

IMPACT OF HIGHER-ORDER UNCERTAINTY

JONATHAN WEINSTEIN AND MUHAMET YILDIZ

ABSTRACT. In some games, the impact of higher-order uncertainty is very large, implying that present economic theories may rely critically on the strong common knowledge assumptions they make. Focusing on normal-form games in which the players' action spaces are compact metric spaces, we show that our key condition, called "global stability under uncertainty," implies that the maximum change in equilibrium actions due to changes in players' beliefs at orders higher than k is exponentially decreasing in k . Therefore, given any need for precision, we can approximate equilibrium actions by specifying only finitely many orders of beliefs.

Key words: higher-order uncertainty, stability, incomplete information.

JEL Numbers: C72, C73.

1. INTRODUCTION

Most economic theories are based on equilibrium analysis of models in which the players' types are simply taken as their beliefs about some underlying uncertainty, such as the marginal cost of a firm or the value of an object for a buyer, and rarely include a player's beliefs about the other players' beliefs about the underlying uncertainty. Using such a type structure implicitly assumes that, conditional on the first-order beliefs about some payoff-relevant uncertainty, all of a player's higher-order beliefs are common knowledge.¹

Date: First Version: October, 2002; This version: September, 2006.

We thank Daron Acemoglu, Eddie Dekel, Glenn Ellison, Sergei Izmalkov, Paul Milgrom, Stephen Morris, Marchiano Siniscalchi, Robert Wilson, an associate editor and two referees, and the seminar participants at MIT and UPenn. We are especially grateful to Drew Fudenberg, whose inquiries led us to our result about dominance-solvability.

¹Here we use the standard terminology: a player's first-order beliefs are his beliefs about the underlying uncertainty; his second-order beliefs are his beliefs about the other players' first-order beliefs, and so on.

There is now an extensive literature, however, that emphasizes that in some games higher-order beliefs have as large an impact on equilibrium behavior as lower-order beliefs (see, e.g., Rubinstein, 1989; Morris, 2002). As Rubinstein illustrates, the equilibria of a game in which a particular piece of information is common knowledge can be profoundly different from the equilibria of games in which this information is mutually known only up to some finite order — no matter how many orders we consider. Moreover, Weinstein and Yildiz (2004) show that higher-order beliefs must have a significant impact whenever two or more actions survive iterated eliminations of actions that are never a strict best reply. Most importantly, when the higher-order beliefs have large impact, standard economic theories may be misleading. For example, it is a central result in mechanism design that in traditional type spaces, generically, all the surplus can be extracted from players (Cremer and Mclean, 1988). This property, however, generically fails within the set of common priors on the universal type space (Neeman, 2004; Heifetz and Neeman, 2006). Likewise, the Coase conjecture may fail when we introduce second-order uncertainty (Feinberg and Skrzypacz, 2005).

The large impact described above is also disturbing because it is hard to believe that we would ever know a player's high-order beliefs with any precision. Without such knowledge, we cannot make accurate predictions when the impact of higher-order uncertainty is large. The predictions of our models are then driven by the arbitrary assumptions we make about higher-order beliefs. Moreover, assuming that higher-order beliefs correspond to higher-order reasoning, such a large impact implies that the bounds of rationality are at least as important as the basic incentives. This may necessitate a change of paradigm for analyzing these problems. Therefore, it is of fundamental importance to classify games in which high-order uncertainty has little impact.

In this paper, we provide a set of sufficient conditions under which high-order uncertainty has little impact. Our main sufficient condition is called “global stability under uncertainty.” It states that the variation in each player's best response is always less than the variation in his beliefs about the others' actions

(according to a suitable metric that we will define), multiplied by a constant b that is less than 1. Global stability under uncertainty is closely related to the standard concept of global stability under myopic best-reply (with certainty) and can be checked via a simple second-order condition when action spaces are one-dimensional.

We consider n -person games in which the action spaces are compact metric spaces and there is some payoff-relevant parameter, the source of uncertainty, that lies in a complete, separable metric space. There are many type spaces that model the players' beliefs about the parameter, using different common-knowledge assumptions. For example, in a textbook type space, players' types are their privately observed signals, and the signals and the parameters have a joint distribution that is assumed to be common knowledge. For any such type space, we can compute the infinite hierarchy of higher-order beliefs of each type (using the assumed joint distribution). Unlike textbook type spaces, the universal type space is defined by letting the set of types be explicitly the set of all possible belief hierarchies. When each type in a standard type space corresponds to a different hierarchy, we can identify this space with a subset of the universal type space. Given any equilibrium in the universal type space, the restriction of strategies to this subset corresponds to an equilibrium in the original type space. To perform our analysis, we fix a (Bayesian) Nash equilibrium of the game with the universal type space, effectively fixing an equilibrium for all embedded type spaces simultaneously.²

Let us also fix a player's beliefs up to a certain order k . Our main result states that, assuming global stability, the maximum variation in the player's equilibrium strategy, as we vary all his higher-order beliefs, is at most b^k times a constant. That means that, if we want to determine the equilibrium action within a certain margin of error (e.g., in order to check the validity of a certain theoretical prediction), we only need to specify finitely many orders of beliefs accurately, where the required number of orders k^* is a logarithmic function of

²When there are "redundant" types (corresponding to the same type in the universal type space) that play different actions in equilibrium, such an equilibrium will not be the restriction of any equilibrium on the universal type space.

the desired precision. In particular, the impact of an erroneous common knowledge assumption at orders higher than k^* will be less than the specified bound. This, of course, also applies to the small type spaces embedded in the universal type space. Moreover, we can put bounds on the change in predictions when we use alternative type spaces that differ only in their (implicit) specifications of higher-order beliefs.

Assuming that best responses are always unique, one can show that global stability implies the contraction property of Nyarko (1997), who investigates the convergence to equilibrium in a general abstract model in the same vein as Townsend (1983). Under his contraction property, Nyarko (1997) shows that the unique equilibrium must be continuous with respect to the usual product topology on the universal type space. Using this, one can further show that the maximum impact of higher-order beliefs must eventually vanish as we consider higher and higher orders of beliefs. Using our global stability assumption, we provide a direct, constructive proof with explicit bounds on these impacts. In this way, we are able to show how fast these impacts vanish as the order of beliefs increases. Our constructive proof also sheds light on why higher-order beliefs must be less important under global stability and how this may be reversed if global stability fails.

2. EXAMPLES WITH LINEAR BEST RESPONSES

We will now illustrate how global stability implies diminishing impact of higher-order uncertainty in two-player games with linear best-response functions, such as the linear Cournot duopoly. In such games, global stability simply requires that the absolute value of the slope with respect to the other player's action be less than 1. (See also Morris and Shin (2002) for examples similar to those in this section.)

2.1. Cournot Duopoly. Consider a Cournot duopoly where the inverse-demand function is given by

$$P = \theta - Q$$

where P is the price of a good, $Q = q_1 + q_2$ where q_i is the supply of firm $i \in N = \{1, 2\}$, and θ is an unknown demand parameter. The costs are zero, so that the payoff function of firm i is

$$u_i(\theta, q_1, q_2) = q_i(\theta - q_1 - q_2).$$

All of the above except θ is common knowledge. Beliefs concerning θ will be modeled by an arbitrary type space, whose details will not affect the derivation below. We use t_i to denote a generic type of player i .

A strategy profile (q_1^*, q_2^*) , where $q_i^* : t_i \mapsto q_i^*(t_i)$ specifies firm i 's supply as a function of its type, is a Bayesian Nash equilibrium iff $q_i^*(t_i)$ maximizes the expected payoff of type t_i given the strategy q_j^* of the other firm. That is, equilibrium action $q_i^*(t_i)$ will maximize the expected payoff

$$E_i [q_i(\theta - q_i - q_j^*)] = q_i(E_i[\theta] - q_i - E_i[q_j^*]),$$

so that

$$(2.1) \quad q_i^*(t_i) = \frac{E_i[\theta]}{2} - \frac{1}{2}E_i[q_j^*],$$

where expectation E_i will be a function of t_i , which we suppress in our notation. For example, $E_i[\theta]$ is the type t_i 's expectation of θ , which could also be written as $E[\theta|t_i]$, and is called the first-order expectation. Type t_i also has an expectation of j 's expectations of θ , namely $E[E[\theta|t_j]|t_i]$, which we shorten to $E_i E_j[\theta]$. This is called his second-order expectation. This extends to higher-order expectations. We will show that (2.1) leads to an expression for $q_i^*(t_i)$ in terms of these higher-order expectations. Our next step is to substitute the identical expression for j , $q_j^*(t_j) = E_j[\theta]/2 - E_j[q_i^*]/2$, into (2.1), obtaining

$$(2.2) \quad q_i^*(t_i) = \frac{E_i[\theta]}{2} - \frac{E_i E_j[\theta]}{4} + \frac{1}{4}E_i E_j[q_i^*].$$

Proceeding iteratively, we obtain

$$q_i^*(t_i) = \frac{E_i[\theta]}{2} - \frac{E_i E_j[\theta]}{4} + \dots + \frac{1}{2^k} \underbrace{E_i E_j E_i \dots E_i}_{k \text{ times}}[\theta] - \frac{1}{2^k} \underbrace{E_i E_j E_i \dots E_i}_{k \text{ times}}[q_j^*]$$

when k is odd; the last term is $+E_i E_j E_i \dots E_j[q_i^*]/2^k$ when k is even. In equilibrium, each firm's supply will always be in $[0, 1]$; hence the absolute value of the

last term is at most $1/2^k$. That is, if we fix the expectations up to k th order, we know the equilibrium action $q^*(t_i)$ up to an error of at most $1/2^k$.

This also implies that we can write the equilibrium action as a convergent series

$$(2.3) \quad q_i^*(t_i) = \frac{E_i[\theta]}{2} - \frac{E_i E_j[\theta]}{4} + \frac{E_i E_j E_i[\theta]}{8} - \frac{E_i E_j E_i E_j[\theta]}{16} + \dots$$

where the coefficient of the k th term is $1/2^k$. The significance of this formula is that the coefficients of expectations decrease exponentially as we go to higher-order expectations.

In this example the slope of the best-reply function with respect to the other player's action was $-1/2$, so that its absolute value was less than 1, satisfying our global stability condition. We now discuss what might happen in a similar example where the slope may be greater than 1, explicitly deriving the higher-order expectations and equilibria within a specific type space.

2.2. A Canonical Game with a Traditional Type Space. Consider a two-player game with a payoff parameter θ on which the players have common prior $\theta \sim N(0, 1)$, where $N(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 . Each player i gets a private signal $x_i = \theta + \varepsilon_i$, where $\varepsilon_i \sim N(0, (1-v)/v)$ for some $v \in (0, 1)$, with θ , ε_1 , and ε_2 all independent. For some $b \geq 0$, let the utility function of a player i be $-(a_i - \theta - ba_j)^2$, where $a_i \in \mathbb{R}$ and $a_j \in \mathbb{R}$ are the actions of each player, so that his best reply is $E[\theta + ba_j | x_i]$. The above is all common knowledge.

Whenever $bv \neq 1$, one can directly compute a Bayesian Nash equilibrium s^* with

$$(2.4) \quad s_i^*(x_i) = \frac{vx_i}{1 - bv}.$$

As in the Cournot example, one can also write any equilibrium s_i^* as a series of higher order expectation:

$$s_i^* = E_i[\theta] + bE_i E_j[\theta] + \dots + b^{k-1} \underbrace{E_i E_j E_i \dots E_i}_{k \text{ times}}[\theta] + b^{k-1} \underbrace{E_i E_j E_i \dots E_i}_{k \text{ times}}[s_j^*]$$

(for k odd.) As a function of x_i , the first-order expectation is $E_i[\theta] = vx_i$, the second-order expectation is $E_i E_j[\theta] = v^2 x_i$, and in general the k th-order expectation is $E_i E_j E_i \dots E_j[\theta] = v^k x_i$. Hence,

$$s_i^*(x_i) = vx_i + bv^2 x_i + \dots + b^{k-1} v^k x_i + b^{k-1} \underbrace{E_i E_j E_i \dots}_{k \text{ times}} [s_j^*].$$

Clearly, when $b < 1$, the contribution of each term $v^k x_i$ decreases exponentially; the resulting series converges to the formula in (2.4). On the other hand, this contribution grows exponentially when $b > 1$, and in the case that $bv > 1$, higher-order terms increase exponentially, yielding a divergent series.³ When $b > 1$, if the model specifies the higher-order beliefs incorrectly, then the model's predictions will be dramatically different than those of the accurate model.

3. MODEL

In the previous section, the linearity of the best-response function greatly simplified the analysis, allowing us to write the equilibrium strategies as explicit functions of higher-order expectations. In general, the best-response of a player depends on his entire belief about the parameters and the other players' actions, not just their expected values. Consequently, equilibrium strategies will depend on all features of a player's higher-order beliefs, rather than just higher-order expectations, and cannot be expressed so simply. Therefore, to analyze the general case, we will introduce higher-order beliefs explicitly. In this section, we will consider a Bayesian game where the type space is the entire universal type space, in which a type is an infinite hierarchy of beliefs. As explained in the introduction, this space contains most textbook type spaces.

We consider a game among players $N = \{1, 2, \dots, n\}$. The source of underlying uncertainty is a payoff-relevant parameter $\theta \in \Theta$, where Θ is a compact metric space with metric d . Each player i has action space A_i , which is a compact metric space with metric d_i , and utility function $u_i : \Theta \times A \rightarrow \mathbb{R}$

³When $bv < 1 < b$, we have a convergent series yielding the seemingly intuitive formula in (2.4), despite the fact that marginal contributions of higher-order expectations increase exponentially.

where $A = \prod_i A_i$. We let $A_{-i} = \prod_{j \neq i} A_j$ with generic member a_{-i} and use the convention that $(\hat{a}_i, a_{-i}) = (a_1, \dots, a_{i-1}, \hat{a}_i, a_{i+1}, \dots, a_n)$.

We now define the players' hierarchy of beliefs about the underlying parameter θ , using the usual construction of the universal type space by Brandenburger and Dekel (1993), a variant of an earlier construction by Mertens and Zamir (1985). We will define our types using the auxiliary sequence $\{X_k\}$ of sets defined inductively by $X_0 = \Theta$ and $X_k = [\Delta(X_{k-1})]^n \times X_{k-1}$ for each $k > 0$.⁴ A player i 's first-order beliefs 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

$$t_i = (t_i^1, t_i^2, \dots)$$

of a player i is a member of $\prod_{k=1}^{\infty} \Delta(X_{k-1})$. Since a player's k th-order beliefs contain information about his lower order beliefs, we need the usual coherence requirements. We write $T = T_1 \times \dots \times T_n$ for the subset of $(\prod_{k=1}^{\infty} \Delta(X_{k-1}))^n$ in which it is common knowledge that the players' beliefs are coherent, where coherence means that the players know their own beliefs and their marginals from different orders agree. We will use the variables $t = (t_1, \dots, t_n)$, $\tilde{t} = (\tilde{t}_1, \dots, \tilde{t}_n) \in T$ as generic type profiles, $t_i, \tilde{t}_i \in T_i$ as generic types, and $t_{-i}, \tilde{t}_{-i} \in T_{-i} \equiv \prod_{j \neq i} T_j$ as generic profiles of types for players other than i .

A strategy of a player i is a measurable mapping $s_i : T_i \rightarrow A_i$, that determines which action $s_i(t_i)$ he would choose given his type t_i . We fix a Bayesian Nash equilibrium $s^* = (s_1^*, s_2^*, \dots, s_n^*)$, which must be such that $s_i^*(t_i)$ maximizes the expected value of $u_i(\theta, a_i, s_{-i}^*(t_{-i}))$ with respect to the belief of t_i on (θ, t_{-i}) at each t_i and for each i , where $s_{-i}^*(t_{-i}) = (s_1^*(t_1), \dots, s_{i-1}^*(t_{i-1}), s_{i+1}^*(t_{i+1}), \dots, s_n^*(t_n))$. Existence and uniqueness of equilibrium for games with unique best responses are guaranteed by global stability, which we define in the next section.

⁴For any metric space X , $\Delta(X)$ denotes the set of probability distributions over the Borel σ -field, endowed with the weak topology. Also, we will always use the product topology on product sets.

Example 1. *Equation 2.3 gives an equilibrium for the Cournot duopoly game in the universal type space, provided that we set action spaces $A_i = [0, 1]$ to ensure that the series converges.*

Example 2. *Consider a second-price auction with private valuations. In order to fit this example into our abstract framework, we set $\Theta = \Theta_1 \times \dots \times \Theta_n$, where θ_i is player i 's valuation, and restrict ourselves to the subset of the universal type space in which each player i assigns probability 1 to some value of θ_i (i.e. he knows his own valuation.). Such types have a weakly dominant action of bidding their own valuation. If we restrict again to the largest belief-closed set of types contained in this subset, these actions give a Bayesian equilibrium on a large belief-closed subset of the universal type space, namely the subset on which it is common knowledge that players know their own valuations. It is necessary to restrict attention to this subset to model private valuations. Our results apply to any subset of the universal type space.*

4. GLOBAL STABILITY UNDER UNCERTAINTY

We are now ready to present our sufficient condition for the diminishing impact of higher order uncertainty: global stability. Under certainty, global stability can be defined by the condition that for any given change in the other players' actions, the resulting change in a player's best response is smaller in magnitude.⁵ We will now extend this notion to the best-response function under uncertainty.

Best Responses. The optimal action of a player i depends on his beliefs about the payoff-relevant variables he does not control: θ and a_{-i} . In a Bayesian Nash equilibrium s , such beliefs are derived from the joint belief of type t_i about θ and the opponents' types t_{-i} , which correspond to actions $s_{-i}(t_{-i})$. For any i and probability distribution $\pi \in \Delta(\Theta \times A_{-i})$, we write $BR_i(\pi) \in A_i$ for the best response, or optimal action, of player i when his beliefs about the

⁵The usual definition may appear different. For instance, in two player games we only need that the product of maximum variations is less than 1. Of course, under this condition, we could rescale our metrics on each strategy space so that our definition is also satisfied.

underlying uncertainty θ and the other players' actions a_{-i} are represented by π . In general, however, there may be multiple best responses, and this allows that in a Bayesian Nash equilibrium two different types may have the same derived beliefs about (θ, a_{-i}) and yet choose different best responses. To simplify the analysis we will therefore assume that in the chosen equilibrium s^* each player uses the same best responses whenever he has the same beliefs about (θ, a_{-i}) , and we let $BR_i(\pi)$ denote a selection from the best response correspondence that agrees with s^* .

Towards stating our global stability condition, for each $i \in N$, define the metric d_{-i} on A_{-i} by

$$d_{-i}(a_{-i}, a'_{-i}) = \max_{j \neq i} d_j(a_j, a'_j).$$

Define the metric \bar{d}_{-i} on $\Delta(\Theta \times A_{-i})$ by

$$(4.1) \quad \bar{d}_{-i}(\pi, \pi') \equiv \inf_{\nu \in \chi_{\pi, \pi'}} \int d_{-i}(a_{-i}, a'_{-i}) d\nu(\theta, a_{-i}, a'_{-i})$$

where

$$(4.2) \quad \chi_{\pi, \pi'} = \{\nu \in \Delta(\Theta \times A_{-i} \times A_{-i}) : \text{marg}_{12}\nu = \pi, \text{marg}_{13}\nu = \pi'\},$$

where marg_{12} denotes the marginal distribution on the first and the second sets, etc. This metric has the crucial property of preserving distances between actions when the actions are embedded in the space $\Delta(\Theta \times A_{-i})$ of beliefs about these actions as point masses, allowing us to sensibly compare variations in actions and the beliefs about them. We are ready to state our central sufficient condition.

Definition 1 (Global Stability under Uncertainty). *We say that global stability under uncertainty holds iff there exists some $b \in [0, 1)$ such that, given any $i \in N$ and any $\pi, \pi' \in \Delta(\Theta \times A_{-i})$ with $\text{marg}_{\Theta}\pi = \text{marg}_{\Theta}\pi'$, we have*

$$(4.3) \quad d_i(BR_i(\pi), BR_i(\pi')) \leq b\bar{d}_{-i}(\pi, \pi').$$

The required condition for global stability is the standard condition for Lipschitz continuity (of each BR_i with respect to the metric \bar{d}_{-i} on $\Delta(\Theta \times A_{-i})$) with the additional requirement that the constant b , which can be thought of as an upper bound on the absolute value of the slope, be less than 1. Our

condition states that when we vary a player's beliefs about the other players' actions, the change in his best response is less than the change in his beliefs. Assuming unique best responses, global stability is a property of the players' best responses (or their utility functions), rather than of an equilibrium.

We will use global stability to bound the changes in a player's equilibrium actions due to changes in his higher-order beliefs. The latter may involve changes in many aspects of his beliefs, such as those regarding correlations between other players' beliefs and the parameter θ , leading to a wide range of (possibly complicated) changes in the joint distribution of (θ, a_{-i}) . Our definition of the metric \bar{d}_{-i} captures the nature of this wide range of possible changes in beliefs in precisely the way necessary to establish our bounds, as we will see in the proof of the main result.

Our next result states that global stability implies that our game is dominance-solvable. Notice that in our game, a strategy of a player i is a function from his entire type space to A_i . Our result refers to *interim* rationalizability, which corresponds to the iterative elimination of action-type pairs where the action is never a weak best reply for the type.⁶

Proposition 1. *Assume that each player has single-valued best response correspondence, and assume global stability under uncertainty. Then, there exists only one rationalizable strategy profile s^* , which is the unique equilibrium.*

For a constructive, direct proof, see our working paper (Weinstein and Yildiz, 2003). An indirect proof is as follows. Firstly, our main theorem shows that the impact of higher-order beliefs vanishes exponentially. But since there are no weak-best responses, the result of Weinstein and Yildiz (2004) then shows that for each type there is at most one rationalizable action. On the other hand,

⁶This differs from ex ante rationalizability, in that we allow different types to have different conjectures about the relationship between other players' types and actions. Interim rationalizability is the appropriate notion for us since we model incomplete information directly, and bypass the ex ante stage completely. Under priors that put positive probability on all of a player's own types, ex ante rationalizability refines interim rationalizability, and hence our results also apply to ex ante rationalizability.

since global stability implies the contraction property of Nyarko (1997), there exists an equilibrium s^* , which will never be eliminated and hence will be the unique rationalizable strategy profile.

The requirement that there is always a unique best response is not superfluous. For example, in the second-price auction with private values, there are multiple best replies, and therefore Proposition 1 does not apply. Indeed, the game has multiple Bayesian Nash equilibria, and hence multiple rationalizable strategies. Yet, arguably the game has only one salient equilibrium, the one that selects the weakly dominant bid for each type. This selection satisfies our global stability condition. Our main result, stated in the next section, applies to all equilibria that yield a selection of the best-response correspondence satisfying global stability.

5. MAXIMUM IMPACT OF HIGHER-ORDER BELIEFS

As argued in the introduction, modelers would prefer not to have to specify the players' higher-order beliefs precisely, and standard economic theories rely on a general common knowledge assumption for all high-order beliefs. It is then very important to determine the accuracy with which we can predict a player's equilibrium behavior if we only know his beliefs up to k th order and have no knowledge of his beliefs at higher orders. We now present our main result, which states that global stability implies at least a certain level of accuracy, which improves exponentially with k .

An intuitive argument for our result is roughly as follows. Suppose that we change a player i 's beliefs only at order $k > 1$ and higher. Since his first $k - 1$ orders of beliefs are fixed, this corresponds to a change in other players' beliefs, as assessed by i , only at orders $k - 1$ and higher. If the change in these players' equilibrium actions due to such changes is bounded by Δ , then the change in equilibrium beliefs of player i about (θ, a_{-i}) is also bounded by Δ according to the metric \bar{d}_{-i} . Then, global stability tells us that the change in the best reply of player i , and hence the change in his equilibrium actions, is bounded by $b\Delta$.

Thus, the maximum impact of higher-order beliefs decreases by a factor of b with each order.

Towards stating our result formally, let

$$(5.1) \quad D_{s^*} = \max_{i \in N} \sup_{t_i, t'_i \in T_i} d_i(s_i^*(t_i), s_i^*(t'_i)) \in \mathbb{R}$$

be the maximum distance between any two equilibrium actions of a player, as we vary his type. We will use this general bound on equilibrium actions to find a bound on a player's actions as we vary his higher-order beliefs, keeping his lower-order beliefs fixed. In our result, one can replace D_{s^*} with any known bound, such as the diameter of the action space, if needed.

Theorem 1. *Assume global stability under uncertainty. Then, there exists $b \in [0, 1)$ such that for any $i \in N$, any $t_i, \tilde{t}_i \in T_i$ and $k \geq 0$, if $t_i^l = \tilde{t}_i^l$ for all $l \leq k$, then*

$$(5.2) \quad d_i(s_i^*(t_i), s_i^*(\tilde{t}_i)) \leq b^k D_{s^*},$$

where D_{s^*} is as defined in (5.1).

Notice that our result relies only on global stability and boundedness of the action space. Under these two assumptions we reach the conclusion that, if we know the beliefs up to a certain order k , we can know the equilibrium action with a maximum error that is an exponentially decreasing function of k , bounding the maximum impact all the higher-order beliefs can have on equilibrium. This is a contribution to the goal set out by Wilson (1987) of “successive reductions in the base of common knowledge required to conduct useful analyses of practical problems.”

In certain cases, a modeler might want to predict the equilibrium behavior within a certain margin of error. For example, checking the validity of certain qualitative predictions of his theories may only require the knowledge of equilibrium actions within a certain margin of error. Theorem 1 tells us how many orders of beliefs he needs to specify. It implies that, given any $\epsilon > 0$ and any

$t \in T$, if we know t up to the order

$$(5.3) \quad k \geq \frac{\log(D_{s^*}) - \log(\epsilon)}{-\log(b)},$$

then we can compute the equilibrium actions up to a maximum error of ϵ . Notice that the expression on the right-hand side is increasing in b and decreasing in ϵ . (We must caution that what we have here is only a theoretical exercise; the order k above may be too high to be useful in practice.)

Our theorem is valid for any two types in the universal type space, which contains all traditional (“small”) type spaces that do not contain two different types with identical belief hierarchies. Hence, our result is also valid for any such traditional type space. (In a small type space, one may be able to find a tighter bound.) More importantly, if there are two such type spaces \hat{T} and \tilde{T} with types $\hat{t}_i \in \hat{T}_i$ and $\tilde{t}_i \in \tilde{T}_i$ whose first k orders of beliefs are identical, then our theorem implies that the equilibrium actions of these types can differ by at most $b^k D_{s^*}$.

Considering a class of games which satisfy our global stability condition, Morris (2002) shows that equilibrium is uniformly continuous only with respect to the topology of uniform convergence—and not with respect to the usual product topology that we use. By requiring uniform continuity, Morris (2002) focuses on worst-case scenarios, while we analyze the continuity of equilibrium within a fixed game. This is the source of the divergence in our conclusions. Under global stability, there is a unique equilibrium. Clearly, in the general case where global stability does not hold, multiplicity of equilibria is a significant phenomenon. In particular, Brandenburger and Dekel (1987) show that one can obtain in equilibrium every iteratively undominated action, by constructing type spaces with redundant types. When global stability holds, however, as we showed, there will be a unique iteratively undominated action. In a complementary approach, Battigalli (1999) and Battigalli and Siniscalchi (2003) introduce a notion of Δ -rationalizability that allows direct analysis of equilibrium implications of explicit common-knowledge assumptions on first-order beliefs and sequential rationality.

6. SUFFICIENT CONDITIONS FOR GLOBAL STABILITY

Next, we present two sets of sufficient conditions for global stability. We omit the proofs for space considerations; they can be found in our working paper. The first set of conditions establishes that global stability under uncertainty is closely linked to the usual global stability in a broad class of games:

Proposition 2. *Assume that the best-response function of each player i takes the form of*

$$(6.1) \quad BR_i(\pi) = f_i(E[g_i(\theta, a_{-i})])$$

where expectation is taken with respect to $\pi \in \Delta(\Theta \times A_{-i})$; $f_i : X \rightarrow A_i$ and $g_i : \Theta \times A_{-i} \rightarrow X$ are two Lipschitz continuous functions defined through some Banach space (X, d_X) with constants α_i and β_i ,⁷ and $b_i \equiv \alpha_i \beta_i < 1$. Then, global stability is satisfied (with $b = \max_{i \in N} b_i$).

The next example presents a general class of games where the above conditions can be easily checked.

Example 3. *For each $i \in N$, take $A_i = [\underline{x}, \bar{x}]$ for some $\underline{x}, \bar{x} \in \mathbb{R}$ and*

$$u_i(\theta, a_i, a_{-i}) = \phi_i(a_i) g_i(\theta, a_{-i}) - c_i(a_i),$$

where $g_i : \Theta \times A_{-i} \rightarrow \mathbb{R}$ is a continuously differentiable function with $|\partial g_i / \partial a_j| < \beta_i$ for each $j \neq i$ and for some $\beta_i \in \mathbb{R}$, and ϕ_i and c_i are twice continuously differentiable functions with $\phi_i' > 0$, $\phi_i'' < 0$, $c_i' > 0$, and $c_i'' \geq 0$. Check that

$$BR_i(\pi) = f_i(E[g_i(\theta, a_{-i})])$$

where $f_i(y)$ is the unique solution to the first order condition $c'(x) / \phi'(x) = y$ when the constraints $\underline{x}_i \leq f_i(y) \leq \bar{x}$ are not binding. Then, by Proposition 2, global stability is satisfied whenever $b \equiv \max_{i \in N, x \in [\underline{x}, \bar{x}]} \beta_i / (c'(x) / \phi'(x))' < 1$.

Our next result presents a simple sufficient condition for global stability in terms of second derivatives of the utility functions.

⁷That is, $d_i(f_i(x), f_i(x')) \leq \alpha_i d_X(x, x')$ and $d_X(g_i(\theta, a_{-i}), g_i(\theta, a'_{-i})) \leq \beta_i d_{-i}(a_{-i}, a'_{-i})$.

Proposition 3. *For each i , assume $A_i \subset \mathbb{R}$, $u_i(\theta, \cdot)$ is twice-continuously differentiable, $u_i(\theta, \cdot, a_{-i})$ is strictly concave, $\partial^2 u_i / \partial a_i^2$ is bounded away from zero, and*

$$(6.2) \quad b_i \equiv \sum_{j \neq i} \frac{\max_{a, \theta} |\partial^2 u_i(\theta, a) / \partial a_i \partial a_j|}{\min_{a, \theta} |\partial^2 u_i(\theta, a) / \partial a_i^2|} < 1.$$

Then, we have global stability under uncertainty whenever (i) $BR_i(\pi)$ is in the interior of A_i for all π , or (ii) A_i is convex.

Notice that, if we drop the max and min in the summand in (6.2), each term becomes the absolute value of the slope of a best-response function with respect to another player's action. Hence, if we instead had a single maximum taken outside the sum, the condition becomes the usual global stability condition under certainty. Under incomplete information, where many different beliefs are possible, we need the stronger condition above to assure global stability. When the second-order derivatives do not vary dramatically, the two conditions will be close.

7. CONCLUSION

Standard economic theories are mostly based on equilibrium analysis of models in which, conditional on only a few low orders of uncertainty, all higher-order beliefs are assumed to be common knowledge. We know, however, that the equilibrium predictions in game theoretical models can be greatly impacted by the assumptions that the model makes about higher-order beliefs. It is disturbing that predictions would be so sensitive to the modeler's precise assumptions about higher-order beliefs, because common sense tells us that it is unlikely that real people consider very high orders of beliefs, so that the modeler's assumptions are inevitably arbitrary and unverifiable. In this paper we presented a sufficient condition, called global stability under uncertainty, which guarantees that the impact of higher-order uncertainty is low. Using the universal type space, which describes every coherent set of beliefs the players can entertain, we have shown under this assumption that if we specify the players' beliefs up to some order k , we will know their equilibrium behavior within a bound that

decreases exponentially in k . To see the implications of our result for game theoretical models, consider an arbitrary model in which our global stability condition is satisfied. If a theoretical prediction in this model would remain true given a margin ϵ of error in equilibrium actions, then the researcher can validate his theory by verifying that first $k(\epsilon)$ orders of beliefs are as specified in his model, where $k(\epsilon)$ grows only logarithmically in ϵ .

APPENDIX A. PROOF OF THEOREM 1

In our proof, we use the following technical lemma, whose proof can be found in our working paper.

Lemma 1. *Let (X, Σ_X) , (Y, Σ_Y) , (Z, Σ_Z) be separable standard Borel spaces, and endow $X \times Y$, $Y \times Z$, $X \times Z$, and $X \times Y \times Z$ with the σ -algebras generated by the corresponding product topologies. Let probability measures P and P' on $X \times Y$ and $X \times Z$, respectively, be such that $\text{marg}_X P = \text{marg}_X P'$. Then, there exists a probability measure \tilde{P} on $X \times Y \times Z$ such that $\text{marg}_{X \times Y} \tilde{P} = P$ and $\text{marg}_{X \times Z} \tilde{P} = P'$.*

Define $\Omega = \Theta \times T$ to be the universal state space. This is the subset of the larger space $\bar{\Omega} = \Theta \times (\prod_{k=1}^{\infty} \Delta(X_{k-1}))^n$ in which coherency is common knowledge. By Brandenburger and Dekel (1993), $\bar{\Omega}$ is a Polish space, yielding a standard separable Borel space, and for every $t = (t_1, \dots, t_n) \in T$ and for every $i \in N$, there exists a probability distribution $\kappa_{t_i} \in \Delta(\bar{\Omega})$ such that

$$(A.1) \quad \text{marg}_{X_{k-1}} \kappa_{t_i} = t_i^k \quad (\forall k),$$

and $\kappa_{t_i}(\Omega) = 1$. Let

$$\beta : (\theta, t) \mapsto (\theta, s_{-i}^*(t_{-i})),$$

and write

$$\pi_{t_i} = \kappa_{t_i} \circ \beta^{-1} \in \Delta(\Theta \times A_{-i})$$

for the joint distribution of the underlying uncertainty and the other players' actions induced by t_i . Notice that $s_i^*(t_i) = BR_i(\pi_{t_i})$.

We will use induction on k . For $k = 0$, (5.2) is true by definition. Fix any $k > 0$, and assume that the result is true for $k - 1$ (for all $j \in N$). Take any t_i and \tilde{t}_i as in

the hypothesis. We have

$$\begin{aligned}
d_i(s_i^*(t_i), s_i^*(\tilde{t}_i)) &= d_i(BR_i(\pi_{t_i}), BR_i(\pi_{\tilde{t}_i})) \\
&\leq b\bar{d}_{-i}(\pi_{t_i}, \pi_{\tilde{t}_i}) \\
\text{(A.2)} \quad &\equiv b \inf_{\nu \in \mathcal{X}_{\pi_{t_i}, \pi_{\tilde{t}_i}}} \int d_{-i}(a_{-i}, a'_{-i}) d\nu(\theta, a_{-i}, a'_{-i}),
\end{aligned}$$

where the inequality is due to global stability and $\mathcal{X}_{\pi_{t_i}, \pi_{\tilde{t}_i}}$ is defined by (4.2). The rest of the proof is devoted to constructing a $\nu \in \mathcal{X}_{\pi_{t_i}, \pi_{\tilde{t}_i}}$ such that, under the induction hypothesis,

$$\text{(A.3)} \quad \int d_{-i}(a_{-i}, a'_{-i}) d\nu(\theta, a_{-i}, a'_{-i}) \leq b^{k-1} D_{s^*}.$$

Combining (A.2) and (A.3), we obtain (5.2).

We will decompose $\bar{\Omega}$ as $\bar{\Omega} = \Theta \times L \times H$ where

$$\text{(A.4)} \quad L = \prod_{l=1}^{k-1} (\Delta(X_{l-1}))^n \quad \text{and} \quad H = \prod_{l=k}^{\infty} (\Delta(X_{l-1}))^n$$

are the spaces of lower and higher-order beliefs. For $k = 1$, we use the convention that L is a singleton, and $l \in L$ can simply be ignored in the following analysis for that case. Note that $X_{k-1} = \Theta \times L$.

By (A.1), we have probability distributions κ_{t_i} and $\kappa_{\tilde{t}_i}$ on $\bar{\Omega}$ such that

$$\text{marg}_{X_{k-1}} \kappa_{t_i} = t_i^k = \tilde{t}_i^k = \text{marg}_{X_{k-1}} \kappa_{\tilde{t}_i},$$

where the second equality is by our hypothesis. Since we have separable standard Borel spaces, by Lemma 1, there exists $\sigma \in \Delta(X_{k-1} \times H \times H)$ such that the marginals of σ on the cross product of X_{k-1} with the first and second copies of H are

$$\text{marg}_{12} \sigma = \kappa_{t_i} \quad \text{and} \quad \text{marg}_{13} \sigma = \kappa_{\tilde{t}_i},$$

respectively.

Now, consider $\nu = \sigma \circ \gamma^{-1} \in \Delta(\Theta \times A_{-i} \times A_{-i})$ where

$$\text{(A.5)} \quad \gamma : (\theta, l, h_1, h_2) \mapsto (\theta, s_{-i}^*(l, h_1), s_{-i}^*(l, h_2)).$$

Notice that the marginal of ν on the first copy of $\Theta \times A_{-i}$ is

$$\text{marg}_{12} \nu = \text{marg}_{12} (\sigma \circ \gamma^{-1}) = (\text{marg}_{12} \sigma) \circ \beta^{-1} = \kappa_{t_i} \circ \beta^{-1} = \pi_{t_i},$$

and similarly $\text{marg}_{13} \nu = \pi_{\tilde{t}_i}$. Therefore, by definition, $\nu \in \mathcal{X}_{\pi_{t_i}, \pi_{\tilde{t}_i}}$.

We now prove (A.3). Take any $(\theta, a_{-i}, a'_{-i}) \in \gamma(X_{k-1} \times H \times H)$. By (A.5), $a_{-i} = s_{-i}^*(\hat{t}_{-i})$ and $a'_{-i} = s_{-i}^*(\bar{t}_{-i})$ for some type profiles $\hat{t} = (l, h_1)$ and $\bar{t} = (l, h_2)$, which agree up to the order $k - 1$. Then, by the induction hypothesis,

$$(A.6) \quad d_{-i}(a_{-i}, a'_{-i}) \leq b^{k-1} D_{s^*}.$$

Since $\text{supp } \nu \subset \gamma(X_{k-1} \times H \times H)$ (by construction), (A.6) implies (A.3).

References

- Battigalli, P., 1999. Rationalizability in games of incomplete information. mimeo.
- Battigalli, P., Siniscalchi, M., 2003. Rationalization and incomplete information. *Advances in Theoretical Economics*. Vol. 3: No. 1, Article 3.
- Brandenburger, A., Dekel, E., 1987. Rationalizability and correlated equilibria. *Econometrica*. 55, 1391-1402.
- Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *J. Econ. Theory*. 59, 189-198.
- Cr emer, J., McLean, R., 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*. 56, 1247-1257.
- Feinberg, Y., Skrzypacz, A., 2005. Uncertainty about uncertainty and delay in bargaining. *Econometrica*. 73, 69-91.
- Heifetz, A., Neeman, Z., 2006. On the generic (im)possibility of full surplus extraction in mechanism design. *Econometrica*. 74, 213-233.
- Mertens, J., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. *Int. J. Game Theory*. 10, 619-632.
- Morris, S., 2002. Typical Types. mimeo.
- Morris, S., Shin, H., 2002. Social value of public information. *Amer. Econ. Rev.* 92, 1521-1534.
- Neeman, Z., 2004. The relevance of private information in mechanism design. *J. Econ. Theory*. 117, 55-77.
- Nyarko, Y., 1997. Convergence in economic models with Bayesian hierarchies of beliefs. *J. Econ. Theory*. 74, 266-296.

Rubinstein, A., 1989. The electronic mail game: Strategic behavior under ‘almost common knowledge’. *American Economic Review*. 79, 385-391.

Townsend, R., 1983. Forecasting the forecasts of others. *J. Polit. Economy*. 9, 546-88.

Weinstein, J., Yildiz, M., 2003. Impact of higher-order uncertainty. MIT Department of Economics Working Paper 03-14.

Weinstein, J., Yildiz, M., 2004. Finite-order implications of any equilibrium. MIT Economics Working Paper No. 04-06.

Wilson, R., 1987. Game-Theoretic analyses of trading processes. In: Bewley, T. (Ed.) *Advances in Economic Theory: Fifth World Congress*, Cambridge UK: Cambridge University Press, 33-70.

NORTHWESTERN UNIVERSITY AND MIT