

Interim Bayesian Nash Equilibrium on Universal Type Spaces for Supermodular Games

Timothy Van Zandt
INSEAD

25 January 2007

We prove the existence of a greatest and a least interim Bayesian Nash equilibrium for supermodular games of incomplete information. There are two main differences from the earlier proofs in Vives (1990) and Milgrom and Roberts (1990): (a) we use the interim formulation of a Bayesian game, in which each player's beliefs are part of his or her type rather than being derived from a prior; (b) we use the interim formulation of a Bayesian Nash equilibrium, in which each player and every type (rather than almost every type) chooses a best response to the strategy profile of the other players. Given also the mild restrictions on the type spaces, we have a proof of interim Bayesian Nash equilibrium for universal type spaces (for the class of supermodular utilities), as constructed, for example, by Mertens and Zamir (1985). We also weaken restrictions on the set of actions.

KEYWORDS: Supermodular games, incomplete information, universal type spaces, interim Bayesian Nash equilibrium.

AUTHOR'S ADDRESS:

Timothy Van Zandt
INSEAD
Boulevard de Constance
77305 Fontainebleau CEDEX
France

Email: tvz@insead.edu
Web: faculty.insead.edu/vanzandt

ACKNOWLEDGMENTS: This paper benefited not merely from general comments but also from specific contributions by several people. I am particularly grateful to Kevin Reffett, who provided me with many helpful hints, discussions, and sources, and pointed out that Theorem 2 is the Tarski-Kantorovich theorem. I had several helpful discussions with Jeffrey Ely and he pointed out Lemma 9 in Ely and Peski (2006), which I extended as Lemma 3. The seminar participants at the HEC Economics Workshop 2005 in Honor of Andreu Mas-Colell, particularly Filipe Martins-da-Rocha, helped me locate the original references for the measurable selection theorems that I use. I also was helped by participants at the Centre d'Economie de la Sorbonne's Séminaire de Microéconomie et Décision—particularly Takashi Kamihigashi. This paper, whose results are used in Van Zandt and Vives (2006), originated from that project and I thank Xavier Vives for our discussions.

Contents

1	Introduction	1
2	Overview of the analysis	2
3	The interim formulation of the Bayesian game	3
4	Ex ante formulation of a Bayesian game	4
5	Topological and order structure of actions and strategies	5
6	The fixed-point theorem	8
7	Existence of the GBR mapping	10
8	Conclusion	15
	Appendix: Summary of order definitions	16
	References	17

1. Introduction

We prove the existence of a greatest and a least *interim* pure-strategy Bayesian Nash equilibrium for supermodular games of incomplete information. Here is a summary of the main result.

Consider the following:

- The *interim* formulation of a Bayesian game, with type spaces and with each individual's beliefs given by a mapping from her set of types to beliefs about the other players' sets of types and the state of nature.
(This is in contrast to the ex ante formulation, in which beliefs are given by a common prior and conditional beliefs.)
- The *interim* formulation of a Bayesian Nash equilibrium, in which each player and each type maximizes her expected payoff.
(This is in contrast to the ex ante formulation as the Nash equilibrium of an ex ante game, in which each player maximizes her expected payoff for almost every type.)

We place no assumptions on each set of types, other than that it is endowed with a sigma-field. Suppose that the following hold for each player:

1. Her set of actions is a compact metric lattice.
2. Her payoff is measurable in the types, continuous in the actions, bounded, and supermodular in own action, and has increasing differences between her own action and the other players' actions.
3. Her interim belief about each event is measurable in her own type.

Then there exist a greatest and a least pure-strategy equilibrium.

Vives (1990) and Milgrom and Roberts (1990) also obtain existence of a pure-strategy Bayesian Nash equilibrium for supermodular games of incomplete information. The main differences are that (a) we use the interim formulation of a Bayesian game, in which each player's beliefs are part of his or her type rather than being derived from a prior; (b) we use the interim formulation of a Bayesian Nash equilibrium, in which each player and every type (rather than almost every type) chooses a best response to the strategy profile of the other players; (c) we assume that the action spaces are compact metric lattices, whereas they assume that each action set is a compact sublattice of Euclidean space. The proof in Milgrom and Roberts (1990) applies a general existence theorem for supermodular games to the ex-ante normal form of the Bayesian game; the proof in Vives (1990) uses a Cournot tâtonnement that is also the basis of the proof in this paper.

Like those authors, we are exploiting the fact that the game is supermodular. Because there is uncertainty and the players are maximizing expected payoff, each player's underlying payoff function must be truly supermodular in own action and have increasing differences between own action and other actions (properties that are preserved by integration) rather than merely satisfy ordinal quasi-supermodularity and single-crossing properties (which are not preserved by integration). For example, a Bayesian game with log-supermodular payoffs need not have a pure-strategy equilibrium.

There are various papers on existence of BNE in *monotone-in-type* pure strategies, in which the type spaces are also partially ordered (unlike in this paper). Since monotone-in-type is a stronger property than we seek here, those papers need additional assumptions on complementarity between actions and types and on monotonicity of beliefs. Leaving these assumptions aside, here is how the models compare. Van Zandt and Vives (2006) use the same set-up as in this paper and rely on this paper's existence result. Athey (2001), McAdams (2003), and Reny (2006) obtain successively more general results for games that might not have strategic complementarities. In this way, the set of games may be more general. In the most general result, Reny (2006), the assumption on the action set is just slightly stronger than in this paper (the difference is that the lattice is locally complete). Reny's assumptions on types and beliefs are more restrictive: the type spaces are a cube in Euclidean space; the game is studied in its ex ante form; and the common prior must be atomless.

We can also compare our results with those for general Bayesian games (i.e., not necessarily supermodular). In the basic theory of static complete-information games, there are general theorems on the existence of mixed-strategy Nash equilibria; strategic complementarities allows one to obtain existence of pure-strategy equilibria. However, for Bayesian games, existence results for mixed-strategy (or behavioral strategy) equilibria require some technical restrictions. The state-of-the-art for games with continuous payoffs remains Milgrom and Weber (1985), who prove existence of equilibrium in behavioral strategies using an ex ante formulation of types, beliefs, and strategies. They assume that the type spaces are complete separable metric spaces and that the actions spaces are compact metric spaces, and then add an equicontinuity assumption on payoffs and an absolute continuity assumption on the prior with respect to the marginal distributions of types. These restrictions are not needed for the existence of pure-strategy Bayesian Nash equilibria in supermodular games.

The title of this paper refers to interim Bayesian Nash equilibrium on *universal type spaces*. The construction of universal type spaces in Mertens and Zamir (1985) and Brandenburger and Dekel (1993) yields (starting with a state space that is a complete separable metric space) type spaces that are complete separable metric spaces. In that sense, such type spaces are already covered by Milgrom and Weber (1985). Heifetz and Samet (1998) extend such construction to arbitrary measure spaces, but such type spaces are already covered by the existence proofs for supermodular games in Vives (1990) and Milgrom and Roberts (1990). However, intrinsic in the construction of universal type spaces is that there is no common prior and that the appropriate notion of equilibrium is interim, as in this paper. It is this combination of universal type spaces and interim Bayesian Nash equilibrium that is new to our existence proof, though it is made possible by the assumption that the game is supermodular and that such complementarity is common knowledge.

2. Overview of the analysis

We prove the existence of a greatest interim BNE equilibrium. The proof for a least interim BNE is omitted because it has analogous steps. In parallel, as a point of compar-

ison, we sketch a proof of existence of a greatest ex ante BNE (different from the proof in Vives (1990)).

We begin by defining a Bayesian game and a Bayesian Nash equilibrium, both for the interim model (Section 3) and the ex ante model (Section 4). We also explain the relationship between the two models.

Section 5 states the topological and order assumptions on actions (that each set of actions is a compact metric lattice) and the implications for each set of strategies (in particular, that the set of strategies is a lattice, that such lattice is typically not complete in the interim model, and that such lattice is complete in the ex ante model if actions sets are Euclidean).

We then proceed to the proofs of existence. Both proofs (interim and ex ante) have broadly two steps:

Step 1. Show, using complementarity assumptions and methods of optimization on lattices, that each player has a greatest best reply (GBR) and that the GBR mapping is an increasing function of the other players' strategies.

Step 2. Apply a lattice fixed-point theorem to the profile of GBR mappings.

We take up these steps in reverse order: Step 2 in Section 6 and Step 1 in Section 7.

Section 8 concludes. An Appendix contains standard order definitions and results, for convenience.

3. The interim formulation of the Bayesian game

We use an interim or incomplete-information formulation of a Bayesian game and of a Bayesian Nash equilibrium (BNE). This formulation is based on type-dependent beliefs (rather than a common prior and conditional beliefs) and interim best replies (rather than ex ante best replies).

The game has the following components.

1. The set of players is N , indexed by $i \in N$.
2. The set of types of player $i \in N$ is T_i , endowed with a sigma-field \mathcal{F}_i .
3. For some models, it is convenient to have a component of the state space that represents (a) residual uncertainty not observed by any player or (b) the set of possible payoffs, so that individual type spaces capture only beliefs. Denote this state space by T_0 and endow it with a sigma-field \mathcal{F}_0 .

Notation. Let $T := \prod_{k \in \{0\} \cup N} T_k$. For $i \in N$, let $T_{-i} := \prod_{k \neq i} T_k$ and let \mathcal{F}_{-i} be the product sigma-field $\bigotimes_{k \neq i} \mathcal{F}_k$.

4. Player i 's type-dependent beliefs are given by a function $p_i: T_i \rightarrow \mathcal{M}_{-i}$, where \mathcal{M}_{-i} is the set of probability measures on $(T_{-i}, \mathcal{F}_{-i})$.

Notation. We denote the probability of a set $F_{-i} \in \mathcal{F}_{-i}$, given beliefs $p_i(t_i)$, by $p_i(F_{-i} | t_i)$. However, the mapping p_i need not represent conditional beliefs derived from a prior on T .

5. The action set of player i is A_i , endowed with a sigma-field to be specified later.

Notation. The set of action profiles is $A := \prod_{i \in N} A_i$. Let $A_{-i} := \prod_{j \neq i} A_j$.

6. The payoff function of player i is $u_i: A \times T \rightarrow \mathbb{R}$.

Type spaces and action sets are non-empty.

A strategy for player i is a measurable function $\sigma_i: T_i \rightarrow A_i$. Let Σ_i denote the set of strategies for player i . Let $\Sigma := \prod_{i \in N} \Sigma_i$ denote the set of strategy profiles and let $\Sigma_{-i} := \prod_{j \neq i} \Sigma_j$ denote the profiles of strategies for players other than i . For notational simplicity, a strategy profile is viewed as a map from T to A , even though it does not depend on T_0 .

A BNE is a strategy profile σ such that each player and each type chooses a best response to the strategy profile of the other players. This is formalized as follows. When player i 's type is t_i and the strategy profile of the other players is σ_{-i} , her expected payoff from choosing action a_i is

$$\pi_i(a_i, t_i; \sigma_{-i}) := \int_{T_{-i}} u_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}) dp_i(t_{-i} | t_i). \quad (1)$$

The integral is over the possible types of the other players (and over T_0) given player i 's beliefs $p_i(\cdot | t_i)$. Let $\varphi_i(t_i; \sigma_{-i})$ be the set of actions for i that maximize this payoff:

$$\varphi_i(t_i; \sigma_{-i}) := \arg \max_{a_i \in A_i} \pi_i(a_i, t_i; \sigma_{-i}). \quad (2)$$

With this notation, we have the following definition of a Bayesian Nash equilibrium.

DEFINITION 1. A Bayesian Nash equilibrium is a strategy profile $\sigma \in \Sigma$ such that, for $i \in N$ and $t_i \in T_i$, $\sigma_i(t_i) \in \varphi_i(t_i; \sigma_{-i})$.

Let $\beta_i: \Sigma_{-i} \rightarrow \Sigma_i$ denote player i 's best-reply correspondence in terms of strategies:

$$\beta_i(\sigma_{-i}) = \{\sigma_i \in \Sigma_i \mid \forall t_i \in T_i: \sigma_i(t_i) \in \varphi_i(t_i; \sigma_{-i})\}. \quad (3)$$

Then, equivalently, a BNE is a strategy profile σ such that $\sigma_i \in \beta_i(\sigma_{-i})$ for $i \in N$.

4. Ex ante formulation of a Bayesian game

For comparison, we provide the following ex ante formulation of a Bayesian game and of a Bayesian Nash equilibrium.

The components of a game are the same, except that the belief mappings are replaced by a common prior μ on T . Strategies are taken to be equivalence classes, modulo being equal μ -a.e.

(We overload notation by using the same symbols to denote corresponding—but not identical—components of the two models. For example, in the interim model, σ_i is a function and Σ_i is a set of functions; in the ex ante model, σ_i is an equivalence class of functions and Σ_i is a set of such equivalence classes.)

In the interim formulation of BNE, each type of each player chooses an action in order to maximize expected utility, given the beliefs for that type. In the ex ante formulation, each player chooses a strategy before observing his type in order to maximize *unconditional* expected utility. That is, $\sigma \in \Sigma$ is an ex ante BNE if, for all $i \in N$, σ_i solves

$$\max_{\sigma'_i \in \Sigma_i} \int_T u_i(\sigma'_i(t_i), \sigma_{-i}(t_{-i}), t) d\mu(t). \quad (4)$$

The interim formulation of a Bayesian game and of a BNE is the correct one for interpreting a game as one of incomplete rather than imperfect information. Furthermore, as we will now remark, the class of interim games is broader than the class of ex ante games (with a mild restriction) and the notion of interim BNE is stronger than the notion of ex ante BNE. Therefore, the results of this paper are stronger than if we had used an ex ante formulation of Bayesian games and BNE (again, with a mild restriction).

Correspondence between interim and ex ante games. Consider an interim game and an ex ante game that have the same components (players, type spaces, utilities), except that the ex ante game has a common prior μ and the interim game has interim beliefs $\{p_i\}_{i \in N}$. We say that the games *correspond* if, for all $F_{-i} \in \mathcal{F}_{-i}$, $p_i(F_{-i} | t_i)$ is a conditional probability of F_{-i} given t_i for the prior μ . This means that p_i is a regular conditional probability for μ conditional on t_i .

The class of interim games is broader than the class of ex ante games, if we restrict attention to the case in which each T_k is a complete separable metric space and the sigma-algebra of T_k is its Borel field. Then, for any prior on T , a regular conditional probability given t_i always exists—that is, an ex ante game always has a corresponding interim game. (See, for example, Dellacherie and Meyer (1978, III.70 and 71).) The regular conditional probabilities are unique up to equivalence; we can take one member of the equivalence class for each player as that player’s beliefs in the interim game. The converse does not hold: an interim game need not have a corresponding ex ante game because a player’s beliefs in an interim game might not be consistent with a prior.

Interim BNE is stronger than ex ante BNE. Consider an ex ante game and an interim game that correspond. For any interim BNE of the interim game, the equivalence class of strategy profiles containing the interim BNE is an ex ante BNE of the ex ante game. The converse does not quite hold. For an ex ante BNE of the ex ante game, any member σ of the equivalence class is an “almost everywhere” interim BNE, meaning that, for every player i and μ -a.e. type $t_i \in T_i$, $\sigma_i(t_i)$ is a best response to σ_{-i} (whereas for an interim BNE, this should hold for *every* $t_i \in T_i$).

REMARK 1. The important commonality of a common prior is that players agree ex ante on which events have probability zero and they do not care what happens on such events. In our definition of an ex ante game, we could have allowed players to have different priors as long as these shared the same null sets, that is, were mutually absolutely continuous. However, this would be a false generalization. From such a game, we could construct an equivalent game with a common prior by letting the common prior be the prior μ_1 of player 1; then, if f_i is the Radon–Nikodým derivative of the prior μ_i of player $i \neq 1$ with respect to μ_1 , we redefine player i ’s utility to be $(a, t) \mapsto u_i(a, t)f_i(t)$.

5. Topological and order structure of actions and strategies

5.1. Assumptions on actions

We impose no further assumptions on the sets of types, but we assume that each set of actions is a lattice and has a compatible compact metrizable topology.

ASSUMPTION 1. For each player i , A_i is a compact metric lattice (its sigma-field is the Borel field).

We use the symbol \geq for all partial orders. Expressions such as “greater than” and “increasing” mean “*weakly* greater than” and “*weakly* increasing”. See the Appendix for standard definitions and results about partial orders and lattices that are used in this paper.

REMARK 2. We use the following properties of a compact metric lattice such as A_i (see, for example, Reny (2006)):

1. The binary operators \sup and \inf from $A_i \times A_i$ to A_i are continuous (this is what defines a topological lattice) and hence measurable.
2. A_i is a complete lattice.
3. Every increasing (resp., decreasing) sequence in A_i converges topologically to its order limit.
4. Any order interval in A_i is closed.

5.2. The order structure of strategies and best replies

Because A_i is a complete lattice, the product space $A_i^{T_i}$ (the set of all functions from T_i to A_i) is a complete lattice under the product (pointwise) partial order: “ $\sigma'_i \geq \sigma_i$ if and only if $\sigma'_i(t_i) \geq \sigma_i(t_i)$ for all $t_i \in T_i$ ”. The supremum of a subset of $A_i^{T_i}$ is the pointwise supremum; e.g., if $\sigma_i, \sigma'_i \in A_i$, then $\sigma_i \vee \sigma'_i$ is the function defined by $t_i \mapsto \sigma_i(t_i) \vee \sigma'_i(t_i)$.

For the set Σ_i of *measurable* functions to be a lattice, i.e., to be a sublattice of $A_i^{T_i}$, we just have to be sure that the pointwise supremum of two measurable functions is measurable. This is true because the lattice operations \sup and \inf are measurable.

For the *ex ante* model, Σ_i is also a lattice, but when defining the partial order we need to add the quantifier “for μ -a.e.”, as in “ $\sigma' \geq \sigma$ if and only if $\sigma'_i(t_i) \geq \sigma_i(t_i)$ for μ -a.e. $t_i \in T_i$ ”.

The greatest-best-reply (GBR) mapping for player i , if well-defined, is the function $\bar{\beta}: \Sigma_{-i} \rightarrow \Sigma_i$ given by

$$\bar{\beta}_i(\sigma_{-i}) := \max \beta_i(\sigma_{-i}).$$

(Such usage of “max” denotes the supremum of the set when this is also a member of the set.) Assuming that $\bar{\beta}_i$ is well-defined for $i \in N$, we define $\bar{\beta}: \Sigma \rightarrow \Sigma$ by $\bar{\beta}(\sigma) := \langle \bar{\beta}_i(\sigma_{-i}) \rangle_{i \in N}$. The greatest fixed point of $\bar{\beta}(\sigma)$, if it exists, is the greatest BNE.

With this notation, we can restate steps 1 and 2 of the proof as follows.

Step 1. Show that $\bar{\beta}_i$ is well-defined and increasing for $i \in N$.

Step 2. Show that $\bar{\beta}: \Sigma \rightarrow \Sigma$ has a fixed point.

5.3. Completeness of the strategy sets

Theorems on optimization and fixed points on a lattice commonly assume that the lattice is complete. In this paper, the lattices we need to worry about are Σ_i for $i \in N$.

For the ex ante model, completeness of Σ_i and hence of Σ holds if we assume that A_i is a compact sublattice of Euclidean space. This follows from Schaefer (1974, Proposition II.8.3); we state the proof of this fact because previous references to it are incomplete.

LEMMA 1. *Let $A \subset \mathbb{R}^n$ be a compact sublattice. Let $(\Omega, \mathcal{F}, \mu)$ be a probability space. Let V be the set of equivalence classes of measurable functions from $(\Omega, \mathcal{F}, \mu)$ to A . Let \geq be the partial order on V defined by “ $f \geq g$ if and only if $f(\omega) \geq g(\omega)$ μ -a.e.” Then V is a complete lattice.*

Proof. Let $L_1(\mathbb{R}^n)$ be the set of integrable functions from $(\Omega, \mathcal{F}, \mu)$ into \mathbb{R}^n , so that $V \subset L_1(\mathbb{R}^n)$.

That V is a lattice was explained in Section 5.2. To show that V is complete, we use Schaefer (1974, Proposition II.8.3), which states that every norm-bounded directed subset of $L_1(\mathbb{R}^n)$ is convergent in the L_1 norm. Let’s recall what this means.

A subset $D \subset V$ is *directed increasing* if, for all $f, g \in D$, there is $h \in D$ such that $h \geq f$ and $h \geq g$. We then say that D converges to $h \in D$ if, for all $\varepsilon > 0$, there is $f \in D$ such that, for $g \in D$ with $g \geq f$, we have $\|h - g\| < \varepsilon$. The definition of directed decreasing set is analogous. Observe that the limit of a convergent directed subset $D \subset L_1(\mathbb{R}^n)$ is the supremum of D .

Let $X \subset V$. Define, recursively, $D_0 = X$ and, for $k = 0, 1, \dots$,

$$D_{k+1} := \{f \vee g \mid f \in D_k, g \in D_k\}.$$

Let $D := \bigcup_{k=0}^{\infty} D_k$. Observe that D is directed increasing; in fact, $f \vee g \in D$ for all $f, g \in D$. Therefore, D is convergent; let h be the limit. Recall that h is the supremum of D . Observe that an upper bound on X is also an upper bound on D and vice versa. Therefore, the supremum of D is the supremum of X . \square

Typically Σ_i *not complete* in the interim model. To see this, suppose that $A_i = \{0, 1\}$, so that all functions from T_i to A_i are indicator functions. Suppose also that all singleton subsets of T_i are measurable and there is a subset $F \subset T_i$ that is not measurable. (For example, $T_i = [0, 1]$, with the Borel field.) Since each singleton is measurable, $\{\mathbf{1}_{\{t_i\}} \mid t_i \in F\}$ is a subset of Σ_i . The obvious candidate for its supremum is the pointwise supremum $\mathbf{1}_F$, but this is not measurable. Any measurable upper bound $\mathbf{1}_G$ is such that $F \subsetneq G$. Then, letting $G' := G \setminus \{a_i\}$ for some $a_i \in G \setminus F$, $\mathbf{1}_{G'}$ is another measurable upper bound that is less than $\mathbf{1}_G$. It follows that the set $\{\mathbf{1}_{\{t_i\}} \mid t_i \in F\}$ of functions has no supremum in Σ_i . Note, however, that any countable set has a supremum and infimum. This fact will be used implicitly in what follows.

The incompleteness of Σ_i is the only distinction between the two models in step 2 of the proof; it requires that, in the interim model, we use an alternate fixed-point theorem that does not assume completeness of the lattice. For step 1, there is another distinction. In the ex ante model, the GBR is a solution to a single optimization problem; in the interim model, the GBR is pieced together from the solutions to many problems, one for each type. This “piecing together” has to be done in a measurable way, which turns out to be tricky—the more so because of the incompleteness of Σ_i as a lattice.

6. The fixed-point theorem

We begin with step 2, which is easier. For now, then, we take as given that $\bar{\beta}_i$ is well-defined and increasing.

6.1. Ex ante model

Consider first the ex ante model. We assume that Σ_i is complete with the following assumption.

ASSUMPTION 2. (Ex ante) For $i \in N$, A_i is a compact sublattice of Euclidean space. The σ -field on A_i is the Borel field.

We could apply Tarski's fixed-point theorem, but we state a related result, due to Abian and Brown (1961), for partially ordered sets. It is more general than what we need, but it provides a better point of comparison with our approach for the interim model.

Let (X, \geq) be a partially-ordered set. A *chain* C is a totally-ordered subset of X . X is *chain-complete* if every chain in X has a supremum. E.g., a complete lattice is chain-complete.

THEOREM 1. (Abian–Brown¹) Suppose that X is a chain-complete partially-ordered set, that $f: X \rightarrow X$ is an increasing function, and that there is $x \in X$ such that $x \leq f(x)$ (e.g., X has a least element). Then f has a greatest fixed point.

COROLLARY 1. (Ex ante) Assume that $\bar{\beta}_i$ is well-defined and increasing for $i \in N$. Then $\bar{\beta}$ has a greatest fixed point and hence the ex ante game has a greatest BNE.

6.2. Interim model

For the interim model, we cannot apply Theorem 1 because Σ_i is not chain complete. However, Σ_i is ω -chain complete; i.e., countable chains have a supremum. It is therefore possible to construct a fixed point using a countable iterative method, yielding the Tarski–Kantorovich Theorem (see Granas and Dugundji (2003, Theorem 2.1.2)).² We state this as Theorem 2 below and provide both intuition and a full proof in order (a) to highlight the proof's practical iterative nature (amenable to numerical calculation, similar to the computation of a fixed point of a contraction mapping) and (b) to show its similarity to the Cournot tâtonnement used in Vives (1990, Theorem 6.1) and Van Zandt and Vives (2006, Lemma 6).

The iteration starts at the greatest element (the greatest strategy profile); one obtains

1. See Markowsky (1976) for a converse. Abian and Brown assume that each well-ordered subset, rather than each totally-ordered subset, has a supremum. However, these two conditions are equivalent because a totally-ordered set has a well-ordered cofinal subset.

2. Heikkilä and Reffett (2006a, 2006b) have related fixed-point theorems that do not assume chain completeness.

a decreasing sequence (of best responses to the previous profile in the sequence). One has to show that the limit of this sequence is a fixed point (a BNE). To obtain the result, one assumes that decreasing sequences have an infimum and that the function (the GBR mapping) is “continuous” with respect to such limits. (These are *assumptions* in the general statement of the fixed-point theorem; they are *results* for the model in this paper to which we apply the fixed-point theorem.)

Note another difference between the Abian–Brown and Tarski–Kantorovich theorems. To merely obtain *existence* of a fixed point, we could use an increasing tatônnement. To get started, we would need an x such that $x \leq f(x)$. We would need a countable version of chain-completeness of X ensure that the tatônnement converged. However, to find the *greatest* fixed point, we have to start at the top and move downward.

DEFINITION 2. A partially ordered set is *downward sequentially complete* if every decreasing sequence has an infimum.

Suppose (X, \geq) and (Y, \geq) are partially-ordered sets that are downward sequentially complete. An increasing function $f: X \rightarrow Y$ is *downward sequentially continuous* if, for every decreasing sequence $\{x_1, x_2, \dots\}$, $\inf\{f(x_n)\} = f(\inf\{x_n\})$.

THEOREM 2. (Tarski–Kantorovich) *Suppose that (X, \geq) is a partially-ordered set that is downward sequentially complete and that has a greatest element. Suppose that $f: X \rightarrow X$ is increasing and downward sequentially continuous. Then f has a greatest fixed point.*

Proof. Let x_0 be the greatest element of X . Define recursively, for $n = 1, 2, \dots$, $x_n = f(x_{n-1})$. Since x_0 is the greatest element of X , $x_1 \leq x_0$. Since f is increasing, $f(x_1) \leq f(x_0)$, i.e., $x_2 \leq x_1$. By induction, $x_n \leq x_{n-1}$ for all n . Since (X, \geq) is downward sequentially complete, the decreasing sequences $\{x_n\}$ and $\{f(x_n)\}$ have an infimum x^* (the same for both since the sequences are the same except that $\{x_n\}$ has an extra first term). Since f is downward sequentially continuous, $f(x^*) = \inf\{f(x_n)\}$, i.e., $f(x^*) = x^*$. Hence, x^* is a fixed point of f .

Suppose that x' is another fixed point of f . Since $x' \leq x_0$, $f(x') \leq f(x_0)$, i.e., $x' \leq x_1$. By induction, $x' \leq x_n$ for all n . Therefore, x' is a lower bound on $\{x_n\}$, so $x' \leq x^*$. Thus, x^* is the greatest fixed point of f . \square

LEMMA 2. (Interim) Σ_i is downward sequentially complete.

Proof. Let $\{\sigma_i^n\}$ be a decreasing sequence in Σ_i . Then, for all t_i , $\sigma_i^*(t_i)$ is a decreasing sequence; hence, it is order and topologically convergent (see Remark 2). Let $\sigma_i^*(t_i)$ be its limit. σ_i^* is the infimum of $\{\sigma_i^n\}$ in $A_i^{T_i}$. We have only to show that σ_i^* is measurable and hence in Σ_i so that it is also the infimum of $\{\sigma_i^n\}$ in Σ_i . This is true because the pointwise limit of a sequence of measurable functions from a measure space to a metric space is measurable. \square

REMARK 3. Lemma 2 does not use A_i 's lattice structure. Instead, each A_i need only be a downward sequentially complete partially-ordered topological space.

COROLLARY 2. (Interim) *Suppose that $\bar{\beta}_i$ is well-defined, increasing, and downward sequentially continuous for $i \in N$. Then $\bar{\beta}$ has a greatest fixed point and hence the interim game has a greatest BNE.*

7. Existence of the GBR mapping

Working backward, we take up step 1, showing that the GBR mapping $\bar{\beta}$ is well-defined and increasing.

7.1. Assumptions

We impose the following assumption for each player in both the interim and the ex ante models.

ASSUMPTION 3. For each player i , $u_i: A \times T \rightarrow \mathbb{R}$

1. is continuous in a ,
2. is measurable in t ,
3. is bounded,
4. is supermodular in a_i ,
5. has increasing differences in (a_i, a_{-i}) .

We also assume the following in the interim model.

ASSUMPTION 4. For each player i , for $F_{-i} \in \mathcal{F}_{-i}$, $t_i \mapsto p_i(F_{-i} | t_i)$ is measurable.

If the interim game has a corresponding ex ante game, then the beliefs in the interim game satisfy this assumption.

7.2. Overview

Fix $i \in N$. That $\bar{\beta}_i$ (if well defined) is decreasing in σ_{-i} is a straightforward consequence of the fact that u_i is supermodular in a_i and has increasing differences in (a_i, a_{-i}) . This is true for both the ex ante and the interim models. The details are standard, so we do not repeat them.

The challenge is to show that $\bar{\beta}_i$ is well-defined, with such challenge due mainly to measurability problems.

However the result is straightforward in the ex ante model. Recall that Σ_i is a complete lattice given the assumption that A_i is a compact sublattice of Euclidean space (Assumption 2). The ex ante utility function, defined as the expected utility given a profile of strategies, is order upper semicontinuous and supermodular given the assumptions on u_i . Then Lemma A.1 immediately tells us that the set of best replies is a complete lattice and that greatest and least ex ante best replies exist. We also omit the details of this now-standard proof.

For the rest of this section, we show the result for the interim model. To simplify notation, we fix $\sigma_{-i} \in \Sigma_{-i}$; we write $\pi_i(a_i, t_i)$ instead of $\pi_i(a_i, t_i; \sigma_{-i})$ for type t_i 's interim expected payoff function; and we write $\varphi_i(t_i)$ instead of $\varphi_i(t_i; \sigma_{-i})$ for type t_i 's optimal actions given σ_{-i} .

We begin, in Section 7.3, by showing that $\bar{\varphi}_i(t_i) := \max \varphi_i(t_i)$ is well-defined for each t_i . Then we have to show that $t_i \mapsto \bar{\varphi}_i(t_i)$ is measurable, so that this strategy is thus $\bar{\beta}_i(\sigma_{-i})$. There are two main steps, given in Section 7.5:

- We use the Measurable Maximum Theorem to show that $\varphi_i: T_i \rightarrow A_i$ is a measurable correspondence.
- Then φ_i has a Castaing representation: a set of measurable selections that is point-wise dense in φ_i . From the Castaing representation, we can construct the GBR in a measurable way.

To apply the Measurable Maximum Theory, we have to demonstrate that π_i is measurable in t_i ; this we do in Section 7.4.

7.3. Each type has a greatest best reply

We check that the interim payoff function $\pi_i(a_i, t_i)$ is well-defined, i.e., that the function that is integrated in equation (1) is measurable and integrable. We do not constrain the beliefs of the player, so integrability is assured by the assumption (3.3) that u_i is bounded.

PROPOSITION 1. *Let $a_i \in A_i$ and $t_i \in T_i$. Then*

$$t_{-i} \mapsto u_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i})$$

is measurable and bounded. Hence, $\pi_i(a_i, t_i)$ is well-defined.

Proof. Since σ_{-i} is measurable and the composition of measurable functions is measurable,

$$t_{-i} \mapsto u_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i})$$

is measurable. This function is also bounded because u_i is bounded. Thus, its expected value $\pi_i(a_i, t_i)$ with respect to the measure $p_i(t_i)$ on T_{-i} is well-defined. \square

The existence of a greatest best response then follows from the supermodularity of u_i in a_i .

PROPOSITION 2. *π_i is supermodular and continuous in a_i . Therefore, for $t_i \in T_i$, $\varphi_i(t_i)$ is a non-empty complete lattice and has a greatest element $\bar{\varphi}_i(t_i)$.*

Proof. Supermodularity and continuity are preserved by integration. (Continuity is also shown in the proof of Proposition 4.) Then apply Lemma A.1. \square

7.4. Interim expected utility is measurable in type

PROPOSITION 3. $\pi_i: A_i \times T_i \rightarrow \mathbb{R}$ is measurable in t_i .

Proof. Fix $a_i \in A_i$ and define $U_i(t_i, t_{-i}) := u_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i})$, so that

$$\pi_i(a_i, t_i) = \int_{T_{-i}} U_i(t_i, t_{-i}) dp_i(t_{-i} | t_i).$$

That is, $\pi_i(a_i, t_i)$ is the expected value of $U_i(t_i, t_{-i})$ when t_i is known and the probability measure on t_{-i} is $p_i(t_i)$.

Ely and Peski (2006, Lemma 9) show that such an expectation is measurable in t_i when (a) U_i is jointly measurable, (b) T_{-i} is a complete separable metric space and \mathcal{F}_{-i} is its Borel field, and (c) $p_i: T_i \rightarrow \mathcal{M}_{-i}$ is measurable when \mathcal{M}_{-i} is endowed with the Borel field of the topology of convergence in measure.

We state and prove, as Lemma 3 below, a version of their lemma that fits our slightly more general assumptions: (a) U_i is jointly measurable, (b) $(T_{-i}, \mathcal{F}_{-i})$ is an arbitrary measurable space, and (c) $t_i \mapsto p_i(F_{-i} | t_i)$ is measurable for all $F_{-i} \in \mathcal{F}_{-i}$.

To apply Lemma 3 to this proof of Proposition 3, we let (X, \mathcal{F}) of the lemma be (T_i, \mathcal{F}_i) ; we let (Y, \mathcal{G}) be $(T_{-i}, \mathcal{F}_{-i})$; we let $U: X \times Y \rightarrow \mathbb{R}$ be $U_i: T_i \times T_{-i} \rightarrow \mathbb{R}$; we let \mathcal{M} be \mathcal{M}_{-i} ; we let $p: X \rightarrow \mathcal{M}$ be $p_i: T_i \rightarrow \mathcal{M}_{-i}$; and we let $\pi(x) := \int_Y U(x, y) dp(y|x)$ be $\pi_i(a_i, t_i) := \int_{T_{-i}} U_i(t_i, t_{-i}) dp_i(t_{-i} | t_i)$. Assumption 4 is that $t_i \mapsto p_i(F_{-i} | t_i)$ is measurable. Furthermore, we have shown that U_i is bounded and measurable (see the proof of Proposition 1). Therefore, according to Lemma 3, $t_i \mapsto \pi_i(a_i, t_i)$ is \mathcal{F}_i -measurable. \square

LEMMA 3. Let (X, \mathcal{F}) and (Y, \mathcal{G}) be two measurable spaces. Let $U: X \times Y \rightarrow \mathbb{R}$ be $\mathcal{F} \otimes \mathcal{G}$ -measurable and bounded. Let \mathcal{M} be the set of probability measures on (Y, \mathcal{G}) . Let $p: X \rightarrow \mathcal{M}$. Suppose that, for $G \in \mathcal{G}$, $x \mapsto p(G|x)$ is \mathcal{F} -measurable. Then $\pi: X \rightarrow \mathbb{R}$, defined by

$$\pi(x) := \int_Y U(x, y) dp(y|x),$$

is \mathcal{F} -measurable.

Proof of Lemma 3. We build up the result by showing that it holds for successively larger classes of U .

Lemma 3 holds if U is an indicator function on a measurable rectangle. Suppose $H \in \mathcal{F} \otimes \mathcal{G}$ is a rectangle, equal to $F \times G$ for some $F \in \mathcal{F}$ and $G \in \mathcal{G}$. Then $\int \mathbf{1}_H dp(y|x) = \mathbf{1}_F p(G|x)$, which is a measurable function of x because $p(G|x)$ is assumed to be a measurable function of x .

Lemma 3 holds if U is any indicator function. This is the most difficult part of the proof. We use the following lemma, which is related to the monotone class theorem. See, for example, Billingsley (1995, p. 42) for a proof of this result.

LEMMA 4. Let \mathcal{P} and \mathcal{L} be two sets of subsets of a set Z . Suppose \mathcal{P} has the following property:

(μ) $A, B \in \mathcal{P} \Rightarrow A \cap B \in \mathcal{P}$.

Suppose \mathcal{L} has the following properties:

(λ_1) $Z \in \mathcal{L}$.

(λ_2) $A \in \mathcal{L} \Rightarrow A^c \in \mathcal{L}$.

(λ_3) If $A_1, A_2, \dots \in \mathcal{L}$ are pairwise disjoint, then $\bigcup_{n=1}^{\infty} A_n \in \mathcal{L}$.

Suppose also that $\mathcal{P} \subset \mathcal{L}$. Then $\sigma(\mathcal{P}) \subset \mathcal{L}$, where $\sigma(\mathcal{P})$ is the sigma-algebra generated by \mathcal{P} .

We apply Lemma 4 as follows. Let $Z = X \times Y$. Let \mathcal{P} be the set of all measurable rectangles in $\mathcal{F} \otimes \mathcal{G}$. Then \mathcal{P} satisfies property (μ). Let \mathcal{L} be the class of sets in $\mathcal{F} \otimes \mathcal{G}$ for which π is \mathcal{F} -measurable when U is an indicator function on the set. Clearly \mathcal{L} satisfies properties λ_1 and λ_2 . It also satisfies property λ_3 , as follows.

For any $H \in \mathcal{F} \otimes \mathcal{G}$, denote by π_H the function π when $U = \mathbf{1}_H$. Let $A_1, A_2, \dots \in \mathcal{L}$ be pairwise disjoint. Let $A_\infty := \bigcup_{n=1}^{\infty} A_n$. By the definition of \mathcal{L} , π_{A_n} is measurable for $n \in \mathbb{N}$. We have to show that π_{A_∞} is measurable. For $m = 1, 2, \dots$, let $B_m := \bigcup_{n=1}^m A_n$. Since $\mathbf{1}_{B_m} = \sum_{n=1}^m \mathbf{1}_{A_n}$, we have that $\pi_{B_m} = \sum_{n=1}^m \pi_{A_n}$ and hence π_{B_m} is also measurable. We now show that π_{A_∞} is the pointwise limit of the sequence $\{\pi_{B_m}\}_{m \in \mathbb{N}}$ and hence is measurable. Fix $x \in X$ and consider $\mathbf{1}_{B_m}$ as a function of y . Then $\{\mathbf{1}_{B_m}\}_{m \in \mathbb{N}}$ is an increasing sequence of \mathcal{G} -measurable functions that converges pointwise to $\mathbf{1}_{A_\infty}$. Therefore, $\{\pi_{B_m}(x) = \int_Y \mathbf{1}_{B_m} dp(y|x)\}_{m \in \mathbb{N}}$ is an increasing sequence that converges to $\pi_{A_\infty}(x) = \int_Y \mathbf{1}_{A_\infty} dp(y|x)$. We have shown that $\{\pi_{B_m}\}$ is an increasing sequence of \mathcal{F} -measurable functions on X that converges pointwise to π_{A_∞} , as desired. Therefore, π_{A_∞} is \mathcal{F} -measurable and $A_\infty \in \mathcal{L}$.

Earlier we showed that π_A is measurable for any measurable rectangle A ; that is, $\mathcal{P} \subset \mathcal{L}$. Therefore $\sigma(\mathcal{P}) \subset \mathcal{L}$. Since $\sigma(\mathcal{P}) = \mathcal{F} \otimes \mathcal{G}$, we have that $\mathcal{L} = \mathcal{F} \otimes \mathcal{G}$. That is, Lemma 3 holds if U is any indicator function.

Lemma 3 holds if U is simple. Suppose U is simple and equal to $\sum_{n=1}^m \alpha_n \mathbf{1}_{H_n}$, where $\{H_1, \dots, H_m\}$ is an $\mathcal{F} \otimes \mathcal{G}$ -measurable partition of $X \times Y$. Then

$$\pi(x) = \sum_{n=1}^m \alpha_n \int \mathbf{1}_{H_n} dp(y|x),$$

which is \mathcal{F} -measurable since we have shown that each function $x \rightarrow \int \mathbf{1}_{H_n} dp(y|x)$ is \mathcal{F} -measurable.

Lemma 3 holds for all measurable U . Since U is measurable, there is a sequence $\{U_n\}_{n \in \mathbb{N}}$ of simple functions converging pointwise to U . We can take all functions in this sequence to have the same bound as U . Then, for each x , $\{y \mapsto U_n(x, y)\}_{n \in \mathbb{N}}$ is a uniformly bounded sequence of functions on Y that converges pointwise to $y \mapsto U(x, y)$. It follows from dominated convergence that

$$\lim_{n \rightarrow \infty} \int_Y U_n(x, y) dp(y|x) = \int_Y U(x, y) dp(y|x).$$

Therefore, if we define $\pi_n(x) := \int_Y U_n(x, y) dp(y|x)$, then $\{\pi_n\}_{n \in \mathbb{N}}$ is a sequence of functions on X that converges pointwise to π . Each π_n is measurable since U_n is simple; therefore π is measurable.

This concludes the proof of Lemma 3. \square

7.5. The type-to-GBR selection is measurable

We thus have shown that π is a Carathéodory function: continuous in a_i and measurable in t_i . Since we also assume that A_i is a compact metric space, it follows from the Measurable Maximum Theorem (e.g., Aliprantis and Border (1999, Theorem 17.18)) that the solution correspondence

$$t_i \longmapsto \arg \max_{\varphi_i} \pi_i(a_i, t_i)$$

is measurable. (Let (Y, \mathcal{G}) be a measurable space and let X be a topological space. A correspondence $\varphi: Y \rightarrow X$ is *measurable* if $\{y \in Y \mid \varphi(y) \cap F \neq \emptyset\} \in \mathcal{G}$ for all closed $F \subset X$.)

Therefore, φ has a Castaing representation (see Castaing and Valadier (1977, Chapter 3)). We construct from it the greatest best reply in a measurable way. We state this step as an independent result.

LEMMA 5. *Let (Y, \mathcal{G}) be a measurable space and let X be a compact metric lattice. Let $\varphi: Y \rightarrow X$ be a measurable correspondence with values that are non-empty, topologically closed, and lattices. Then, for all $y \in Y$, $\varphi(y)$ contains a greatest element $\bar{\varphi}(y)$ and the function $\bar{\varphi}: Y \rightarrow X$ is measurable.*

Proof. We have that φ is measurable, X is a Polish space, and φ has non-empty closed values. Therefore, φ has a Castaing representation: A sequence $\{f_n: Y \rightarrow X\}_{n \in \mathbb{N}}$ of measurable selections such that $\varphi(y) = \text{cl} \{f_n(y) \mid n \in \mathbb{N}\}$ for all $y \in Y$.

Define recursively, for $n \in \mathbb{N}$, $\bar{f}_n(y) = \sup\{f_n(y), \bar{f}_{n-1}(y)\}$ for $y \in Y$. Since $\varphi(y)$ is a lattice, \bar{f}_n is a selection of φ . Because \sup is measurable for a topological lattice, \bar{f}_n is measurable. Since $\{\bar{f}_n(y)\}_{n \in \mathbb{N}}$ is an increasing sequence, it converges topologically to its least upper bound $\bar{f}(y)$, which is in $\varphi(y)$ since $\varphi(y)$ is closed. Since $\bar{f}: Y \rightarrow X$ is the pointwise limit of $\{\bar{f}_n \mid n \in \mathbb{N}\}$, it is measurable—that is, it is a measurable selection of φ . We have left to show that $\bar{f}(y) = \sup \varphi(y)$.

Since order intervals are closed, $\{x \in X \mid x \leq \bar{f}(y)\}$ is a closed set that contains the dense subset $\{f_n(y) \mid n \in \mathbb{N}\}$ of $\varphi(y)$ and thus contains $\varphi(y)$. Thus, $\bar{f}(y)$ is an upper bound on $\varphi(y)$ and is thus the greatest element of $\varphi(y)$. \square

COROLLARY 3. *$t_i \rightarrow \bar{\varphi}_i(t_i)$ is a measurable selection of φ_i , so it is the greatest element of $\beta_i(\sigma_{-i})$.*

7.6. The GBR mapping is downward sequentially continuous

First we establish (topological) sequential continuity of $\pi_i(a_i, t_i; \sigma_{-i})$ in both a_i and σ_{-i} .

PROPOSITION 4. *Let $i \in N$ and $t_i \in T_i$. If $\{a_i^n\}_{n \in \mathbb{N}}$ and $\{\sigma_{-i}^n\}_{n \in \mathbb{N}}$ are sequences in A_i and Σ_{-i} , respectively, such that $\{a_i^n\}_{n \in \mathbb{N}}$ converges to a_i^∞ and $\{\sigma_{-i}^n\}_{n \in \mathbb{N}}$ converges*

pointwise (for all t_{-i}) to σ_{-i}^∞ , then

$$\lim_{n \rightarrow \infty} \pi_i(a_i^n, t_i; \sigma_{-i}^n) = \pi_i(a_i^\infty, t_i; \sigma_{-i}^\infty). \quad (5)$$

In particular, $\pi_i(a_i, t_i; \sigma_{-i})$ is continuous in a_i .

Proof. Because u_i is continuous and bounded, the sequence of measurable functions

$$t_{-i} \mapsto u_i(a_i^n, \sigma_{-i}^n(t_{-i}), t_i, t_{-i})$$

is bounded and converges pointwise to

$$t_{-i} \mapsto u_i(a_i^\infty, \sigma_{-i}^\infty(t_{-i}), t_i, t_{-i}).$$

Thus, by dominated convergence, their integral with respect to $p_i(t_i)$ converges:

$$\lim_{n \rightarrow \infty} \int_{T_{-i}} u_i(a_i^n, \sigma_{-i}^n(t_{-i}), t_i, t_{-i}) dp_i(t_{-i} | t_i) = \int_{T_{-i}} u_i(a_i^\infty, \sigma_{-i}^\infty(t_{-i}), t_i, t_{-i}) dp_i(t_{-i} | t_i).$$

This is equation (5). □

PROPOSITION 5. $\bar{\beta}_i$ is downward sequentially continuous.

Proof. Let $\{\sigma_{-i}^n\}_{n \in \mathbb{N}}$ be a decreasing sequence in Σ_{-i} , with infimum σ_{-i} . Then $\{\bar{\beta}_i(\sigma_{-i}^n)\}$ is a decreasing sequence (since $\bar{\beta}_i$ is increasing); let σ_i be its infimum. The fact that $\sigma_i \in \beta_i(\sigma_{-i})$ follows from the sequential continuity of π_i shown in Proposition 4. Specifically, for each $t_i \in T_i$ and each $a_i \in A_i$, $\pi_i(a_i, t_i; \sigma_{-i}^n) \leq \pi_i(\bar{\beta}_i(\sigma_{-i}^n)(t_i), t_i; \sigma_{-i}^n)$, and such inequality is preserved in the limit: $\pi_i(a_i, t_i; \sigma_{-i}) \leq \pi_i(\sigma_i(t_i), t_i; \sigma_{-i})$. Thus, $\sigma_i(t_i)$ is an optimal response by type i to σ_{-i} for each type $t_i \in T_i$. Hence, $\sigma_i \in \beta_i(\sigma_{-i})$. We just have to show that $\sigma_i \geq \bar{\beta}_i(\sigma_{-i})$, so that indeed σ_i is the greatest best response and is equal to $\bar{\beta}_i(\sigma_{-i})$. For all n , $\sigma_{-i} \leq \sigma_{-i}^n$ and hence $\bar{\beta}_i(\sigma_{-i}) \leq \bar{\beta}_i(\sigma_{-i}^n)$. Thus, $\bar{\beta}_i(\sigma_{-i})$ is also a lower bound on $\{\bar{\beta}_i(\sigma_{-i}^n)\}$ and hence is less than or equal to the greatest lower bound σ_i . □

8. Conclusion

We have thus proved the following.

THEOREM 3. Consider the interim formulation of a Bayesian game as stated in Section 3. Assume the following for each player i :

1. A_i is a compact metric lattice.
2. u_i is bounded, is measurable in t_i , is continuous and supermodular in a_i , and has increasing differences in (a_i, a_{-i}) .
3. For $F_{-i} \in \mathcal{F}_{-i}$, $t_i \mapsto p_i(F_{-i} | t_i)$ is measurable.

Then the game has a greatest and a least interim Bayesian Nash equilibrium.

In Section 7, we showed that the greatest-best-reply, $\bar{\beta}_i$, is well defined for each player i ; furthermore, it is an increasing function and it is decreasing sequentially continuous. (This part used all the assumptions of Theorem 3.) Therefore, according to Corollary 2, $\bar{\beta}$ has a greatest fixed point; this is the greatest equilibrium of the game.

In games of perfect information with quasi-supermodular payoffs and strategic complementarities, it is possible to dispense with the assumption that the action set is topologically compact and rely instead on the completeness of the lattice to ensure that each player always has a best response. We can also rely on completeness of the lattice to obtain that each type has a best response. However, the compactness of the strategy sets played a role when showing that the greatest best response is measurable in type.³ Specifically, it is needed to apply the Measurable Maximum Theorem to ensure that the type-to-best-replies correspondence is measurable. Without compactness, it is still easy to show that such correspondence has a measurable graph; however, this weaker condition is not sufficient to obtain a Castaing representation of the correspondence. That said, it remains an open question whether there might be alternative way to relax topological compactness.

Appendix: Summary of order definitions

For the convenience of the reader and to fix some notation and terminology that may vary from author to author, we include a few definitions about order and lattices. Throughout, we use terms like “greater than” and “increasing” to mean “weakly greater than” and “weakly increasing”.

Let (X, \geq) be a partially ordered set.

Let $D \subset X$. The greatest and least elements of D , when they exist, are denoted $\max D$ and $\min D$, respectively. A *supremum* (resp., *infimum*) of D is a least upper bound (resp., greatest lower bound); it is denoted $\sup D$ (resp., $\inf D$).

Let (T, \geq) be another partially ordered set. A function $f: X \rightarrow T$ is *increasing* if, for x, y in X , $x \geq y$ implies that $f(x) \geq f(y)$.

A functional $g: X \times Y \rightarrow \mathbb{R}$ has *increasing differences* in (x, t) if $g(x', t) - g(x, t)$ is increasing in t for $x' > x$ or, equivalently, if $g(x, t') - g(x, t)$ is increasing in x for $t' > t$.

(X, \geq) is a *lattice* if any two elements have a supremum and an infimum. A lattice (X, \geq) is *complete* if every non-empty subset has a supremum and an infimum.

A functional $g: X \rightarrow \mathbb{R}$ on a lattice X is *supermodular* if, all x, y in X , $g(\inf(x, y)) + g(\sup(x, y)) \geq g(x) + g(y)$.

Supermodularity is a stronger property than increasing differences: If T is also a lattice and if g is supermodular on $X \times T$, then g has increasing differences in (x, t) .

If X is the product of linearly ordered sets X_1, \dots, X_k , then X is a lattice and $g: X \rightarrow \mathbb{R}$ is supermodular if and only if g has increasing differences in (x_i, x_j) for $i \neq j$.

A *chain* $C \subset X$ is a totally ordered subset of X . A function $f: X \rightarrow \mathbb{R}$ is *order upper semicontinuous* if $\lim_{x \in X, x \downarrow \inf(C)} f(x) \leq f(\inf(C))$ and $\lim_{x \in X, x \uparrow \sup(C)} f(x) \geq f(\sup(C))$ for

3. Sigma-compactness would suffice.

any chain C .

The main comparative-statics tool applied in this paper is the following; see Milgrom and Roberts (1990, Section 1).

LEMMA A.1. *Let X be a complete lattice and let T be a partially ordered set. Let $u: X \times T \rightarrow \mathbb{R}$ be a function that is supermodular and upper order continuous on the lattice X for each $t \in T$. Let $\varphi(t) = \arg \max_{x \in X} u(x, t)$. Then $\varphi(t)$ is a non-empty complete sublattice for all t ; hence $\max \varphi(t)$ and $\min \varphi(t)$ exist.*

Assume also that u has increasing differences in (x, t) . Then $t \mapsto \sup \varphi(t)$ and $t \mapsto \inf \varphi(t)$ are increasing selections of φ .

References

- Abian, Smbat and Arthur B. Brown (1961). A Theorem on Partially Ordered Sets with Applications to Fixed Point Theorems. *Canadian Journal of Mathematics*, 13, 78–82.
- Aliprantis, Charalambos D. and Kim C. Border (1999). *Infinite Dimensional Analysis*. Berlin: Springer-Verlag.
- Athey, Susan (2001). Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information. *Econometrica*, 69, 861–889.
- Billingsley, Patrick (1995). *Probability and Measure*. New York: John Wiley & Sons.
- Brandenburger, Adam and Eddie Dekel (1993). Hierarchies of Beliefs and Common Knowledge. *Journal of Economic Theory*, 59, 189–198.
- Castaing, Charles and Michel Valadier (1977). *Convex Analysis and Measurable Multifunctions*. Berlin: Springer-Verlag.
- Dellacherie, Claude and Paul-Andr e Meyer (1978). *Probabilities and Potential*. Amsterdam: North-Holland.
- Ely, Jeffrey and Marcin Peski (2006). Hierarchies of Beliefs and Interim Rationalizability. *Theoretical Economics*, 1, 19–65.
- Granas, Andrzej and James Dugundji (2003). *Fixed Point Theory*. Springer-Verlag.
- Heifetz, Aviad and Dov Samet (1998). Topology-Free Typology of Beliefs. *Journal of Economic Theory*, 82, 324–341.
- Heikkil a, Seppo and Kevin Reffett (2006a). Fixed Point Theorems and Their Applications to Theory of Nash Equilibria. *Nonlinear Analysis*, 64(1415–1436).
- Heikkil a, Seppo and Kevin Reffett (2006b). Fixed Point Theorems for Multivalued Mappings in Ordered Topological Spaces with Applications to Game Theory. University of Oulu and Arizona State University.
- Markowsky, George (1976). Chain-Complete Posets and Directed Sets with Applications. *Algebra Univers*, 6, 53–68.
- McAdams, David (2003). Isotone Equilibrium in Games of Incomplete Information. *Econometrica*, 71, 1191–1214.
- Mertens, Jean-Francois and Shmuel Zamir (1985). Formulation of Bayesian Analysis for Games with Incomplete Information. *International Journal of Game Theory*, 14, 1–29.
- Milgrom, Paul and John Roberts (1990). Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities. *Econometrica*, 58, 1255–1277.
- Milgrom, Paul and Robert Weber (1985). Distributional Strategies for Games with Incomplete Information. *Mathematics of Operations Research*, 10, 619–632.
- Reny, Philip J. (2006). On the Existence of Monotone Pure Strategy Equilibria in Bayesian Games. University of Chicago.
- Schaefer, Helmut H. (1974). *Banach Lattices and Positive Operators*. Berlin: Springer-Verlag.

Van Zandt, Timothy and Xavier Vives (2006). Monotone Equilibria in Bayesian Games of Strategic Complementarities. *Journal of Economic Theory*. Forthcoming.

Vives, Xavier (1990). Nash Equilibrium with Strategic Complementarities. *Journal of Mathematical Economics*, 19, 305–321.