. *For any equilibrium s^*, any player i, and any t_i and $k \leq \infty$, $A^k_i[s^*, t_i] \subseteq S^k_i[t_i]$.*

Proof. For $k = 0$, the proposition is true by definition. Assume that it is true for some $k - 1 \geq 0$, i.e., for any i, and for for each t_{-i}, $A^{k-1}_{-i}[s^*, t_{-i}] \subseteq S^{k-1}_{-i}[t_{-i}]$. Now, $T_{-i} \subset L \times H$ where $L = \left(\prod_{l=1}^{k-1} \hat{\Delta}(X_{l-1})\right)^{n-1}$ and $H = \left(\prod_{l=k}^{\infty} \hat{\Delta}(X_{l-1})\right)^{n-1}$ are the spaces of lower and higher-order beliefs with generic members l and h, respectively.[6] Now take any $\tilde{t}_i$ with $\tilde{t}^m_i = t^m_i$ for all $m \leq k$. Clearly,

$$\text{marg}_{\Theta \times L} \kappa_{\tilde{t}_i} = \text{marg}_{\Theta \times L} \kappa_{t_i}. \tag{A.1}$$

For each $(\theta, l) \in \Theta \times L$ with $\text{marg}_{\Theta \times L} \kappa_{\tilde{t}_i}(\theta, l) \equiv \kappa_{\tilde{t}_i}(\{(\theta, l)\} \times H) > 0$, let $P(\cdot|\theta, l, s^*_{-i})$ be the the probability distribution on A_{-i} induced by the belief $\tilde{t}_i$ and s^*_{-i} conditional on (θ, l), i.e.,

$$P(a_{-i}|\theta, l, s^*_{-i}) = \kappa_{\tilde{t}_i}\left(\{(\theta, l, h) | s^*_{-i}(l, h) = a_{-i}\}\right) / \kappa_{\tilde{t}_i}(\{(\theta, l)\} \times H) \tag{A.2}$$

at each a_{-i}. Since there are only countably many (θ, l) with $\text{marg}_{\Theta \times L} \kappa_{\tilde{t}_i}(\theta, l) > 0$, there exists $\sigma_{-i} \in \Delta(M_{-i})$ such that $\sigma_{-i}(\{s_{-i} | s_{-i}(\theta, l, h) = a_{-i}\}) = P(a_{-i}|\theta, l, s^*_{-i})$ for each such (θ, l), and $\sigma_{-i}(\{s_{-i} | s_{-i}(\theta, t_{-i}) = \hat{s}_{-i}(\theta, t_{-i})\}) = 1$ otherwise for some fixed $\hat{s}_{-i} \in \hat{S}^{k-1}_{-i}$. Note that according to σ_{-i}, at each (θ, l) in the support and for

[6] $T_{-i} \neq L \times H$ because of the coherency requirement, which has no impact on the rest of the proof. (κ_{t_i} and $\kappa_{\tilde{t}_i}$ put probability 1 on the subset T_{-i}.)

each a_{-i}, the probability that a_{-i} is played is always $P\left(a_{-i}|\theta,l,s^*_{-i}\right)$. By induction hypothesis, we then have

$$\sigma_{-i}\left(\left\{s_{-i}|s_{-i}(\theta,l,h)\in S^{k-1}_{-i}[(l,h)]\right\}\right)=1. \tag{A.3}$$

Moreover, for each $\left(\bar\theta,a_{-i}\right)$ with $t^1_i\left(\bar\theta\right)=\tilde t^1_i\left(\bar\theta\right)>0$,

$$\begin{aligned}
\pi\left((\theta,a_{-i})|t_i,\sigma_{-i}\right) &= \int 1_{\{\theta=\bar\theta\}}P\left(a_{-i}|\theta,l,s^*_{-i}\right)d\kappa_{t_i}(\theta,l,h)\\
&= \int 1_{\{\theta=\bar\theta\}}P\left(a_{-i}|\theta,l,s^*_{-i}\right)\kappa_{t_i}\left(\{(\theta,l)\}\times H\right)d\mathrm{marg}_{\Theta\times L}\kappa_{t_i}(\theta,l)\\
&= \int 1_{\{\theta=\bar\theta\}}\kappa_{\tilde t_i}\left(\left\{(\theta,l,h)|s^*_{-i}(l,h)=a_{-i}\right\}\right)d\mathrm{marg}_{\Theta\times L}d\kappa_{\tilde t_i}(\theta,l)\\
&= \pi\left((\theta,a_{-i})|\tilde t_i,s^*_{-i}\right),
\end{aligned}$$

where $1_{\{\theta=\bar\theta\}}$ is the indicator function for $\left\{\theta=\bar\theta\right\}$; the first equality is obtained by integrating over σ_{-i} appropriately; the second equality is due to the fact that $P\left(a_{-i}|\theta,l,s^*_{-i}\right)$ does not depend on h, and the third equality is due to (A.1) and (A.2). When $t^1_i(\theta)=\tilde t^1_i(\theta)=0$, we trivially have $\pi\left((\theta,a_{-i})|t_i,\sigma_{-i}\right)=\pi\left((\theta,a_{-i})|\tilde t_i,s^*_{-i}\right)=0$. Hence,

$$s^*_i\left(\tilde t_i\right)\in BR_i\left(\pi\left(\cdot|\tilde t_i,s^*_{-i}\right)\right)=BR_i\left(\pi\left(\cdot|t_i,\sigma_{-i}\right)\right)\subset S^k_i[t_i],$$

where the last inclusion is by (A.3) and definition of $S^k_i[t_i]$. □

Proof of Lemma 2. It follows from the following lemma. □

Lemma 5. *For any nice game and for any i, t_i, k, the following are true.*

1. $S^k_i[t_i]=\left[\underline{a}^k,\bar a^k\right]$ *for some* $\underline{a}^k_i,\bar a^k_i\in A_i$*, which depend on* t_i*.*
2. *For each* $a^k_i\in S^k_i[t_i]$*, there exists* $\hat s_{-i}\in\hat S^{k-1}_{-i}$ *such that*

$$BR_i\left(\pi\left(\cdot|t_i,\hat s_{-i}\right)\right)=\{a^k_i\}.$$

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., $S^{k-1}_j[t_j]$ is a closed interval in $A_j=[0,1]$ for each j. This implies that $\hat S^{k-1}_{-i}$ is a closed, convex metric space (with product topology).[7] Moreover, by the Maximum Theorem, $BR_i(\pi(\cdot|t_i, s_{-i}))$ is an upper-semi-continuous function of s_{-i}. But by the strict concavity assumption, $BR_i(\pi(\cdot|t_i, s_{-i}))$ is singleton, and hence the function $\beta_i(\cdot; t_i)$ that maps each $s_{-i} \in \hat{S}^{k-1}_{-i}$ to the unique member of $BR_i(\pi(\cdot|t_i, s_{-i}))$ is continuous. Since $\hat{S}^{k-1}_{-i}$ is compact and convex, this implies that $\beta_i\left(\hat{S}^{k-1}_{-i}; t_i\right)$ is compact and connected, and hence it is convex as it is unidimensional. That is, $\beta_i\left(\hat{S}^{k-1}_{-i}; t_i\right) = \left[\underline{a}^k, \bar{a}^k\right]$ for some $\underline{a}^k_i, \bar{a}^k_i \in A_i$. We **claim** that $\beta_i\left(\hat{S}^{k-1}_{-i}; t_i\right) = S^k_i[t_i]$. This readily proves part 1. Part 2 follows from the definition of $\beta_i\left(\hat{S}^{k-1}_{-i}; t_i\right)$.

Towards proving our **claim**, for each $(\theta, t_{-i}) \in \mathrm{supp}\kappa_{\mathrm{t_i}}$ and for each $s_{-i} \in \hat{S}^{k-1}_{-i}$, define function $U_i(\cdot|\theta, t_{-i}, s_{-i})$ by setting $U_i(a_i|\theta, t_{-i}, s_{-i}) = u_i(\theta, a_i, s_{-i}(\theta, t_{-i}))$ at each a_i. Clearly, U_i is strictly concave, and for each $\sigma_{-i} \in \Delta\left(\hat{S}^{k-1}_{-i}\right)$, the expected payoff of type t_i is

$$\int U_i(a_i|\theta, t_{-i}, s_{-i})\, d\kappa_{t_i}(\theta, t_{-i})\, d\sigma_{-i}(s_{-i}). \tag{A.4}$$

Now, take any $a_i > \bar{a}^k_i$. Then, for each (θ, t_{-i}, s_{-i}), by definition of $\bar{a}^k_i$ and strict concavity of $U_i(\cdot|\theta, t_{-i}, s_{-i})$, we have $U_i(a_i|\theta, t_{-i}, s_{-i}) < U_i\left(\bar{a}^k_i|\theta, t_{-i}, s_{-i}\right)$. It then follows from (A.4) that $\bar{a}^k_i$ yields higher expected payoff than a_i for each $\sigma_{-i} \in \Delta\left(\hat{S}^{k-1}_{-i}\right)$, and thus $a_i \notin S_i[t_i]$. Similarly, $a_i \notin S_i[t_i]$ for each $a_i < \underline{a}^k_i$. □

Proof of Lemma 3. Part 2 follows from the definitions, and Part 3 follows from Part 2 and the fact that $D(A^\infty[s^*, t]) \le D\left(A^k[s^*, t]\right)$ for each k. To prove Part 1, take any $\epsilon > 0$ and any sequence $\epsilon_k > 0$ that converges to 0. For each k, there exists $\tilde{t}[k]$ such that $\tilde{t}^m[k] = t^m$ for each $m \le k$ and $d\left(s^*\left(\tilde{t}[k]\right), s^*(t)\right) \ge D\left(A^k[s^*, t]\right)/2 - \epsilon_k$ so that

$$0 \le D\left(A^k[s^*, t]\right) \le 2d\left(s^*\left(\tilde{t}[k]\right), s^*(t)\right) + 2\epsilon_k. \tag{A.5}$$

[7] Proof: Firstly, $\prod_{(\theta, t_j, j)} S^{k-1}_j[t_j]$ is a compact space by Tychonoff's theorem. But the space of all measurable functions $f : \Theta \times T_{-i} \to \mathbb{R}^{N\setminus\{i\}}$ is closed. Hence, the intersection of these two spaces, namely $\hat{S}^{k-1}_{-i}$, is compact. Convexity of $\hat{S}^{k-1}_{-i}$ follows from the facts that measurability is preserved under point-wise multiplication and addition and that the range is convex.

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 $k \to \infty$. Hence, if s^* is continuous at t, then as $k \to \infty$, $s^*\left(\tilde{t}[k]\right) \to s^*(t)$, and thus the right hand side of (A.5) converges to 0. That is, $D\left(A^k[s^*, t]\right) \to 0$, showing by part 2 that s^* is continuous with respect to higher-order beliefs at t. □

Lemma 6. *For any finite-action game, the following propositions are ordered with logical implication in the following decreasing order.*

(1) σ^* *is continuous with respect to product topology at* t.
(2) σ^* *is continuous with respect to higher-order beliefs at* t.
(3) σ^* *is weakly continuous with respect to higher-order beliefs at* t.
(4) $A^\infty[\sigma^*, t] \subseteq$ *supp*$(\sigma^*(t))$.

Moreover, (2) implies that $|A^\infty[\sigma^*, t]| \leq 1$.

Proof. Since mixed strategies can be considered as pure strategies with values in $\Delta(A_i)$, by Lemma 3.1, (1) implies (2). To show that (2) implies (3), for each $B \subseteq A$, write $\Sigma_B = \{\alpha \in \Delta(A) \,|\, \mathrm{supp}(\alpha) \subseteq B\}$, which is a compact set. Then, for each disjoint B and C, $d_{B,C} = \min\{d(\alpha_B, \alpha_C) \,|\, \alpha_B \in \Sigma_B, \alpha_C \in \Sigma_C\} > 0$. Write $d_{\min}$ for the minimum of $d_{B,C}$ among all disjoint B and C. Clearly, if $d(\alpha, \alpha') < d_{\min}$, then the supports of α and α' have non-empty intersection. But (2) implies that there exists k such that whenever $\tilde{t}^m = t^m$ for each $m \leq k$, $d\left(\sigma^*\left(\tilde{t}\right), \sigma^*(t)\right) < d_{\min}/2$, whence $\mathrm{supp}\left(\sigma^*\left(\tilde{t}\right)\right) \cap \mathrm{supp}(\sigma^*(t)) \neq \emptyset$—hence (3). (3) implies (4) because for each $a \in A^\infty[\sigma^*, t]$ and each k, there exists $\tilde{t}$ such that $\tilde{t}^m = t^m$ for each $m \leq k$ and $\mathrm{supp}\left(\sigma^*\left(\tilde{t}\right)\right) = \{a\}$. The last statement in the lemma also follows from the latter observation because $\sigma^*(t)$ cannot be arbitrarily close to two distinct pure action profiles. □

References

[1] Aumann, R. and A. Brandenburger (1995): "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63, 1161-1180.

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

[3] Battigalli, P. and M. Siniscalchi (2003): "Strong Belief and Forward-Induction Reasoning," *Journal of Economic Theory*, 106, 356-391.

[4] Battigalli, P. and M. Siniscalchi (2003) "Rationalization and Incomplete Information," *Advances in Theoretical Economics*, Vol. 3: No. 1, Article 3.

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

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

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

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

[9] Brandenburger, A. and Keisler (2000): "Epistemic Conditions for Iterated Admissibility," mimeo (www.stern.nyu.edu/~abranden).

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

[11] Feinberg, Y. (2002): "Subjective Reasoning in Dynamic Games," mimeo.

[12] Feinberg, Y. and A. Skrzypacz (2002): "Uncertainty about uncertainty and delay in bargaining," mimeo.

[13] Fudenberg, D., D. Kreps, and D. Levine (1988): "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory*, 44, 354-380.

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

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

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

[17] Kreps, D. and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50-4, 863-894.

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

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

[20] Monderer, D. and D. Samet (1997): "Proximity of information in games with incomplete information," *Mathematics of Operations Research*, 21, 707-725.

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

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

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

[24] Milgrom, P. and J. Roberts (1990): "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58-6, 1255-1277.

[25] Milgrom, P. and R. Weber (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research*, 10, 619-632.

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

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

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

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

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

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

[32] 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.

MIT

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