

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 be misleading as these theories assume common knowledge of the type structure after specifying the first or the second orders of beliefs. Focusing on normal-form games in which the players' strategy spaces are compact metric spaces, we show that our key condition, called "global stability under uncertainty," implies a variety of results to the effect that the impact of higher-order uncertainty is small. Our central result states that, under global stability, the maximum change in equilibrium strategies 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 strategies by specifying only finitely many orders of beliefs.

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

JEL Numbers: C72, C73.

1. INTRODUCTION

Most economic theories are based on equilibrium analysis of models in which the players' types (following Harsanyi (1967)) 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

Date: First Version: October, 2002; This version: March, 2003.

We thank Daron Acemoglu, Glenn Ellison, Sergei Izmalkov, Paul Milgrom, Robert Wilson and the seminar participants at MIT and UPenn.

knowledge.¹ Even the literature on global games (Carlsson and van Damme (1993)) and on forecasting others' forecasts (Townsend (1983)) makes this assumption (in a finite-dimensional space of payoff uncertainty.)²

There is now an extensive literature, however, that emphasizes that in some games higher-order uncertainty has as large an impact on equilibrium behavior as lower-order uncertainty (see Rubinstein (1989), Kajii and Morris (1998) and Morris (2002)). As Rubinstein (1989) 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. Most importantly, when the higher-order beliefs have large impact, the present economic theories may be misleading.³ This large impact 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. 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 would 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 broad 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 the embedding metric defined later), multiplied by a constant b that is less than 1. Under certain continuity assumptions, we show that

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

²For an illustration of how a model with such an assumption can be deceptive regarding the impact of higher-order uncertainty, see Section 2.3.

³For example, the Coase conjecture may fail when we introduce second-order uncertainty as shown by Feinberg and Skrzypacz (2002).

global stability under uncertainty is closely related to the standard concept of global stability of best-response correspondence (under certainty). For games with one-dimensional strategy spaces, we further provide a simple second-order condition that guarantees global stability under uncertainty.

We consider finite-person games in which the strategy spaces are compact metric spaces and there is some payoff-relevant source of uncertainty that comes from a complete, separable metric space. We work in universal type space, where the players' types are their entire hierarchy of beliefs about the underlying uncertainty, allowing players to entertain any coherent set of beliefs. We fix a (Bayesian) Nash equilibrium of this game. Note that, since every type space can be embedded in universal type space, this corresponds to fixing an equilibrium for all type spaces simultaneously. 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 behavior 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, where the required number of orders k^* is a logarithmic function of 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 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."

We have so far focused on the maximum change in a player's equilibrium strategy due to any change in his higher-order beliefs. We also investigate the relationship of the change in strategy to the size of the change in beliefs. Towards this goal, firstly, we define an "embedding metric" on beliefs at each order (as well as on beliefs about the other players' actions). This metric has the crucial property of preserving the distances in lower-order beliefs when they are embedded in the space of higher-order beliefs as point masses, allowing us to sensibly compare variations at different orders. We ask how much a

player’s strategy varies as we vary his belief at some order k and keep all his other beliefs fixed. (To be able to do this without violating the coherency of his beliefs, we need an independence assumption about the different orders of beliefs, an assumption that is satisfied in traditional “independent private value” environments.) Now we can define the marginal impact of a change in k th-order beliefs as the variation in equilibrium strategies divided by the size of this change in beliefs as measured by our embedding metric. We show that, under global stability and the independence assumption, the marginal impact of changes in k th-order beliefs is at most b^k times a constant. This formalizes our notion that, under global stability, the marginal impact of higher-order beliefs decreases exponentially. In that case, precision in lower-order beliefs will be much more important than the precision in higher-order beliefs in approximating a problem. It also follows that the players’ equilibrium behavior would not change much if they formed erroneous higher-order beliefs. These assertions may all sound very natural; we should emphasize that they may easily fail when global stability does not hold. In particular, with linear best-responses, the marginal impact of k th-order beliefs actually increases exponentially in k whenever global stability does not hold.

It also follows from our assumptions that equilibrium behavior is continuous with respect to the product topology on type space that comes from the embedding metric. Under the assumption that the best-response correspondence is singleton and satisfies global stability under uncertainty, this further implies that the correspondence that takes each type to the set of all equilibrium actions for that type is lower semi-continuous.

Although there is a sizeable literature on the impact of higher-order uncertainty following Rubinstein (1989), the focus of most studies has been relaxation of common knowledge and lower semi-continuity of equilibrium in the worst-case scenarios, such as approximating common knowledge with common p -beliefs (Monderer and Samet (1989)), robustness of equilibrium against (possibly substantial) payoff uncertainty with small probability (Fudenberg, Kreps, and Levine (1988) and Kajii and Morris (1997)), and strong topologies under

which equilibrium is lower semi-continuous uniformly over all games (Monderer and Samet (1997) and Kajii and Morris (1998)). Most closely related to our work, Morris (2002) analyzes the impact of higher-order uncertainty within a model with linear best responses, reaching the conclusion that impact of higher-order beliefs can be arbitrarily large if we require a uniform bound over all games. Our focus differs in two ways. Firstly, we measure the impact of higher-order uncertainty within a single game (dropping the uniformity requirement). Second, while our sufficient condition implies continuity of best response, most of these papers analyze matrix games and naturally use the supremum metric on the mixed strategies, when the best response is generically discontinuous.

The outline of the paper is as follows. In the next section, we illustrate the relation between stability and dampening impact of higher-order beliefs using games with linear best responses. In Section 3, we present our basic model with independence assumption and introduce the embedding metric; we introduce global stability in Section 4 and provide sufficient conditions and examples for it in Section 5. Our major results are presented in Section 6 with independence assumption, and our main result is extended beyond this assumption in Section 7. Section 8 concludes. Some proofs are relegated to the Appendix.

2. EXAMPLES WITH LINEAR BEST RESPONSES

We will now show how dampening impact of higher-order uncertainty is equivalent to stability in games with linear best-response functions, such as the linear Cournot duopoly. This illustrates the close relationship between these two concepts which we will establish in a broader context in the later sections.

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

$$P = 1 - Q$$

where P is the price of a good and $Q = q_1 + q_2$ where q_i is the supply of firm $i \in N = \{1, 2\}$. The marginal cost of firm i is denoted by c_i , so that its payoff

function is

$$u_i(q_1, q_2) = q_i(1 - q_1 - q_2 - c_i).$$

The inverse-demand and payoff functions are common knowledge.

Each firm knows its own marginal cost. If we assumed that the marginal costs were common knowledge, then we would have the classical complete-information case. We could also allow incomplete information by assuming that (c_1, c_2) is drawn from a commonly known distribution, representing the beliefs of j about c_i conditional on its own cost c_j . If we further assumed that c_1 and c_2 were independently distributed, then this would correspond to the assumption that the firms' beliefs about the other firms' cost are common knowledge. In this paper, we do not make such strong informational assumptions; we want to allow variations in all levels of uncertainty. Firm j has a probability distribution t_j^1 on c_i , representing its beliefs about c_i . Firm i has also a probability distribution t_i^2 on t_j^1 , representing i 's beliefs about j 's beliefs about c_i . In general, firm i has probability distribution t_i^k on t_j^{k-1} , representing k th-order beliefs of firm i . Firm i 's type is the entire list $t_i = (c_i, t_i^1, t_i^2, \dots)$.

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 an 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 strategy q_i^* will maximize the expected payoff

$$E_i [q_i (1 - q_i - q_j^*(t_j) - c_i)] = q_i (1 - q_i - E_i [q_j^*(t_j)] - c_i),$$

where expectation E_i will be determined by its beliefs (t_i^1, t_i^2, \dots) at all levels, as $q_j^*(t_j)$ depends on the entire type t_j . This implies that

$$(2.1) \quad q_i^* = \frac{1 - c_i}{2} - \frac{1}{2} E_i [q_j^*(t_j)].$$

Of course, we also have

$$(2.2) \quad q_j^* = \frac{1 - c_j}{2} - \frac{1}{2} E_j [q_i^*(t_i)].$$

Substituting (2.2) in (2.1), we can obtain

$$(2.3) \quad q_i^* = \frac{1 - c_i}{2} - \frac{1 - E_i [c_j]}{4} + \frac{1}{4} E_i E_j [q_i^*].$$

A further substitution of (2.1) in (2.3) would yield

$$q_i^* = \frac{1 - c_i}{2} - \frac{1 - E_i[c_j]}{4} + \frac{1 - E_i E_j[c_i]}{8} - \frac{1}{8} E_i E_j E_i[q_j^*].$$

Here $E_i[c_j]$ depends only on t_i^1 , the beliefs of i about the cost of j , $E_i E_j[c_i]$ depends only on t_i^2 , the beliefs of i about the beliefs of j about the cost of i , and $E_i E_j E_i[q_j^*]$ depends on the third and all higher-order beliefs. In general,

$$q_i^* = \frac{1 - c_i}{2} - \frac{1 - E_i[c_j]}{4} + \frac{1 - E_i E_j[c_i]}{8} - \dots + \frac{1}{2^k} \underbrace{E_i E_j E_i \dots E_i}_{k \text{ times}}[q_j^*]$$

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 beliefs up to k th order, we know the equilibrium strategy q^* up to an error of at most $1/2^k$.

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

$$q_i^* = \frac{1 - c_i}{2} - \frac{1 - E_i[c_j]}{4} + \frac{1 - E_i E_j[c_i]}{8} - \frac{1 - E_i E_j E_i[c_j]}{16} + \dots$$

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.

2.2. General Case with Linear Best Responses. The analysis above can be easily generalized to the case with linear best-response functions

$$BR_i = a_i + b E_i[s_j]$$

where a_i is the underlying parameter for player i (such as $(1 - c_i)/2$) and s_j is the (unknown) action of player j . Now, the equilibrium strategies satisfy

$$(2.4) \quad s_i^* = a_i + b E_i[a_j] + b^2 E_i E_j[a_i] + \dots + b^k \underbrace{E_i E_j E_i \dots E_i}_{k \text{ times}}[s_j^*]$$

when k is odd. The absolute value of the coefficients will decrease exponentially, resulting in a convergent infinite series as above, if and only if $|b| < 1$.

- Note that this corresponds precisely to the stability of the equilibrium of the complete information game under the best-response correspondence.
- When the equilibrium is unstable, the impact of higher-order beliefs in equilibrium is actually higher than that of lower-order beliefs, and one must know the higher-order beliefs to an impossibly high level of precision in order to predict behavior.
- Our derivation in this section relies only on the formation of higher-order expectations—not on the particular type space used. Hence it applies to any type structure.
- We are only able to use the substitution trick here to derive a simple formula because of the linearity of the best-response function. In the general case a player's best response depends on the details of the entire distribution (as noted by Morris (2000)) and there is no direct relationship between a player's best responses under certainty and uncertainty, rendering such elementary analysis impossible and requiring the more sophisticated tools of the following sections.

Note also that Morris and Shin (2003) and Morris (2002) obtain specific examples with linear best responses similar to ours in this section. They focus on different issues; Morris and Shin (2003) focus on the role of public information while Morris (2002) focus on the large impact of higher-order expectations in the worst-case scenario (when the slope of the best response approaches 1).

2.3. A Traditional Type Structure. We have ex ante $a \sim N(0, 1)$, and each player i gets a private signal $x_i = a + \varepsilon_i$ where $\varepsilon_i \sim N(0, (1 - v)/v)$ for some $v \in (0, 1)$ and a , ε_1 , and ε_2 are all independent. For each i , assume $BR_i = E[a + bs_j|x_i] = a_i + bE[s_j|x_i]$ for some $b \geq 0$, where $a_i \equiv E[a|x_i] = vx_i$. The above is all common knowledge.

Check that, whenever $bv \neq 1$, we have a Bayesian Nash equilibrium s^* with

$$(2.5) \quad s_i^* = \frac{vx_i}{1 - bv}.$$

When $bv < 1$, equilibrium seems intuitive. When $bv > 1$, however, counterintuitively the coefficient of x_i is negative and hence s_i^* is decreasing in x_i . Now,

write s_i^* as a series of higher order expectation as in (2.4). Since the k th-order expectation of a is $E_i E_j E_i \dots E_j [a] = v^k x_i$, we have

$$s_i^* = vx_i + bv^2 x_i + b^2 v^3 x_i + \dots + b^k \underbrace{E_i E_j E_i \dots E_i}_{k \text{ times}} [s_j^*].$$

Firstly, notice that when $bv > 1$, higher-order terms increase exponentially, yielding a divergent series. This explodingly large impact of higher-order uncertainty, however, does not appear in the directly computed formula in (2.5). Second, when $bv < 1 < b$, we have a convergent series yielding seemingly intuitive formula in (2.5), despite the fact that marginal contributions of higher-order expectations increase exponentially. This is only because our single-dimensional type space forced the variations in higher-order expectations to decrease exponentially,⁴ compensating the increases in marginal contributions. But in the approximated real-life situation, the players will probably have higher-order doubts about this model. In that case, their higher-order expectations may vary significantly, leading to dramatically different behavior (under the equilibrium of more accurate model). In that case, the model's predictions about the behavior will be misleading, *and* considerations about higher-order beliefs within the model will yield a false sense of robustness.

3. MODEL WITH INDEPENDENCE

We consider a game among players $N = \{1, 2, \dots, n\}$. The source of underlying uncertainty is a payoff-relevant parameter $a \in A$ where (A, d) is a compact Polish space (i.e., a complete and separable metric space), where d is a metric on set A . (In the Cournot example above $a = (c_1, c_2) \in [0, 1]^2$.) Each player i has action space S_i , which is a compact metric space, and utility function $u_i : A \times S \rightarrow \mathbb{R}$ where $S = \prod_i S_i$.

Notation. Given any list X_1, \dots, X_n of sets, write $X_{-i} = \prod_{j \neq i} X_j$, $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) \in X_{-i}$, and $(x_i, x_{-i}) = (x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n)$. Likewise, for any family of functions $f_j : X_j \rightarrow Y_j$, $j \in N$, we define $f_{-i} : X_{-i} \rightarrow Y_{-i}$ by $f_{-i}(x_{-i}) = (f_j(x_j))_{j \neq i}$. Given any metric space (X, d) , write $\Delta(X)$ for

⁴This is a general phenomenon (see Samet (1998).)

the space of probability distributions on X , suppressing the fixed σ -algebra on X which at least contains all open sets; when we use product spaces, we will always use the product σ -algebra. We write d_i for the metric on S_i for each $i \in N$ and define the metric d_{-i} on S_{-i} by

$$d_{-i}(s_{-i}, s'_{-i}) = \max_{j \neq i} d_j(s_j, s'_j).$$

We now define the players' hierarchy of beliefs about the underlying parameter a . We confine ourself to the belief structures where a player's beliefs are independent from his own beliefs at other orders. We do this because we want to be able to (i) vary a player's k th-order beliefs without worrying about the coherency of his beliefs and (ii) measure the impact of this change on equilibrium strategies without worrying about its impact through the changes in the player's beliefs at other orders. (The independence assumption will be dropped in our main result.)

We define the beliefs (or type) of a player i inductively. His first order beliefs (about a) are represented by a probability distribution $t_i^1 \in \Delta_1 \equiv \Delta(A)$ on A . His k th-order beliefs (about t_{-i}^{k-1}) are represented by a probability distribution $t_i^k \in \Delta_k \equiv \Delta(\Delta_{k-1}^{n-1})$ on Δ_{k-1}^{n-1} . The type of a player i is the list

$$t_i = (t_i^1, t_i^2, t_i^3, \dots)$$

of all these probability distributions. We write T_i for the set of all possible types t_i of player i . We also write $T = \prod_i T_i$ for the set of all type profiles t . His beliefs are represented by the product measure $t_i^1 \times t_i^2 \times t_i^3 \times \dots$ of his beliefs $(t_i^1, t_i^2, t_i^3, \dots)$ at each order; that is, given any $\prod_{k=0}^{\infty} X_k \subset A \times T_{-i}$, the probability that he assigns to the event $\{(a, t_{-i}) \in \prod_{k=0}^{\infty} X_k\}$ is $\prod_{k=1}^{\infty} t_i^k(X_{k-1})$. (Here, of course, we have used the independence assumption.) We write $t \setminus \tilde{t}_i^k$ for the belief structure obtained by changing t_i^k to \tilde{t}_i^k in t ; $t \setminus \tilde{t}_{-i}^k$ and $t_i \setminus \tilde{t}_i^k$ are defined similarly.

Example 1. (Independent private value environment) *Take any incomplete-information game with payoffs $u_i(s; \theta_i)$ for each i where each $\theta_i \in \Theta_i$ is independently distributed with some probability distribution P_i and privately known by player i , and this is common knowledge. This game can be embedded in our*

framework, by taking $A = \cup_i \Theta_i$, $t_i^1 = \delta_{\theta_i}$, $t_i^2 = \hat{t}_i^2 \equiv P_{-i} \circ \xi^{-1}$ where $P_{-i} = \prod_{j \neq i} P_j$ and $\xi : \theta_{-i} \mapsto \prod_{j \neq i} \delta_{\theta_j}$, and taking $t_i^k = \hat{t}_i^k \equiv \delta_{\hat{t}_{-i}^{k-1}}$ for each $k > 2$, where δ_x denotes the measure that puts probability 1 on $\{x\}$.

A strategy of a player i is a mapping

$$t_i \mapsto s_i(t_i) \in S_i,$$

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 $E[u_i(a, s_i, s_{-i}^*(t_{-i})) | t_i]$ of $u_i(a, s_i, s_{-i}^*)$ under the probability distribution $t_i^1 \times t_i^2 \times t_i^3 \times \dots$ at each t_i and for each i .

Embedding metric. Throughout the paper, we will need a measure of the distances between probability distributions. We therefore introduce the following metric, which we will call *embedding metric*. Let (X, d) be any metric space. Given any $\mu, \mu' \in \Delta(X)$, we first write

$$(3.1) \quad \Delta_{\mu, \mu'} = \{\nu \in \Delta(X \times X) \mid \text{marg}_1 \nu = \mu, \text{marg}_2 \nu = \mu'\}$$

for the set of all joint probability distributions with marginals μ and μ' , where marg_i is the marginal distribution on the i th copy of X . Now we define our embedding metric d on $\Delta(X)$ by setting

$$(3.2) \quad d(\mu, \mu') = \inf_{\nu \in \Delta_{\mu, \mu'}} E_\nu [d(x_1, x_2)],$$

where E_ν is the expectation operator with respect to ν and (x_1, x_2) is a generic member of $X \times X$. It is easy to verify that this is an extension in the following sense: if μ and μ' are point masses at x and x' , respectively, then $d(\mu, \mu') = d(x, x')$ — thus the notational convenience of using d for both metrics. An equivalent definition is given by

$$(3.3) \quad d(\mu, \mu') = \inf_{Y \sim \mu, Y' \sim \mu'} E[d(Y, Y')]$$

where \inf is taken over all pairs Y and Y' of X -valued random variables with distributions μ and μ' , respectively, and coming from the same probability space, and E is the expectation operator on this space.

The embedding metric has the following property of preserving Lipschitz continuity; the proof is in the Appendix. Notice in the lemma that $\mu \circ f^{-1}$ is the distribution of $f(Y)$ for a random variable $Y \sim \mu$.

Lemma 1. *Let (X, d_X) and (Z, d_Z) be two metric spaces, and $f : X \rightarrow Z$ be such that*

$$d_Z(f(x), f(x')) \leq \lambda d_X(x, x') \quad (\forall x, x')$$

for some λ . Let also d_X and d_Z be the embedding metrics on $\Delta(X)$ and $\Delta(Z)$, respectively. Then,

$$d_Z(\mu \circ f^{-1}, \mu' \circ f^{-1}) \leq \lambda d_X(\mu, \mu') \quad (\forall \mu, \mu').$$

4. STABILITY AND HIGHER-ORDER UNCERTAINTY

We are now ready to present our sufficient condition for the dampening impact of higher order uncertainty: stability of equilibrium under the best-response function. The global stability of equilibrium is usually defined by the condition that the variation in the best response is less than the variation in the other players' strategies under certainty.⁵ We will first extend this notion to the best response function under uncertainty, which is not directly related to the best response function under certainty.

Best Responses. Given any player i and any probability distribution π on $A \times S_{-i}$, we write $BR_i(\pi)$ for the best response of player i when his beliefs about the underlying uncertainty a and the other players' actions s_{-i} are represented by π . Notice that we are taking the best response to be a function rather than a correspondence. Under certain conditions (e.g., when the strategy spaces are convex and utilities are strictly quasi-concave in own strategy), the best-response correspondence will indeed be singleton. In general, however, there may be multiple best responses. In those cases we will assume that the equilibrium uses a single consistent selection from the best-response correspondence. In the former case, the global stability defined below will be a property

⁵The usual definition appears to be 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.

of the game, while in the latter case, it will be a property of equilibrium. Under the independence assumption, we will have $\pi = t_i^1 \times \mu$ for some $t_i^1 \in \Delta(A)$ and $\mu \in \Delta(S_{-i})$. In that case, we will write $BR_i(t_i^1, \mu)$ instead of $BR_i(\pi)$. When it does not lead to any confusion, we will sometimes suppress some of the arguments (e.g., write $BR_i(\mu)$ when t_i^1 is fixed) or write it in the form of $BR_i(a, s_{-i}; t_i)$, denoting the best response of player i when his type is t_i , where a and s_{-i} are random variables.

Global Stability under uncertainty. We say that *global stability under uncertainty holds* iff there exists $b \in [0, 1)$ such that, given any $i \in N$, $t_i^1 \in \Delta(A)$, and any $\mu, \mu' \in \Delta(S_{-i})$,

$$d_i(BR_i(\mu), BR_i(\mu')) \leq b d_{-i}(\mu, \mu'),$$

where d_{-i} is the embedding metric on $\Delta(S_{-i})$.

The required condition for global stability is the standard condition for Lipschitz continuity (of each BR_i with respect to the embedding metric defined on $\Delta(S_{-i})$) 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. Of course, this is the same as saying that for each i there is a $b_i \in [0, 1)$ satisfying the above condition, since we can take $b = \max\{b_1, \dots, b_n\}$.

Global stability is sufficient to guarantee that the impact of higher-order beliefs on equilibrium is diminishing. This is formally expressed in the next result. Consider a change in a player i 's k th-order beliefs from t_i^k to $\tilde{t}_i^k = t_i^k \circ \phi^{-1}$, so that i believes that the other players' $k - 1$ st order beliefs have changed according to some mapping ϕ . The next result states that, in that case, the change in equilibrium strategy of player i can be at most b times the expected maximum change in the other players' equilibrium strategies due to the change in their $k - 1$ st order beliefs, under the original beliefs of i .

Proposition 1. *Assume global stability under uncertainty for some $b \in [0, 1)$. Then, given any t , $k > 1$, i and any measurable function $\phi : \Delta_{k-1}^{n-1} \rightarrow \Delta_{k-1}^{n-1}$,*

$$d_i(s_i^*(t_i \setminus t_i^k \circ \phi^{-1}), s_i^*(t_i)) \leq bE [d_{-i}(s_{-i}^*(t_{-i} \setminus \phi(t_{-i}^{k-1})), s_{-i}^*(t_{-i})) | t_i].$$

Proof. Let μ and μ' be the distributions of s_{-i}^* under t_i and $t_i \setminus t_i^k \circ \phi^{-1}$, respectively. Clearly, s_{-i}^* and $s_{-i}^* \circ \phi$ are two random variables coming from the same state space T_{-i} and have the distributions μ and μ' , respectively, under t_i . Therefore,

$$\begin{aligned} d_i(s_i^*(t_i \setminus t_i^k \circ \phi^{-1}), s_i^*(t_i)) &= d_i(BR_i(t_i^1, \mu), BR_i(t_i^1, \mu')) \\ &\leq \inf_{s_{-i} \sim \mu, s'_{-i} \sim \mu'} b_i E[d_{-i}(s_{-i}, s'_{-i})] \\ &\leq b_i E[d_{-i}(s_{-i}^*, s_{-i}^* \circ \phi) | t_i] \\ &= b_i E[d_{-i}(s_{-i}^*(t_{-i} \setminus \phi(t_{-i}^{k-1})), s_{-i}^*(t_{-i})) | t_i]. \end{aligned}$$

□

5. SUFFICIENT CONDITIONS FOR STABILITY

In this section we present two sets of sufficient conditions for global stability under uncertainty. Both sets of conditions are closely related to global stability under certainty. We first present a general class of games where global stability under uncertainty is closely related to global stability under certainty. This class is characterized by Assumption 1a.

Assumption 1a. Best-response function of player i takes the form of

$$(5.1) \quad BR_i(t_i^1, \mu) = f_i(E[g_i(a, s_{-i})])$$

where expectation is taken with respect to $t_i^1 \times \mu \in \Delta(A \times S_{-i})$; $f_i : X \rightarrow S_i$ and $g_i : A \times S_{-i} \rightarrow X$ are two Lipschitz continuous functions defined through some Banach space (X, d_X) ; i.e., there exist α_i and β_i such that $d_i(f_i(x), f_i(x')) \leq \alpha_i d_X(x, x')$ and $d_X(g_i(a, s_{-i}), g_i(a, s'_{-i})) \leq \beta_i d_{-i}(s_{-i}, s'_{-i})$.

Note that the functional form in (5.1) is satisfied whenever u_i is analytical and the optimization problem has an interior solution. (The Taylor expansion for the first order condition would imply such a functional form, where $E[g_i]$ is the vector of all moments.) The more substantial part of this assumption is that f_i and g_i are Lipschitz continuous. Under certainty, Assumption 1a yields

a best response function

$$BR_i(a, s_{-i}) = f_i(g_i(a, s_{-i})) \equiv h_i(a, s_{-i}).$$

Our equilibrium would be stable under the best response correspondence if $d_i(h_i(a, s_{-i}), h_i(a, s'_{-i})) \leq b_i d_{-i}(s_{-i}, s'_{-i})$ at each a for some $b_i < 1$. The latter condition is slightly weaker than the following assumption.

Assumption 1b. For each $i \in N$, we have $b_i \equiv \alpha_i \beta_i < 1$.

Proposition 2. *Assumptions 1a and 1b imply global stability under uncertainty.*

Proof. In the Appendix. □

That is, under Assumption 1a, global stability under uncertainty is implied by the existence of α_i 's and β_i 's that satisfy Assumption 1b. Moreover, whenever f or g is the identity, global stability under certainty and uncertainty will be equivalent. Hence, there is a close link between these two concepts. Although Assumption 1 might not be easy to check in general, our next example presents a general class of games where these conditions can be easily checked.

Example 2. For each $i \in N$, take $S_i = [\underline{x}, \bar{x}]$ for some $\underline{x}, \bar{x} \in \mathbb{R}$ and

$$u_i(a, s_i, s_{-i}) = \phi_i(s_i) g_i(a, s_{-i}) - c_i(s_i),$$

where $g_i : A \times S_{-i} \rightarrow \mathbb{R}$ is a continuously differentiable function with $|\partial g_i / \partial s_j| < \beta_i$ 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 $c_i'' \geq 0$. Note that g_i is Lipschitz continuous with parameter β_i with respect to the changes in s_{-i} . Check that

$$BR_i(t_i^1, \mu) = f_i(E[g_i(a, s_{-i})])$$

where $f_i(z)$ is \underline{x} if $z < c'(\underline{x}) / \phi'(\underline{x})$, \bar{x} if $z > c'(\bar{x}) / \phi'(\bar{x})$, and it is the unique solution x to the first order condition $c'(x) / \phi'(x) = z$ otherwise. By the inverse-function theorem, f_i is also Lipschitz continuous with parameter $\alpha_i = 1/\gamma_i$ where $\gamma_i = \min_{x \in [\underline{x}, \bar{x}]} (c'(x) / \phi'(x))' > 0$. Therefore, global stability is satisfied whenever $b \equiv \max_{i \in N} \beta_i / \gamma_i < 1$.

Focusing on games where the agents' strategy spaces are one-dimensional, our next result presents a simple sufficient condition for global stability, and hence for dampening impact of higher order uncertainty, in terms of second derivatives of the utility functions.

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

$$(5.2) \quad b_i \equiv \max_a \sum_{j \neq i} \frac{\max_s |\partial^2 u_i(a, s) / \partial s_i \partial s_j|}{\min_s |\partial^2 u_i(a, s) / \partial s_i^2|} < 1.$$

Then, we have global stability under uncertainty whenever (i) $BR_i(t_i^1, \mu)$ is in the interior of S_i for all $t_i^1 \times \mu$, or (ii) S_i is convex.

Proof. In the Appendix. □

Example 3. *Consider Cournot oligopoly with linear inverse-demand function P and arbitrary cost function c_i with $c_i'' \geq 0$ (where both P and c_i may depend on parameter a .) Check that $|\partial^2 u_i / \partial s_i \partial s_j| = |P'|$ and $|\partial^2 u_i / \partial s_i^2| = 2|P'| + c_i''$, so that*

$$b_i = \max \frac{|P'|}{2|P'| + c_i''} \leq \frac{1}{2},$$

yielding global stability.

6. EQUILIBRIUM IMPACT WITH INDEPENDENCE

In this section, using the embedding metric defined above, we will put a natural metric on the type space, which will allow us to compare variations in different orders of the type space. We will show that, under the previously stated conditions, variations in higher-order beliefs have a lower impact on equilibrium behavior than comparable variations in lower-order beliefs.

6.1. Embedding metric on beliefs. We now apply the embedding-metric construction inductively to define our embedding metric on beliefs of each player

at each order. First, for $k = 1$, we extend d to $\Delta_1 = \Delta(A)$ by setting

$$d(t_i^1, \tilde{t}_i^1) = \inf_{a \sim t_i^1, a' \sim \tilde{t}_i^1} E[d(a, a')]$$

at each $t_i^1, \tilde{t}_i^1 \in \Delta_1$ and to Δ_1^{n-1} by setting

$$d(t_{-i}^1, \tilde{t}_{-i}^1) = \max_{j \neq i} d(t_j^1, \tilde{t}_j^1)$$

at each $t_{-i}^1, \tilde{t}_{-i}^1 \in \Delta_1^{n-1}$. For any $k > 1$, we extend d to Δ_k by setting

$$d(t_i^k, \tilde{t}_i^k) = \inf_{Y \sim t_i^k, Y' \sim \tilde{t}_i^k} E[d(Y, Y')],$$

where Y and Y' take values in Δ_{k-1}^{n-1} (whose generic member is t_{-i}^{k-1}), and to Δ_k by setting

$$d(t_{-i}^k, \tilde{t}_{-i}^k) = \max_{j \neq i} d(t_j^k, \tilde{t}_j^k)$$

at each $t_{-i}^k, \tilde{t}_{-i}^k \in \Delta_1^{n-1}$.

6.2. Dampening impact of higher-order uncertainty. Assuming global stability, we will now find an upper bound for the change in equilibrium strategy caused by a change in any k th-order beliefs. When we consider comparable changes (according to d) at all orders k , this bound will be decreasing exponentially in k .

Proposition 4. *Assume that, for each $i \in N$, $BR_i(\cdot, \mu)$ is Lipschitz continuous uniformly on μ , i.e.,*

$$(6.1) \quad d_i(BR_i(t_i^1; \mu), BR_i(\tilde{t}_i^1; \mu)) \leq \alpha d_i(t_i^1, \tilde{t}_i^1) \quad (\forall \mu, t_i^1, \tilde{t}_i^1)$$

for some $\alpha \in \mathbb{R}$. Assume also global stability under uncertainty for parameter b . Then, in the model with independence, for any i , t_i , k , and any \tilde{t}_i^k ,

$$(6.2) \quad d_i(s_i^*(t_i), s_i^*(t_i \setminus \tilde{t}_i^k)) \leq \alpha b^{k-1} d(t_i^k, \tilde{t}_i^k).$$

The conclusion can be spelled out as follows: Change the beliefs of a player at some order k while all the other beliefs are fixed. The change in the equilibrium strategy due to this change in the beliefs is at most an exponentially decreasing function of k times the change in the beliefs according to our embedding metric.

In other words, the bound of the rate of change in equilibrium strategy as a function of k th-order belief is exponentially decreasing in k .

Proof. Firstly, for $k = 1$, (6.2) is just (6.1). Now assume that (6.2) holds at some $k - 1$, i.e., for all $j \in N$, \hat{t} , and \tilde{t}_j^{k-1} ,

$$(6.3) \quad d_j(s_j^*(\hat{t}_j), s_j^*(\hat{t}_j \setminus \tilde{t}_j^{k-1})) \leq \alpha b^{k-2} d(\hat{t}_j^{k-1}, \tilde{t}_j^{k-1}).$$

For any fixed t and $i \in N$, let us define $f : \Delta_{k-1}^{n-1} \rightarrow S_{-i}$ by setting

$$f(\hat{t}_{-i}^{k-1}) = BR_{-i}(t \setminus \hat{t}_{-i}^{k-1})$$

at each $\hat{t}_{-i}^{k-1} \in \Delta_{k-1}^{n-1}$. Fix $\hat{t}_i = t_i$, so that our induction hypothesis (6.3) becomes

$$d(f(\hat{t}_{-i}^{k-1}), f(\tilde{t}_{-i}^{k-1})) \leq \alpha b^{k-2} d(\hat{t}_{-i}^{k-1}, \tilde{t}_{-i}^{k-1}) \quad (\forall \hat{t}_{-i}^{k-1}, \tilde{t}_{-i}^{k-1}).$$

Then, by Lemma 1, for any \tilde{t}_i^{k-1} ,

$$d(t_i^k \circ f^{-1}, \tilde{t}_i^k \circ f^{-1}) \leq \alpha b^{k-2} d(t_i^k, \tilde{t}_i^k).$$

Notice that $t_i^k \circ f^{-1}$ and $\tilde{t}_i^k \circ f^{-1}$ are the distributions of s_{-i}^* under t_i and $t_i \setminus \tilde{t}_i^k$. Therefore, by global stability,

$$d_i(s_i^*(t_i), s_i^*(t_i \setminus \tilde{t}_i^k)) \leq b d(t_i^k \circ f^{-1}, \tilde{t}_i^k \circ f^{-1}) \leq \alpha b b^{k-2} d(t_i^k, \tilde{t}_i^k) \leq \alpha b^{k-1} d(t_i^k, \tilde{t}_i^k).$$

□

In case we only know a player's beliefs up to k th order and have no knowledge of his beliefs at higher orders, the following result tells us the accuracy with which we can predict his equilibrium behavior. This is important because, as argued in the Introduction, modelers would prefer not to have to specify the players' higher-order beliefs. This result might be thought to be a corollary to Proposition 4; it can be obtained simply by adding the effects of changes at $k + 1$ st and all higher orders. The validity of this infinite summation, however, will be established only when we prove Proposition 8, which is a more general form of this proposition and only assumes global stability and boundedness of

the strategy space.⁶ Notice also that our present result does not refer to any topology on the type space —although we used our embedding metric to reach this result.

Proposition 5. *Under the assumptions and the notation of Proposition 4, let $D_A = \max_{a,a' \in A} d(a, a')$. Let t_i, \tilde{t}_i be such that $t_i^l = \tilde{t}_i^l$ for all $l \leq k$ for some $k > 1$. Then, in the model with independence,*

$$d_i(s_i^*(t_i), s_i^*(\tilde{t}_i)) \leq b^k \alpha D_A / (1 - b).$$

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 strategies within a certain margin of error. Proposition 5 tells us how many orders of uncertainty he needs to specify. It implies that, given any $\epsilon > 0$ and any $t \in T$, if we know t up to the order

$$(6.4) \quad k \geq \frac{\log(\epsilon) - \log(\alpha D_A / (1 - b))}{\log(b)},$$

then we can compute the equilibrium strategies up to a maximum error of ϵ . Notice that the expression on the right-hand side is increasing in b and decreasing in ϵ .

6.3. Continuity in product topology. Many authors emphasized that equilibrium strategy is not continuous with respect to the product topology on type space and introduced stronger topologies, such as the topology of uniform convergence, in order to make the equilibrium strategies continuous (see Monderer and Samet (1996) and Kajii and Morris (1998)). (The equilibrium correspondence fails to be lower semi-continuous.) These authors require uniform convergence over all games, in essence focusing on the worst-case games, such as e-mail game which has high dependence on higher order beliefs. Moreover, they

⁶This infinite summation would give us a proof only if we had continuity at infinity, i.e., for any sequence $\{t[l]\}_{l \in \mathbb{N}}$ of types such that $t[l]$ is identical to some fixed type t at the first l orders, $\lim_{l \rightarrow \infty} d^{s^*}(t[l], t) = 0$. Proposition 8 implies the latter statement as an immediate corollary.

consider the games with discrete strategy-spaces, where the best-response correspondence cannot usually have any continuous selection, which is needed for global stability. Here we fix a game, and ask whether the equilibrium strategies of this game are continuous with respect to a product topology. Our next result answers this question in the affirmative for games satisfying global stability under uncertainty and for the product topology on type space generated by the embedding metric on beliefs.

Note that this topology is the topology of pointwise convergence under the embedding metric. That is, equilibrium strategy s_i^* is continuous with respect to this topology iff, for any sequence $\{t_{i,m}\}_{m \in \mathbb{N}}$ of types,

$$[t_{i,m}^k \rightarrow \hat{t}_i^k \quad \forall k \in \mathbb{N}] \Rightarrow [s_i^*(t_{i,m}) \rightarrow s_i^*(\hat{t}_i)],$$

where convergence of beliefs at each order is according to the embedding metric. Also, because the space of beliefs is compact under the embedding metric, this topology is metrized by the metric d_b (called a Fréchet metric) defined by

$$d_b(t_i, \tilde{t}_i) = \sum_{k=1}^{\infty} b^{k-1} d(t_i^k, \tilde{t}_i^k),$$

where b is any number in $(0, 1)$. Our next result states that, under global stability, the equilibrium strategy is Lipschitz continuous with respect to a Fréchet metric, and hence it is continuous in the product topology.

Proposition 6. *Under the assumptions and the notation of Proposition 4, in the model with independence, for each $i \in N$, the equilibrium strategy s_i^* of player i is Lipschitz continuous with respect to d_b . In that case, s_i^* is continuous with respect to the product topology on type space generated by the embedding metric on beliefs at each order.*

Proof. Fix any two types t_i and \tilde{t}_i of player i . For each $k \in \mathbb{N}$, define the type $t_{i,k}$ by setting

$$t_{i,k}^l = \begin{cases} t_i^l & \text{if } l \leq k, \\ \tilde{t}_i^l & \text{otherwise} \end{cases}$$

at each order l . We have

$$\begin{aligned} d_i(s_i^*(t_i), s_i^*(\tilde{t}_i)) &\leq \sum_{k=1}^{\infty} d_i(s_i^*(t_{i,k}), s_i^*(t_{i,k-1})) = \sum_{k=1}^{\infty} d_i(s_i^*(t_{i,k}), s_i^*(t_{i,k} \setminus \tilde{t}_i^k)) \\ &\leq \sum_{k=1}^{\infty} \alpha b^{k-1} d(t_i^k, \tilde{t}_i^k) = \alpha d_b(t_i, \tilde{t}_i), \end{aligned}$$

where α is as defined in Proposition 4, proving the result. To see the first inequality, note that we can change \tilde{t}_i to t_i by changing \tilde{t}_i^k to t_i^k one at a time. Hence, by Proposition 5, for each $\epsilon > 0$, there exists an integer l such that $d_i(s_i^*(t_i), s_i^*(\tilde{t}_i)) \leq \sum_{k=1}^l d_i(s_i^*(t_{i,k}), s_i^*(t_{i,k-1})) + \epsilon \leq \sum_{k=1}^{\infty} d_i(s_i^*(t_{i,k}), s_i^*(t_{i,k-1})) + \epsilon$. Since ϵ is arbitrary, this yields the inequality. The next equality is by definition; the next inequality Proposition 4, and the last equality is by definition. \square

Corollary 1. *Assume that players' best responses are always unique and satisfy global stability under uncertainty. Define equilibrium correspondence Σ^* by setting $\Sigma^*(t) = \{s^*(t) \mid s^* \text{ is a Bayesian Nash Equilibrium}\}$ at each $t \in T$. Then, under the assumptions of Proposition 4, Σ^* is lower semi-continuous with respect to the product topology on type space generated by the embedding metric on beliefs at each order.*

Proof. Take any t , any $s^*(t) \in \Sigma^*(t)$, and sequence $t(n)$ that converges to t in the topology above. By definition $s^*(t)$ is the value of a Bayesian Nash equilibrium s^* at t . Then, by Proposition 6, $s^*(t(n)) \in \Sigma^*(t(n))$ converges to $s^*(t)$. \square

7. WITHOUT INDEPENDENCE

We will now define the universal type space without imposing independence. We will show that our main result, namely Proposition 5, generalizes to this structure.

General model (without independence). The independence assumption was built into our previous model to allow for simpler notation and a clearer consideration of the effects of changing beliefs at a single order. In order to allow for the general case, we will need to consider the usual (and more complicated)

construction of the universal type space by Brandenburger and Dekel (1993), a variant of an earlier construction by Mertens and Zamir (1985). The meanings of k th-order beliefs in our two models are not parallel, as in the new model the k th-order belief will contain information about all lower orders as well. We will define our types using the auxiliary sequence $\{X_k\}$ of sets defined inductively by $X_0 = A$ and $X_k = [\Delta(X_{k-1})]^n \times X_{k-1}$ for each $k > 0$. We endow each X_k with the weak topology and the σ -algebra generated by this topology, yielding a standard separable Borel space as A is a Polish space. A player i 's first order beliefs are represented by a probability distribution τ_i^1 on X_0 , second order beliefs (about all players' first order beliefs and the underlying uncertainty) are represented by a probability distribution τ_i^2 on X_1 , etc. Therefore, a type τ_i of a player i is a member of $\prod_{k=1}^{\infty} \Delta(X_{k-1})$. Since a player's k th-order beliefs now contain information about his lower order beliefs, we need the usual coherence requirements. We write \mathcal{T} 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, i.e., the players know their own beliefs and their marginals from different orders agree. We will use the variables $\tau, \tilde{\tau} \in \mathcal{T}$ as generic type profiles. The rest of the model in Section 3 is unchanged.

Dropping the independence assumption causes two complications. First, since a player's higher-order beliefs contain information about his lower-order beliefs, we can no longer vary a player's belief at order k and keep his beliefs at order $l > k$ constant —as in Proposition 4— without violating the coherency requirements. Instead, we allow all the beliefs at all orders higher than k to vary (as in Proposition 5). Second, since the other players' actions are now (possibly) perceived to be correlated with the underlying uncertainty, we need to extend our definition of global stability to allow such correlation. To do this, let \bar{d}_{-i} be any metric on $A \times S_{-i}$ such that

$$(7.1) \quad \bar{d}_{-i}((a, s_{-i}), (a, s'_{-i})) = d_{-i}(s_{-i}, s'_{-i})$$

for each $a \in A$ and $s_{-i}, s'_{-i} \in S_{-i}$, so that the metric d_{-i} on S_{-i} is preserved when S_{-i} is embedded in $A \times S_{-i}$. Extend also \bar{d}_{-i} to $\Delta(A \times S_{-i})$ using the embedding metric as before. We are able to leave the metric \bar{d}_{-i} only partially specified

since global stability is related only to responsiveness of the best response with respect to the changes in the other players' strategies. Thus we will be able to prove that as long as the inequality below is satisfied for some \bar{d}_{-i} our results will hold.

Global Stability under uncertainty in general model. We say that *global stability under uncertainty* holds iff there exist some $b \in [0, 1)$ and some embedding metric \bar{d}_{-i} on $\Delta(A \times S_{-i})$ satisfying (7.1) and such that, given any $i \in N$ and any $\pi, \pi' \in \Delta(A \times S_{-i})$ with $\text{marg}_A \pi = \text{marg}_A \pi'$, we have

$$(7.2) \quad d_i(BR_i(\pi), BR_i(\pi')) \leq b \bar{d}_{-i}(\pi, \pi').$$

The next proposition extend Propositions 2 and 3 to the present set up. One can also check that Example 2 of Section 5 remains valid under the new definition, while we will have global stability under the new definition in Example 3 whenever $\max_a |P'| \leq 2 \min_a |P'|$.

Proposition 7. (a) *Assumptions 1a and 1b imply global stability under uncertainty.* (b) *For each i , assume $S_i \subset \mathbb{R}$, $u_i(a, \cdot)$ is twice-continuously differentiable, $u_i(a, \cdot, s_{-i})$ is strictly concave, $\partial^2 u_i / \partial s_i^2$ is bounded away from zero, and*

$$(7.3) \quad b_i \equiv \sum_{j \neq i} \frac{\max_{s,a} |\partial^2 u_i(a, s) / \partial s_i \partial s_j|}{\min_{s,a} |\partial^2 u_i(a, s) / \partial s_i^2|} < 1.$$

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

Proof. In the Appendix. □

We are now ready to state our main result, which extends Proposition 5 to our general model.

Proposition 8. *Let $D_S = \max_{i \in N} \sup_{\pi, \pi' \in \Delta(A \times S_{-i})} d_i(BR_i(\pi), BR_i(\pi')) \in \mathbb{R}$. Let also $\tau, \tilde{\tau} \in \mathcal{T}$ be such that $\tau_i^l = \tilde{\tau}_i^l$ for all $l \leq k$ for some $k \geq 0$. Assume global stability under uncertainty for parameter b . Then, in the general model,*

$$(7.4) \quad d_i(s_i^*(\tau_i), s_i^*(\tilde{\tau}_i)) \leq b^k D_S.$$

Notice that our result assumes only global stability and boundedness of the strategy 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 play within a bound of error that is an exponentially decreasing function of k , bounding the maximum impact all the higher-order beliefs can have on equilibrium. Our result does not refer to any topology on the type space. Finally, D_S is chosen as a bound on the variations in equilibrium outcomes. If there are other known bounds on the equilibrium outcomes (perhaps due to some support restrictions), then we can replace D_S with these bounds. In the remainder of the section we prove our proposition. We start with the following technical lemma.

Lemma 2. *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 $\text{marg}_{X \times Z} \tilde{P} = P'$.*

Proof. In the Appendix. □

Proof of Proposition 8. Define $\Omega = A \times \mathcal{T}$ to be the universal state space. This is the subset of the larger space $\bar{\Omega} = A \times (\prod_{k=1}^{\infty} \Delta(X_{k-1}))^n$ 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 $\tau = (\tau_1, \dots, \tau_n) \in \mathcal{T}$ and for every $i \in N$, there exists a probability distribution $\kappa_{\tau_i} \in \Delta(\bar{\Omega})$ such that

$$(7.5) \quad \text{marg}_{X_{k-1}} \kappa_{\tau_i} = \tau_i^k \quad (\forall k),$$

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

$$\beta : (a, \tau) \mapsto (a, s_{-i}^*(\tau_{-i})),$$

and write

$$\pi_{\tau_i} = \kappa_{\tau_i} \circ \beta^{-1} \in \Delta(A \times S_{-i})$$

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

We will use induction on k . For $k = 0$, this is true by definition. Fix any $k > 0$, and assume that the result is true for $k - 1$. Take any τ and $\tilde{\tau}$ as in the hypothesis. We have

$$\begin{aligned}
 d_i(s_i^*(\tau_i), s_i^*(\tilde{\tau}_i)) &= d_i(BR_i(\pi_{\tau_i}), BR_i(\pi_{\tilde{\tau}_i})) \\
 &\leq b\bar{d}_{-i}(\pi_{\tau_i}, \pi_{\tilde{\tau}_i}) \\
 (7.6) \qquad \qquad \qquad &\equiv b \inf_{\nu \in \Delta_{\pi_{\tau_i}, \pi_{\tilde{\tau}_i}}} E_\nu [\bar{d}_{-i}((a, s_{-i}), (a', s'_{-i}))],
 \end{aligned}$$

where the inequality is due to global stability and $\Delta_{\pi_{\tau_i}, \pi_{\tilde{\tau}_i}}$ is defined by (3.1). The rest of the proof is devoted to constructing a $\nu \in \Delta_{\pi_{\tau_i}, \pi_{\tilde{\tau}_i}}$ such that, under the induction hypothesis,

$$(7.7) \qquad E_\nu [\bar{d}_{-i}((a, s_{-i}), (a', s'_{-i}))] \leq b^{k-1} D_S.$$

Combining (7.6) and (7.7), we obtain (7.4).

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

$$(7.8) \qquad L = \prod_{l=1}^{k-1} (\Delta(X_{l-1}))^n \text{ and } H = \prod_{l=k}^{\infty} (\Delta(X_{l-1}))^n$$

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

By (7.5), we have probability distributions κ_{τ_i} and $\kappa_{\tilde{\tau}_i}$ on $\bar{\Omega}$ such that

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

where the second equality is by our hypothesis. Since we have separable standard Borel spaces, by Lemma 2, 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

$$\text{marg}_{12} \sigma = \kappa_{\tau_i} \qquad \text{and} \qquad \text{marg}_{13} \sigma = \kappa_{\tilde{\tau}_i},$$

respectively.

Now, consider $\nu = \sigma \circ \gamma^{-1} \in \Delta((A \times S_{-i})^2)$ where

$$(7.9) \quad \gamma : (a, l, h_1, h_2) \mapsto (\beta(a, l, h_1), \beta(a, l, h_2)).$$

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

$$\text{marg}_1 \nu = \text{marg}_1 (\sigma \circ \gamma^{-1}) = (\text{marg}_{12} \sigma) \circ \beta^{-1} = \kappa_{\tau_i} \circ \beta^{-1} = \pi_{\tau_i},$$

and similarly $\text{marg}_2 \nu = \pi_{\bar{\tau}_i}$. Therefore, by definition, $\nu \in \Delta_{\pi_{\tau_i}, \pi_{\bar{\tau}_i}}$.

We now prove (7.7). Write $I \equiv \gamma(X_{k-1} \times H \times H)$ and take any $((a, s_{-i}), (a', s'_{-i})) \in I$. By (7.9), we have $a = a'$, and hence by (7.1),

$$(7.10) \quad \bar{d}_{-i}((a, s_{-i}), (a', s'_{-i})) = d_{-i}(s_{-i}, s'_{-i}).$$

But by (7.9), $s_{-i} = s_{-i}^*(\hat{\tau}_{-i})$ and $s'_{-i} = s_{-i}^*(\bar{\tau}_{-i})$ for some type profiles $\hat{\tau} = (l, h_1)$ and $\bar{\tau} = (l, h_2)$, which agree up to the order $k - 1$ by (7.8). Then, by the induction hypothesis,

$$(7.11) \quad d_{-i}(s_{-i}, s'_{-i}) \leq b^{k-1} D_S.$$

Combining (7.10) and (7.11), we obtain

$$(7.12) \quad \bar{d}_{-i}((a, s_{-i}), (a', s'_{-i})) \leq b^{k-1} D_S.$$

Since $\text{supp} \nu \subset I$ (by construction), (7.12) implies (7.7). \square

8. CONCLUSION

Present economic theories are mostly based on equilibrium analysis of models in which only a few low orders of uncertainty are properly specified. We know, however, that in some games higher-order uncertainty may have a large impact in equilibrium, changing equilibrium behavior profoundly. In this paper we presented a sufficient condition, namely global stability under uncertainty, which guarantees that the impact of higher-order uncertainty is low. Using the universal type space, in which players can entertain any coherent set of beliefs, 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 (cf. Proposition 8). That is, if a theoretical prediction requires knowledge of the strategies within a margin ϵ of error, then the researcher

can validate his theory by specifying first $k(\epsilon)$ orders of beliefs, where $k(\epsilon)$ is a logarithmic function of ϵ . Under a further independence assumption we also formalize our notion that, under stability, the marginal impact of higher-order uncertainty is (exponentially) decreasing in the order (cf. Propositions 1 and 4).⁷ That is, the problem must be approximated using lower-order uncertainty rather than higher-order uncertainty; this may be reversed when stability does not hold, as the impact of higher-order uncertainty may grow exponentially. In the latter case, we believe that accurate prediction using traditional analysis will be impossible.

Our study has two important limitations which require further research, imposed by our use of normal-form representation and the solution concept of (unrestricted) Bayesian Nash equilibrium. Firstly, many theories are based on extensive-form representations and use refinements, such as sequential rationality (Selten (1974), Kreps and Wilson (1982)). Their predictions are often driven by these refinements when equilibrium itself does not have any predictive power in their games. It is then crucial to extend our analysis to such a framework, using extensive-form constructions, such as Battigalli and Siniscalchi (1999). Secondly, although we are motivated by the wide use of equilibrium analysis, the concept of equilibrium has only very limited epistemic support (see Aumann and Brandenburger (1995) and Dekel and Gul (1997)). Hence, extension of our work for better-founded concepts, such as rationalizability, will be useful; our results would also be complemented by extensions of Milgrom and Roberts' (1990) results to universal type space, bounding rationalizable strategies by equilibria.

APPENDIX A. OMITTED PROOFS

A.1. Proof of Lemma 1. Take any $\mu, \mu' \in \Delta(X)$, and fix any $\epsilon > 0$. By definition of $d_X(\mu, \mu')$, there exists $\nu \in \Delta_{\mu, \mu'}$ such that

$$(A.1) \quad E_\nu [d_X(x_1, x_2)] \leq d_X(\mu, \mu') + \epsilon.$$

⁷It also follows from these assumptions that the equilibrium strategy is continuous in player's type with respect to a product topology (cf. Proposition 6).

Define $\bar{f} : X^2 \rightarrow Z^2$ by $\bar{f}(x_1, x_2) = (f(x_1), f(x_2))$. Then, by definition, $\nu \circ \bar{f}^{-1} \in \Delta_{\mu \circ f^{-1}, \mu' \circ f^{-1}}$. Hence,

$$\begin{aligned} d_Z(\mu \circ f^{-1}, \mu' \circ f^{-1}) &\leq E_{\nu \circ \bar{f}^{-1}}[d_Z(z_1, z_2)] \\ &= E_\nu[d_Z(f(x_1), f(x_2))] \\ &\leq E_\nu[\lambda d_X(x_1, x_2)] = \lambda E_\nu[d_X(x_1, x_2)] \\ &\leq \lambda d_X(\mu, \mu') + \lambda \epsilon; \end{aligned}$$

since ϵ is arbitrary, the result follows. [Here, the first inequality is by (3.2); the next equality is by change of variables, the next inequality is by the hypothesis, and the last inequality is by (A.1).] \square

A.2. Proof of Proposition 2. Take any $i \in N$, $t_i^1 \in \Delta(A)$, and any $\mu, \mu' \in \Delta(S_{-i})$. Recall that $d_{-i}(\mu, \mu') \equiv \inf_{s_{-i} \sim \mu, s'_{-i} \sim \mu'} E[d_{-i}(s_{-i}, s'_{-i})]$. Take any random variable a with distribution $t_i^1 \in \Delta(A)$, and any two random variables s_{-i} and s'_{-i} with distributions μ and μ' , respectively. By Assumption 1, we have $BR_i(t_i^1, \mu) = f_i(E[g_i(a, s_{-i})])$ and $BR_i(t_i^1, \mu') = f_i(E[g_i(a, s'_{-i})])$. Hence,

$$\begin{aligned} d_i(BR_i(\mu), BR_i(\mu')) &= d_i(f_i(E[g_i(a, s_{-i})]), f_i(E[g_i(a, s'_{-i})])) \\ &\leq \alpha_i d_X(E[g_i(a, s_{-i})], E[g_i(a, s'_{-i})]) \\ &\leq \alpha_i (E[d_X(g_i(a, s_{-i}), g_i(a, s'_{-i}))]) \\ &\leq \alpha_i \beta_i E[d_{-i}(s_{-i}, s'_{-i})] = b_i E[d_{-i}(s_{-i}, s'_{-i})], \end{aligned}$$

where the first and the last inequalities are due to Assumption 1, and the second inequality is by triangle inequality. Since s_{-i} and s'_{-i} are arbitrary, this yields

$$d_i(BR_i(\mu), BR_i(\mu')) \leq b_i d_{-i}(\mu, \mu').$$

A.3. Proof of Proposition 3. Take any $t_i^1 \in \Delta(A)$ and $\mu, \mu' \in \Delta(S_{-i})$. We will assume $BR_i(t_i, \mu)$ and $BR_i(t_i, \mu')$ are in the interior of S_i . (When S_i is convex, we can take $BR_i(t_i, \mu)$ and $BR_i(t_i, \mu')$ as the unconstrained optima, as in that case the variations in the constrained optima are if anything less than the variations in unconstrained optima.) We write

$$U^i(s; t_i^1) = \int u_i(a, s) dt_i^1(a)$$

and write U_i^i , U_{ii}^i , and U_{ij}^i for the first and second order partial derivatives of U^i with respect to s_i , and the cross partial with respect s_i and s_j , respectively. Firstly, since $BR_i(\mu)$ and $BR_i(\mu')$ are in the interior, the first order conditions for optimization problems with μ and μ' yield

$$(A.2) \quad E [U_i^i (BR_i(\mu), s_{-i})] = 0$$

and

$$(A.3) \quad E [U_i^i (BR_i(\mu'), s'_{-i})] = 0,$$

respectively. Let

$$J = E [U_i^i (BR_i(\mu), s'_{-i})]$$

be the value of the derivative at $BR_i(\mu)$ for the optimization problem with μ' . We will now find upper and lower bounds for $|J|$, and these bounds will yield (??). First we find an upper bound:

$$\begin{aligned} |J| &= |E [U_i^i (BR_i(\mu), s'_{-i})]| \\ &= |E [U_i^i (BR_i(\mu), s'_{-i}) - U_i^i (BR_i(\mu), s_{-i})]| \\ &\leq E [|U_i^i (BR_i(\mu), s'_{-i}) - U_i^i (BR_i(\mu), s_{-i})|] \\ &\leq E \left[\sum_{j \neq i} \max_s |U_{ij}^i (s; t_i^1)| d_{-i} (s_{-i}, s'_{-i}) \right] \\ (A.4) \quad &= \sum_{j \neq i} \max_s |U_{ij}^i (s; t_i^1)| E [d_{-i} (s_{-i}, s'_{-i})]. \end{aligned}$$

Here the first equality is by definition, the second equality is by (A.2), and the following inequality is by the triangle inequality. To derive the penultimate inequality, we write $U_i^i (BR_i(\mu), s'_{-i}) - U_i^i (BR_i(\mu), s_{-i})$ as the some of the changes that we would get by changing each coordinate in turn, and apply the mean value theorem to each, obtaining

$$\begin{aligned} |U_i^i (BR_i(\mu), s'_{-i}) - U_i^i (BR_i(\mu), s_{-i})| &\leq \sum_{j \neq i} \max_s |U_{ij}^i (s; t_i^1)| |s_j - s'_j| \\ &\leq \sum_{j \neq i} \max_s |U_{ij}^i (s; t_i^1)| d_{-i} (s_{-i}, s'_{-i}), \end{aligned}$$

where the last inequality is by our definition of the metric d_{-i} . To find our lower bound, we write

$$\begin{aligned}
|J| &= |E [U_i^i (BR_i (\mu), s'_{-i})]| \\
&= |E [U_i^i (BR_i (\mu), s'_{-i}) - U_x^i (BR_i (\mu'), s'_{-i})]| \\
&= E [|U_i^i (BR_i (\mu), s'_{-i}) - U_i^i (BR_i (\mu'), s'_{-i})|] \\
&\geq E \left[\min_s |U_{ii}^i (s; t_i^1)| |BR_i (\mu) - BR_i (\mu')| \right] \\
\text{(A.5)} \quad &= \min_s |U_{ii}^i (s; t_i^1)| |BR_i (\mu) - BR_i (\mu')|.
\end{aligned}$$

Here the first and the second equalities are by definition and (A.3), respectively. The third equality is crucial; we have equality here because $U_i^i (\cdot; s'_{-i})$ is strictly decreasing, and hence $U_i^i (BR_i (\mu), s'_{-i}) - U_i^i (BR_i (\mu'), s'_{-i})$ never changes its sign. The inequality in the next line is again by the mean value theorem, and the last equality is because the term inside the expectation is a constant. Combining (A.4) and (A.5) and observing that $d_{-i} (\mu, \mu') \equiv \inf_{s_{-i} \sim \mu, s'_{-i} \sim \mu'} E [d_{-i} (s_{-i}, s'_{-i})]$ and that s_{-i} and s'_{-i} are arbitrary, we obtain

$$|BR_i (\mu) - BR_i (\mu')| \leq d_{-i} (\mu, \mu') \sum_{j \neq i} \frac{\max_s |U_{ij}^i (s; t_i^1)|}{\min_s |U_{ii}^i (s; t_i^1)|}.$$

Check that $\max_s |U_{ij}^i (s; t_i^1)| \leq \int \max_s |\partial^2 u_i (a, s) / \partial s_i \partial s_j| dt_i^1 (a)$ and $\min_s |U_{ii}^i (s; t_i^1)| \geq \int \min_s |\partial^2 u_i (a, s) / \partial s_i^2| dt_i^1 (a)$. Therefore,

$$\begin{aligned}
\sum_{j \neq i} \frac{\max_s |U_{ij}^i (s; t_i^1)|}{\min_s |U_{ii}^i (s; t_i^1)|} &\leq \sum_{j \neq i} \frac{\int \max_s |\partial^2 u_i (a, s) / \partial s_i \partial s_j| dt_i^1 (a)}{\int \min_s |\partial^2 u_i (a, s) / \partial s_i^2| dt_i^1 (a)} \\
&\leq \sum_{j \neq i} \int \frac{\max_s |\partial^2 u_i (a, s) / \partial s_i \partial s_j|}{\min_s |\partial^2 u_i (a, s) / \partial s_i^2|} dt_i^1 (a) \\
&\leq \max_a \sum_{j \neq i} \frac{\max_s |\partial^2 u_i (a, s) / \partial s_i \partial s_j|}{\min_s |\partial^2 u_i (a, s) / \partial s_i^2|} < 1,
\end{aligned}$$

completing the proof.

A.4. Proof of Proposition 7. The proof of part (b) is very similar to the proof of Proposition 3 above. We will prove part (a).

Under Assumptions 1a and 1b, take any $i \in N$. Firstly, if $\beta_i = 0$, then $g_i(a, s_{-i}) = g_i(a, s'_{-i}) = \tilde{g}_i(a)$ for each (a, s_{-i}, s'_{-i}) , hence, for each π, π' with $\text{marg}_A \pi = \text{marg}_A \pi'$, we have $BR_i(\pi) = f_i(E_\pi(\tilde{g}_i(a))) = f_i(E_{\pi'}(\tilde{g}_i(a))) = BR_i(\pi')$, yielding $d_i(BR_i(\pi), BR_i(\pi')) = 0 \leq b_i \bar{d}_{-i}(\pi, \pi')$ for any \bar{d}_{-i} .

Now assume that $\beta_i > 0$. Since g_i is continuous and $A \times S_{-i}$ is compact, there exists $M_i > 0$ such that

$$(A.6) \quad d_X(g_i(a, s_{-i}), g_i(a', s'_{-i})) \leq M_i \quad (\forall a, s_{-i}, a', s'_{-i}).$$

Define a metric $d_{A,i}$ on A by setting $d_{A,i}(a, a') = M_i/\beta_i$ at each distinct a, a' , and define \bar{d}_{-i} on $A \times S_{-i}$ by

$$\bar{d}_{-i}((a, s_{-i}), (a', s'_{-i})) = d_{A,i}(a, a') + d_{-i}(s_{-i}, s'_{-i}).$$

Now, take any two random variables $(a, s_{-i}) \sim \pi$ and $(a', s'_{-i}) \sim \pi'$ that come from the same probability space and write p for the probability that $a \neq a'$. Note that

$$(A.7) \quad E[\bar{d}_{-i}((a, s_{-i}), (a', s'_{-i}))] = pM_i/\beta_i + E[d_{-i}(s_{-i}, s'_{-i})].$$

Moreover, we have

$$\begin{aligned} d_i(BR_i(\pi), BR_i(\pi')) &\leq \alpha_i E[d_X(g_i(a, s_{-i}), g_i(a, s'_{-i}))] \\ &= \alpha_i E[d_X(g_i(a, s_{-i}), g_i(a, s'_{-i})) : a \neq a'] \\ &\quad + \alpha_i E[d_X(g_i(a, s_{-i}), g_i(a, s'_{-i})) : a = a'] \\ &\leq \alpha_i p M_i + \alpha_i \beta_i E[d_{-i}(s_{-i}, s'_{-i}) : a = a'] \\ &\leq \alpha_i p M_i + \alpha_i \beta_i E[d_{-i}(s_{-i}, s'_{-i})] \\ &= b_i (pM_i/\beta_i + E[d_{-i}(s_{-i}, s'_{-i})]) \\ &= b_i E[\bar{d}_{-i}((a, s_{-i}), (a', s'_{-i}))], \end{aligned}$$

where the first inequality is derived as in the proof of Proposition 2, the next equality is by additivity, the next equality is by (A.6) and the Lipschitz continuity of g_i , the next inequality is by the non-negativity of d_{-i} , and the last two equalities are by definition of b_i and (A.7). Since $(a, s_{-i}) \sim \pi$ and $(a', s'_{-i}) \sim \pi'$ are arbitrary, this shows that $d_i(BR_i(\pi), BR_i(\pi')) \leq b_i E[\bar{d}_{-i}(\pi, \pi')]$.

A.5. Proof of Lemma 2. Let $\hat{P} \equiv \text{marg}_X P = \text{marg}_X P'$. Since we have separable standard Borel spaces, there exists conditional probability $P(\cdot|\cdot) : (\Sigma_{X \times Y}) \times (X \times Y) \rightarrow [0, 1]$ with respect to the σ -field $\Sigma_X \times \{Y\}$, and we simply write $P(B|x)$ for $P(X \times B | (x, y))$ where y can be chosen arbitrarily. We define $P'(C|x)$ similarly for each $C \in \Sigma_Z$. Notice that $P(\cdot|x)$ and $P'(\cdot|x)$ are probability distributions on (Y, Σ_Y) and (Z, Σ_Z) , respectively.⁸ For each $x \in X$, let

$$\tilde{P}_x \equiv P(\cdot|x) \times P'(\cdot|x)$$

be the product measure of $P(\cdot|x)$ and $P'(\cdot|x)$ on $Y \times Z$, and define probability measure \tilde{P} by setting

$$\tilde{P}(F) = \int \tilde{P}_x(F_x) d\hat{P}(x)$$

at each measurable set $F \subseteq X \times Y \times Z$ where

$$F_x = \{(y, z) \in Y \times Z | (x, y, z) \in F\}.$$

Notice that, for any rectangle $A \times B \times C \in \Sigma_X \times \Sigma_Y \times \Sigma_Z$,

$$\tilde{P}(A \times B \times C) = \int \chi_A(x) P(B|x) P'(C|x) d\hat{P}(x),$$

where χ_A denotes the characteristic function of A .

Now we show that \tilde{P} satisfies the statement of the lemma. For each $A \in \Sigma_X$ and $B \in \Sigma_Y$, we have

$$\begin{aligned} \text{marg}_{X \times Y} \tilde{P}(A \times B) &\equiv \tilde{P}(A \times B \times Z) \\ &= \int \chi_A(x) P(B|x) P'(Z|x) d\hat{P}(x) \\ &= \int \chi_A(x) P(B|x) d\hat{P}(x) \\ &\equiv P(A \times B). \end{aligned}$$

Since the probability measures $\text{marg}_{X \times Y} \tilde{P}$ and P agree on the π -system of all rectangles $A \times B$, which generates the entire σ -field on $X \times Y$, by Dynkin's π - λ Theorem they are equal. This is similarly true for $\text{marg}_{X \times Z} \tilde{P}$ and P' .

⁸See Parthasaraty (1967) for the results of probability theory in this proof.

REFERENCES

- [1] Aumann and Brandenburger (1995): "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63, 1161-1180.
- [2] Battigalli and Siniscalchi (1999): "Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic games," *Journal of Economic Theory*, 88, 188-230.
- [3] Brandenburger, A. and E. Dekel (1993): "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189-198.
- [4] Carlsson, H. and E. van Damme (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61, 989-1018.
- [5] Dekel, E. and F. Gul (1997): "Rationality and Knowledge in Game Theory," in Kreps, D. and K. Wallis (eds): *Advances in Economics and Econometrics: Theory and Applications*, Cambridge University Press, UK.
- [6] Feinberg, Y. and A. Skrzypacz (2002): "Uncertainty about uncertainty and delay in bargaining," mimeo.
- [7] Fudenberg, D., D. Kreps, and D. Levine (1988): "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory*, 44, 354-380.
- [8] Harsanyi, J. (1967): "Games with Incomplete Information played by Bayesian Players. Part I: the Basic Model," *Management Science* 14, 159-182.
- [9] Kajii, A. and S. Morris (1997): "The Robustness of Equilibria to Incomplete Information," *Econometrica*, 65, 1283-1309.
- [10] Kajii, A. and S. Morris (1998): "Payoff Continuity in Incomplete Information Games," *Journal of Economic Theory*, 82, 267-276.
- [11] Kreps, D. and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50-4, 863-894.
- [12] Mertens and Zamir (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 10, 619-632.
- [13] Monderer, D. and D. Samet (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170-190.
- [14] Monderer, D. and D. Samet (1997): "Proximity of information in games with incomplete information," *Mathematics of Operations Research*, 21, 707-725.
- [15] Morris, S. (2002): "Typical Types," mimeo.
- [16] Morris, S. and H. Shin (2002): "Social Value of Public Information," *American Economic Review*, 92, 1521-1534.
- [17] Milgrom, P. and J. Roberts (1990): "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58-6, 1255-1277.
- [18] Parthasarathy, K. (1967): *Probability Measures on Metric Spaces*, Academic Press, New York.

- [19] Rubinstein, A. (1989): “The Electronic Mail Game: Strategic Behavior Under ‘Almost Common Knowledge,’” *The American Economic Review*, Vol. 79, No. 3, pp. 385-391.
- [20] Samet, D. (1998): “Iterated Expectations and Common Priors,” *Games and Economic Behavior*, 24, 131-141.
- [21] Selten, R. (1975): “Re-examination of the perfectness concept for equilibrium points in extensive games,” *International Journal of Game Theory*, 4, 25-55.
- [22] Townsend, R. (1983): “Forecasting the Forecasts of Others,” *Journal of Political Economy*, 9-4, 546-88.
- [23] 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>