Universal knowledge–belief structures

Martin Meier

Instituto de Análisis Económico – CSIC, Barcelona, Spain

Received 24 May 2006

Available online 28 June 2007

Abstract

In this paper, we provide a notion of structure preserving maps (i.e. knowledge–belief morphisms) between knowledge–belief spaces. Then we show that—under the condition that the knowledge operators of the players in a knowledge–belief space operate only on measurable subsets of the space—there is a unique (up to isomorphism) universal knowledge–belief space to which every knowledge–belief space can be mapped by a unique knowledge–belief morphism.

© 2007 Elsevier Inc. All rights reserved.

JEL classification: C70; C72; D80; D82; D83

Keywords: Belief; Knowledge; Uncertainty; Universal type space; Universal knowledge space; Games of incomplete information

1. Introduction

Type spaces in the sense of Harsanyi (1967, 1968a, 1968b) and knowledge spaces are the most important tools to understand games of incomplete information. Player's uncertainty in a type space is represented by a σ -additive probability measure over the space. The knowledge spaces, respectively Harsanyi type spaces, used in applications are usually finite. Regardless, they typically do not contain enough states so as to represent all the potential states of mind that the players could possibly have about the interaction at hand. This poses the question as to whether the missing states prevent a correct analysis of the problem.

Mertens and Zamir (1985) showed that this problem does not arise for Harsanyi type spaces. They showed that, under suitable topological assumptions on the type spaces there is a “largest” Harsanyi type space which “contains” all possible states of the world. That is, they showed

E-mail address: martin.meier@iae.csic.es.

the existence of a *universal Harsanyi type space* to which every Harsanyi type space can be mapped by a *morphism*. A morphism is a map that preserves the state of nature and the beliefs of the players. Therefore, the analysis carried out in a finite or otherwise restrictive Harsanyi type space could be transferred, intact, to the universal Harsanyi type space. The proof of the existence of a universal Harsanyi type space was extended to more general topological cases by Brandenburger and Dekel (1993), Heifetz (1993), Mertens et al. (1994), and finally to the general measure theoretic case by Heifetz and Samet (1998b). Battigalli and Siniscalchi (1999) extended the framework of Brandenburger and Dekel to type spaces with beliefs conditioned on observations about nature and proved the existence of a universal type space for this case. Meier (2006) proved the existence of a universal Harsanyi type space if beliefs are described by finitely additive probability measures.

Knowledge (see for example Aumann, 1976, 1999a or Heifetz, 1997) is usually described by knowledge spaces. For each player, there is an information partition of the underlying space of states of the world: For each state, the player knows an event if and only if the partition member that contains that state is a subset of the event. Then, given an event (that is a subset of the underlying space of states of the world) one can consider the set of states in which the player knows the event. This set is itself an event and so one has derived, for each player, a knowledge operator that maps events to events. Such a knowledge operator satisfies certain properties: For example, a player knows what he knows, he knows what he does not know, and what he knows obtains. Alternatively, one could start directly with one such operator for each player. This would basically be just another way to write down the structure: One can show that each such knowledge operator is derived from a partition, where the partition can be defined in terms of the operator.

Heifetz and Samet (1998a) have shown that, unfortunately, unlike for Harsanyi type spaces, there is no universal knowledge space. This result was extended by Meier (2005) to the more general context of information structures (or “Kripke structures,” as they are called in Modal Logic), that can be viewed as generalizations of knowledge spaces that describe non-probabilistic beliefs.

A result in a similar vein is the impossibility result of Brandenburger and Keisler (2006) in the related context of possibility models. They associate each such structure with a first-order logic. The idea behind this approach is that the players would use this language to reason about each other and nature. What is desired then, is that for each possible belief of player i (which is a subset of the product of the space of states of nature and the space of types of the other player j that is definable by a first order formula) there is a type of player i who has exactly this belief. Unfortunately, to ask for beliefs-completeness in this sense is too much. Such a structure cannot exist (given at least 2 states of nature). The reason behind this is that first order logic can express self-referential statements which could be thought of as the formalization of a more complicated version of the antinomy of the liar, but involving two players: “Ann believes that Bob assumes that Ann believes that Bob’s assumption is wrong.” Whether Ann believes that Bob’s assumption is wrong or right, both cases lead to a contradiction.

In this framework of possibility models, Mariotti et al. (2005) were able to construct a universal structure by imposing topological restrictions on the possibility correspondences of the players. Their universal possibility model is also complete with respect to the beliefs defined in their setting, where a possible belief of a player is compact subset of the product of the space of states of nature and the space of types of the other player(s). In the universal model, for every such compact subset there exists a type of the player whose possibility set is exactly this compact subset. Note that by the impossibility result of Brandenburger and Keisler not every subset which is definable by a formula of first-order logic is such a compact subset. Indeed, complements of

such sets are also first-order definable, while the complement of a compact set might not be compact.

We remark that possibility models do not have all the nice properties of knowledge. For example, the truth axiom is violated. A possibility model is a product structure. That is, to every type t_1 of player 1 that occurs in some state of the world and every type t_2 of player 2 that occurs in some other state of the world there is a state of the world where player 1 is of type t_1 and at the same time the type of player 2 is t_2 . Such a property is of course incompatible with the truth axiom: In a knowledge space, if type t_1 of player 1 knows something that is incompatible with type t_2 of player 2, then there cannot exist a state of the world in which player 1 is of type t_1 and player 2 is of type t_2 .

If we do not assume that there is a common prior of which the beliefs of the players are posteriors, the beliefs of the players in a Harsanyi type space are subjective, and apart from some consistency conditions (so that the beliefs of a type of a player can be represented by a σ -additive probability measure, and that the players “know” their own beliefs) there is no restriction imposed upon the players’ beliefs. However, such a probabilistic assessment of events is necessary if one wants to study the behavior of utility maximizing agents (players).

The subjectivity of probabilistic beliefs is in sharp contrast to the properties of knowledge, where if a player knows an event in a state this event must obtain in that state. Hence, we believe that it is desirable to bring these two approaches together within one structure where one could exploit the strengths of both approaches.

Aumann (1999b) has constructed a canonical knowledge–belief space, that is, a knowledge–belief space which contains all states that can be described as maximal satisfied sets of formulas in the language of some modal logic. The only difference of his definition of a knowledge–belief space to the definition here is that the knowledge part in Aumann (1999b) is described by partitions of the state space, one for each player, and that the knowledge operators are derived from those partitions.

Fagin and Halpern (1994) also have a semantics for knowledge and probabilistic belief that is similar to Aumann’s and the one described here. The non-probabilistic part of their model consists of a Kripke structure. The probabilistic part is modeled in the following way: For each state and each player there is a probability space, whose underlying space is a subset of the state space of the Kripke structure. In their models, it is not guaranteed that the set of states where a player believes that some event E occurs with probability (at least) p is a measurable event. They present some properties sufficient to guarantee the measurability of such events. One condition sufficient to guarantee measurability is that, in each state all the players have the same probabilistic beliefs. However, in Game Theory and Economics one clearly wants to allow that in a state the beliefs of the players differ from each other.

Still, an appropriate definition of structure preserving maps, that is, morphisms, has to be provided and some appropriate conditions have to be found such that there is some hope to construct a universal knowledge–belief space under these conditions.

In this paper, we provide a notion of structure preserving maps (that is, knowledge–belief morphisms) between knowledge–belief spaces.

Then we show that under the assumption that the knowledge operators of the players in a knowledge–belief space operate only on measurable subsets of the space, we are able to construct a universal knowledge–belief space to which every knowledge–belief space can be mapped by a unique knowledge–belief morphism. This universal knowledge–belief space is unique up to isomorphism. The proof of the existence of the universal knowledge–belief space applies techniques which are similar to the ones used by Heifetz and Samet (1998b).

Why should the knowledge operators of the players just operate on measurable sets and not on all subsets of the space? The justification for this is that we think of events as those sets of states that the players can describe, and only those can be the objects of their reasoning. In view of this interpretation a statement saying “player i knows that the actual state of the world is in E ,” where E is an entity of states he cannot represent in his mind, is meaningless. Of course, it might well be that in *some* knowledge–belief spaces all subsets of the space of states of the world can be described by the players (for example in the finite knowledge–belief spaces), but we do not want to assume this in general.

Why do we not follow the standard procedure and start (like Aumann, 1999b) with partitions for the players and derive the players’ knowledge operators from them? Here, we do not restrict ourselves to finitely or countably many players and to a separable space of states of nature (or countably many atomic formulas). We do not know of any natural, and not too restrictive condition on the partitions of the players that would guarantee that the derived knowledge operators send measurable sets to measurable sets (except asking for this property directly, but then, why bother to introduce partitions in the first place?). Note that, for example, measurability of the partition members is not enough: There might be uncountably many partition members P_i that are contained in a certain event E and then the set $K_i E := \bigcup \{P_i \mid P_i \subseteq E\}$ might not be measurable.

2. Preliminaries

Unless otherwise stated, θ denotes functions from the set of states of the world to the set of states of nature, μ and ν denote measures, φ, χ, ψ expressions, and ω sets of expressions.

If not stated otherwise, we keep the following

Convention 1. If (M, Σ) is a measurable space, then $\Delta(M, \Sigma)$ denotes the space of probability measures on (M, Σ) . We consider this space as a measurable space endowed with the σ -field Σ_Δ generated by all the sets $\{\mu \in \Delta(M, \Sigma) \mid \mu(E) \geq p\}$, where $E \in \Sigma$ and $p \in [0, 1]$.

Remark 1. Let (M', Σ') and (M, Σ) be measurable spaces and let $f: M' \rightarrow M$ be measurable.

- (1) If μ' is a probability measure on (M', Σ') , then $\mu'(f^{-1}(\cdot))$ (that is $\mu'(f^{-1}(E))$, for $E \in \Sigma$) is a probability measure on (M, Σ) .
- (2) If $\Delta_f: \Delta(M', \Sigma') \rightarrow \Delta(M, \Sigma)$ is defined by $\Delta_f(\mu') := \mu'(f^{-1}(\cdot))$, for $\mu' \in \Delta(M', \Sigma')$, then Δ_f is measurable, since we have $\Delta_f(\mu')(E) \geq p$ iff $\mu'(f^{-1}(E)) \geq p$, for $E \in \Sigma$.

Definition 1. Let (M, Σ) be a measurable space. A function $K: \Sigma \rightarrow \Sigma$ is called a *knowledge operator* iff it satisfies the following properties:

- (1) $K(E) \subseteq E$,
- (2) $E \subseteq F$ implies $K(E) \subseteq K(F)$,
- (3) $\neg K(E) \subseteq K(\neg K(E))$,
- (4) $\bigcap_{n \in \mathbb{N}} K(E_n) \subseteq K(\bigcap_{n \in \mathbb{N}} E_n)$.

As is well known, positive introspection follows from the other properties of the knowledge operator. For a proof, see Aumann (1999a, p. 270).

Remark 2. Let K be a knowledge operator on the measurable space (M, Σ) . Then we have

$$K(E) \subseteq K(K(E)), \quad \text{for all } E \in \Sigma.$$

For the rest of this section, we fix a non-empty set of players I , a non-empty set of states of nature S , and, unless otherwise stated, a σ -field Σ_S on S , such that for all $s, s' \in S$ with $s \neq s'$ there is an $E \in \Sigma_S$ such that $s \in E$ and $s' \notin E$. Note that, apart from being non-empty, we do not impose any restriction on the cardinality of I and S .

We define now knowledge–belief spaces, i.e. the objects which we will study in this paper. They are Harsanyi type spaces which are endowed with additional knowledge operators, one for each player:

Definition 2. A knowledge–belief space on S for player set I is a 5-tuple

$$\underline{M} := \langle M, \Sigma, (K_i)_{i \in I}, (T_i)_{i \in I}, \theta \rangle,$$

where

- (1) M is a non-empty set,
- (2) Σ is a σ -field on M ,
- (3) K_i is a knowledge operator on (M, Σ) , for $i \in I$,
- (4) for $i \in I$: T_i is a $\Sigma - \Sigma_\Delta$ -measurable function from M to $\Delta(M, \Sigma)$, the space of probability measures on (M, Σ) ,
- (5) for $m \in M$ and $E \in \Sigma$: $m \in K_i(E)$ implies $T_i(m)(E) = 1$,
- (6) for $m \in M$ and $E \in \Sigma$: $[T_i(m)] \subseteq E$ implies $m \in K_i(E)$, where $[T_i(m)] := \{m' \in M \mid T_i(m') = T_i(m)\}$,
- (7) θ is a $\Sigma - \Sigma_S$ -measurable function from M to S .

The definition of the knowledge–belief spaces given here exhibits an asymmetry in that the knowledge part is modeled with knowledge operators, while the belief part is described by type functions and not by p -belief operators for $p \in [0, 1]$.

We explained already the reason for choosing knowledge operators instead of information partitions. We could certainly describe the beliefs part with p -belief operators. But to point out that the knowledge–belief spaces are Harsanyi type spaces with some additional knowledge structure, the beliefs part is still represented by type functions. Nevertheless, in Definition 6, p -belief operators will be defined.

We define now the beliefs preserving maps between knowledge–belief spaces.

Definition 3. Let $\underline{M}' = \langle M', \Sigma', (K'_i)_{i \in I}, (T'_i)_{i \in I}, \theta' \rangle$ and $\underline{M} = \langle M, \Sigma, (K_i)_{i \in I}, (T_i)_{i \in I}, \theta \rangle$ be knowledge–belief spaces on S for player set I .

A function $f : M' \rightarrow M$ is a *knowledge–belief morphism* if it satisfies the following conditions:

- (1) f is $\Sigma' - \Sigma$ -measurable,
- (2) for all $m' \in M'$

$$\theta'(m') = \theta(f(m')),$$

- (3) for all $E \in \Sigma$, and $i \in I$

$$K'_i(f^{-1}(E)) = f^{-1}(K_i(E)),$$

(4) for all $m' \in M'$, $E \in \Sigma$, and $i \in I$

$$T_i(f(m'))(E) = T'_i(m')(f^{-1}(E)).$$

Definition 4. A knowledge–belief morphism is a *knowledge–belief isomorphism*, if it is one-to-one, onto, and the inverse function is also a knowledge–belief morphism.

An easy check shows:

Remark 3. Knowledge–belief spaces on S for player set I , as objects, and knowledge–belief morphisms, as morphisms, form a category.

Definition 5. A knowledge–belief space $\underline{\Omega}$ on S for player set I is *universal*¹ if for every knowledge–belief space $\underline{M}$ on S for player set I there is a *unique* knowledge–belief morphism from $\underline{M}$ to $\underline{\Omega}$.

Remark 4. Universal knowledge–belief spaces on S for player set I are unique up to knowledge–belief isomorphism.

3. The universal knowledge–belief space in terms of expressions

Again, for this section, unless otherwise stated, we fix a non-empty player set I , and a measurable space of states of nature (S, Σ_S) such that for all $s, s' \in S$ with $s \neq s'$ there is an $E \in \Sigma_S$ such that $s \in E$ and $s' \notin E$.

Given these data, we define *kb-expressions* (allowing also for infinite conjunctions) which are natural generalizations of the *expressions* defined by Heifetz and Samet (1998b). Expressions are defined in a similar fashion as, for example, the formulas of the probability logic of Heifetz and Mongin (2001). Analogous to Heifetz and Samet (1998b), given a knowledge–belief space on S for player set I and a state of the world in this knowledge–belief space, we define the *kb-description* of this state as the set of those kb-expressions that are true in this state of the world. Then, we show that the set of all kb-descriptions constitutes a knowledge–belief space (Proposition 4) and that this knowledge–belief space is the universal knowledge–belief space (Theorem 1).

We define now the already mentioned *p*-belief operators.

Definition 6. For a knowledge–belief space $\langle M, \Sigma, (K_i)_{i \in I}, (T_i)_{i \in I}, \theta \rangle$ on S for player set I , $i \in I$, $E \in \Sigma$, and $p \in [0, 1]$ define

$$B_i^p(E) := \{m \in M \mid T_i(m)(E) \geq p\}.$$

Note that $B_i^p(E) = T_i^{-1}(\{\mu \in \Delta(M, \Sigma) \mid \mu(E) \geq p\})$ and that $B_i^p(E) \in \Sigma$, if $E \in \Sigma$.

Now, a language is introduced, the formulas of which are called kb-expressions. These kb-expressions are formulas of an infinitary modal language, where the atomic formulas are measurable sets of states of nature instead of primitive propositions.

¹ We use here the term “universal knowledge–belief space” although, in terms of category theory the term “terminal knowledge–belief space” would be the adequate one, since the universal knowledge–belief space is a terminal object in the category of knowledge–belief spaces. However, we use the former terminology for consistency with the existing literature on interactive epistemology.

Definition 7. Given a measurable space of states of nature (S, Σ_S) and a non-empty player set I , the set Φ of *kb-expressions* is the least set such that:

- (1) every $E \in \Sigma_S$ is an expression,
- (2) if φ is an expression, then $\neg\varphi$ is an expression,
- (3) if φ is an expression, then $k_i(\varphi)$ is an expression, for $i \in I$,
- (4) if φ is an expression, then $b_i^p(\varphi)$ is an expression, for $i \in I$ and $p \in [0, 1]$,
- (5) if Ψ is a non-empty set of expressions with $|\Psi| \leq \aleph_0$, then $\bigwedge_{\varphi \in \Psi} \varphi$ is an expression.

If Ψ is a non-empty set of expressions with $|\Psi| \leq \aleph_0$, then we set $\bigvee_{\varphi \in \Psi} \varphi := \neg \bigwedge_{\varphi \in \Psi} \neg\varphi$.

Definition 8. Let $\underline{M} := \langle M, \Sigma, (K_i)_{i \in I}, (T_i)_{i \in I}, \theta \rangle$ be a knowledge–belief space on S for player set I . Define

- (1) $E^{\underline{M}} := \theta^{-1}(E)$, for $E \in \Sigma_S$,
- (2) $(\neg\varphi)^{\underline{M}} := M \setminus \varphi^{\underline{M}}$, for $\varphi \in \Phi$,
- (3) $(k_i(\varphi))^{\underline{M}} := K_i(\varphi^{\underline{M}})$, for $\varphi \in \Phi$ and $i \in I$,
- (4) $(b_i^p(\varphi))^{\underline{M}} := B_i^p(\varphi^{\underline{M}})$, for $\varphi \in \Phi$, $i \in I$ and $p \in [0, 1]$,
- (5) $(\bigwedge_{\varphi \in \Psi} \varphi)^{\underline{M}} := \bigcap_{\varphi \in \Psi} \varphi^{\underline{M}}$, for an at most countable Ψ such that $\emptyset \neq \Psi \subseteq \Phi$.

So, defined as above, kb-expressions define measurable subsets of M . It is easy to check that $(\bigvee_{\varphi \in \Psi} \varphi)^{\underline{M}} = \bigcup_{\varphi \in \Psi} \varphi^{\underline{M}}$, for Ψ such that $\emptyset \neq \Psi \subseteq \Phi$ and $|\Psi| \leq \aleph_0$.

If no confusion may arise, we sometimes omit—with some abuse of notation—the superscript $\underline{M}$.

Definition 9. For a knowledge–belief space $\underline{M} := \langle M, \Sigma, (K_i)_{i \in I}, (T_i)_{i \in I}, \theta \rangle$ on S for player set I and $m \in M$ define $D(m)$, the *kb-description* of m , as

$$D(m) := \{\varphi \in \Phi \mid m \in \varphi^{\underline{M}}\}.$$

The following proposition shows that kb-morphisms preserve the description of a state.

Proposition 1. Let $\langle M, \Sigma, (K_i)_{i \in I}, (T_i)_{i \in I}, \theta \rangle$ and $\langle N, \Sigma^N, (K_i^N)_{i \in I}, (T_i^N)_{i \in I}, \theta^N \rangle$ be knowledge–belief spaces on S for player set I and let $f: M \rightarrow N$ be a kb-morphism. Then, for all $m \in M$

$$D(f(m)) = D(m).$$

We start now with the construction of the universal knowledge–belief space.

The first step is to define the set of states of the world and to endow this set with a σ -field.

Definition 10. Define Ω to be the set of all kb-descriptions of states of the world in knowledge–belief spaces on S for player set I . For $\varphi \in \Phi$ define

$$[\varphi] := \{\omega \in \Omega \mid \varphi \in \omega\}.$$

Obviously, we have $\Omega \setminus [\varphi] = [\neg\varphi]$ and $\bigcap_{\psi \in \Psi} [\psi] = [\bigwedge_{\psi \in \Psi} \psi]$, where φ is a kb-expression and Ψ is a non-empty set of kb-expressions with $|\Psi| \leq \aleph_0$. It follows that:

Remark 5. The set

$$\Sigma_{\Omega} := \{[\varphi] \mid \varphi \in \Phi\}$$

is a σ -field on Ω .

Lemma 1. For every knowledge–belief space $\underline{M}$ on S for player set I and for every $\varphi \in \Phi$, the kb-description map $D: M \rightarrow \Omega$ satisfies

$$D^{-1}([\varphi]) = \varphi^{\underline{M}}.$$

Note that Lemma 1 implies that D is measurable.

Now, for every $i \in I$, we define an operator K_i^* on $(\Omega, \Sigma_{\Omega})$.

Definition 11. For $\varphi \in \Phi$ define

$$K_i^*([\varphi]) := [k_i(\varphi)].$$

Proposition 2. For every knowledge–belief space $\underline{M} := \langle M, \Sigma, (K_i)_{i \in I}, (T_i)_{i \in I}, \theta \rangle$ on S for player set I , for every $\varphi \in \Phi$ and every $i \in I$, the kb-description map $D: M \rightarrow \Omega$ satisfies

$$K_i(D^{-1}([\varphi])) = D^{-1}(K_i^*([\varphi])).$$

Next, we endow our space with type functions, one for each player.

Proposition 3. For every $i \in I$ there exists a function

$$T_i^*: \Omega \rightarrow \Delta(\Omega, \Sigma_{\Omega})$$

such that for every knowledge–belief space $\underline{M}$ on S for player set I with kb-description map D and every $m \in M$

$$T_i^*(D(m))(\cdot) = T_i(m)(D^{-1}(\cdot)).$$

We also need a function θ^* which tells us which state of nature corresponds to a given state of the world.

Lemma 2. There is a measurable function $\theta^*: \Omega \rightarrow S$ such that for every knowledge–belief space $\underline{M}$ on S for player set I and every $m \in M$

$$\theta^*(D(m)) = \theta(m).$$

Note that there might be many states of the world to which a given state of nature corresponds to.

Now, we check that the K_i^* are indeed knowledge operators.

Lemma 3. K_i^* is a knowledge operator on $(\Omega, \Sigma_{\Omega})$, for every $i \in I$.

A first success is that we have constructed a knowledge–belief space, which is our candidate for being the universal knowledge–belief space.

Proposition 4. *The space*

$$\langle \Omega, \Sigma_\Omega, (K_i^*)_{i \in I}, (T_i^*)_{i \in I}, \theta^* \rangle$$

is a knowledge–belief space on S for player set I .

The following lemma guarantees, together with Proposition 1, that the kb-morphism from a knowledge–belief space to the universal knowledge–belief space is unique. Not surprisingly, this morphism is the corresponding kb-description map.

Lemma 4. *The kb-description map*

$$D: \Omega \rightarrow \Omega$$

is the identity.

Finally we obtain our desired main result.

Theorem 1. *The space*

$$\langle \Omega, \Sigma_\Omega, (K_i^*)_{i \in I}, (T_i^*)_{i \in I}, \theta^* \rangle$$

is a universal knowledge–belief space on S for player set I .

Proof. According to Lemma 1, for every knowledge–belief space $\underline{M}$ on S for player set I , the kb-description map $D: M \rightarrow \Omega$ is measurable, and according to Proposition 2, Proposition 3 and Lemma 2, D is a kb-morphism. It remains to show that it is the unique knowledge–belief morphism from $\underline{M}$ to $\underline{\Omega}$. But this is clear by Proposition 1 and Lemma 4. $\square$

4. Discussion

Comparing the existence result here with the non-existence results of Heifetz and Samet (1998a), Meier (2005), or the impossibility result of Brandenburger and Keisler (2006), one might ask, what is the difference, why does it work this time?

As already mentioned, the reason behind the impossibility result of Brandenburger and Keisler is that the language they use is so powerful that, in some sense, it allows to talk about formulas of the language itself. Clearly, the language used here to construct the universal knowledge–belief space is much weaker. Hence such self-referential statements cannot be made with the kb-expressions used in this paper.

The impossibility results of Heifetz and Samet (1998a) or Meier (2005) do not mention a language explicitly. The assumption that the knowledge operators operate on all subsets of the knowledge spaces leads to the following circularity: Informally, a universal space contains representations of all possible profiles of knowledge/belief-types of the players, especially those that could be defined within that space. Hence a minimal requirement in order to have a chance to construct a universal space is that one is able to find a space that has at least as many states as there are possible profiles of types of the players. However, if all subsets of the space are measurable, there are always more possible beliefs than states, so one would have to pass to a larger space, but in this larger space again more possible beliefs would be definable than the cardinality of states this space would contain, and hence this enlargement procedure would never come to

an end. The restriction of the knowledge operators to measurable events in this paper makes it possible to solve this problem.

Heifetz and Samet (1998b, p. 339) mention that their existence proof does also apply whenever beliefs are described by monotonic and continuous set-functions and not only when beliefs are sigma-additive probability measures. However, in the case of knowledge the carrier functions that represent knowledge are not continuous with respect to increasing sequences of events (a player may not know any event in an increasing sequence of events, but may nevertheless know their union).

One important difference between their construction and ours is that they manage to restrict themselves to expressions of finite length. This is so, because their Lemma 4.5 implies that the beliefs of a player (in a state) about events in the σ -algebra generated by the expressions are already determined by the beliefs about the expressions of finite length (note that in general the events defined by expressions of finite length form only an algebra but not a σ -algebra). As the above mentioned discontinuity indicates, in the case of knowledge this would probably not work (the “knowledge version” of Lemma 4.5 of Heifetz and Samet (1998b) does not exist), but it is also not needed: Since we allow for expressions of countable length, if in some state a player does not know any event in an increasing sequence of events each of which is defined by an expression, but does nevertheless know their union, our language is rich enough to express this fact.

Our language is closed under negation, countable conjunction (and also countable disjunction) and each measurable event in the universal knowledge–belief space is defined by some expression. Hence the kb-description of a state in some knowledge–belief space contains all the information about this state that needs to be taken into account in the construction of the universal knowledge–belief space. This in turn implies that if Heifetz and Samet (1998b) would have chosen to allow also for expressions of countable length, there would have been no need to prove Lemma 4.5 in their paper.

5. Proofs

Proof of Remark 4. If $\underline{\Omega}$ and $\underline{U}$ are universal knowledge–belief spaces (on the same space of states of nature and for the same player set, of course), then there are knowledge–belief morphisms $f: \underline{U} \rightarrow \underline{\Omega}$ and $g: \underline{\Omega} \rightarrow \underline{U}$. It is easy to check, that the composite of two knowledge–belief morphisms is also a knowledge–belief morphism and that the identity is always a knowledge–belief morphism from a knowledge–belief space $\underline{\Omega}$ to itself. By the uniqueness, it follows that $g \circ f = \text{id}_{\underline{U}}$ and therefore f is one-to-one and g is onto, and $f \circ g = \text{id}_{\underline{\Omega}}$ and therefore g is one-to-one and f is onto. f and g are knowledge–belief morphisms and $f = g^{-1}$ and $g = f^{-1}$. $\square$

Proof of Proposition 1. We show by induction on the formation of the expressions that $m \in \varphi^M$ iff $f(m) \in \varphi^N$

- Let $E \in \Sigma_S$. We have $\theta^N(f(m)) = \theta(m)$, so $f(m) \in E^N$ iff $m \in E^M$.
- We have

$$f(m) \in (\neg\varphi)^N \quad \text{iff} \quad f(m) \notin \varphi^N \quad \text{iff} \quad m \notin \varphi^M \quad \text{iff} \quad m \in (\neg\varphi)^M.$$

- Let Ψ be a non-empty set of expressions with $|\Psi| \leq \aleph_0$. Then

$$f(m) \in \left(\bigwedge_{\varphi \in \Psi} \varphi \right)^N \quad \text{iff for all } \varphi \in \Psi: f(m) \in \varphi^N,$$

which is by induction hypothesis the case iff for all $\varphi \in \Psi: m \in \varphi^M$, which is the case iff $m \in (\bigwedge_{\varphi \in \Psi} \varphi)^M$.

- We have

$$\begin{aligned} f(m) \in (k_i(\varphi))^N &\quad \text{iff} \quad f(m) \in K_i^N(\varphi^N) \quad \text{iff} \quad m \in f^{-1}(K_i^N(\varphi^N)) \quad \text{iff} \\ m \in K_i(f^{-1}(\varphi^N)). \end{aligned}$$

By the induction hypothesis, we have $f^{-1}(\varphi^N) = \varphi^M$. Hence it follows that

$$f(m) \in (k_i(\varphi))^N \quad \text{iff} \quad m \in (k_i(\varphi))^M.$$

- We have

$$f(m) \in (b_i^p(\varphi))^N \quad \text{iff} \quad T_i^N(f(m))(\varphi^N) \geq p \quad \text{iff} \quad T_i(m)(f^{-1}(\varphi^N)) \geq p.$$

By the induction hypothesis: $f^{-1}(\varphi^N) = \varphi^M$. Hence $T_i(m)(f^{-1}(\varphi^N)) = T_i(m)(\varphi^M)$. We have $T_i(m)(\varphi^M) \geq p$ iff $m \in (b_i^p(\varphi))^M$. It follows that

$$f(m) \in (b_i^p(\varphi))^N \quad \text{iff} \quad m \in (b_i^p(\varphi))^M. \quad \square$$

Proof of Lemma 1. Clear by the definition of $[\varphi]$. $\square$

Proof of Proposition 2. By Lemma 1 we have

$$K_i(D^{-1}([\varphi])) = K_i(\varphi^M) = (k_i(\varphi))^M = D^{-1}([k_i(\varphi)]) = D^{-1}(K_i^*([\varphi])). \quad \square$$

Proof of Proposition 3. For $\omega \in \Omega$ choose a knowledge–belief space $\underline{M}$ on S for player set I and $m \in M$ such that $D(m) = \omega$. For $[\varphi] \in \Sigma_\Omega$ define

$$T_i^*(\omega)([\varphi]) := T_i(m)(D^{-1}([\varphi])).$$

We have

$$T_i(m)(D^{-1}([\varphi])) = T_i(m)(\varphi^M) = \sup\{p \mid b_i^p(\varphi) \in D(m)\},$$

so $T_i^*(\omega)([\varphi])$ depends just on $D(m)$ and is well defined. By Remark 1, we have

$$T_i(m)(D^{-1}(\cdot)) \in \Delta(\Omega, \Sigma_\Omega). \quad \square$$

Proof of Lemma 2. Let

$$d_0(m) := \{E \in \Sigma_S \mid m \in \theta^{-1}(E)\}.$$

Obviously, $d_0(m) = D(m) \cap \Sigma_S$. By the properties of (S, Σ_S) , we have for all $s \in S: \{s\} = \bigcap_{s \in E \in \Sigma_S} E$. It follows for every knowledge–belief space $\underline{M}'$ on S for player set I and $m' \in M'$ that

$$\theta(m') = s \quad \text{iff} \quad d_0(m') = \{E \mid s \in E\}.$$

For $\omega \in \Omega$ choose a knowledge-belief space $\underline{M}$ on S for player set I and $m \in M$, such that $D(m) = \omega$. Define now $\theta^*(\omega) := \theta(m)$. Since $\theta(m)$ just depends on $D(m)$, $\theta^*(\omega)$ is well defined.

It remains to show that θ^* is measurable: Let $E \in \Sigma_S$. We have

$$\theta^*(D(m)) \in E \quad \text{iff} \quad m \in \theta^{-1}(E) \quad \text{iff} \quad E \in D(m) \quad \text{iff} \quad D(m) \in [E].$$

It follows that $\theta^{*-1}(E) = [E]$. $\square$

Proof of Lemma 3. Let $(\chi_n)_{n \in \mathbb{N}}$, $\varphi, \psi \in \Phi$ and $\underline{M}$ be a knowledge-belief space and $m \in M$.

- (1) Let $D(m) \in K_i^*([\varphi])$. Then, we have $m \in K_i(\varphi^{\underline{M}})$ and therefore $m \in \varphi^{\underline{M}}$. It follows that $D(m) \in [\varphi]$.
- (2) Let $[\varphi] \subseteq [\psi]$ and $D(m) \in K_i^*([\varphi])$. We have $\varphi^{\underline{M}} = D^{-1}([\varphi]) \subseteq D^{-1}([\psi]) = \psi^{\underline{M}}$ and $m \in K_i(\varphi^{\underline{M}})$. It follows that $m \in K_i(\psi^{\underline{M}})$ and therefore $D(m) \in K_i^*([\psi])$.
- (3) Let $D(m) \in \Omega \setminus K_i^*([\varphi])$. We have $m \in M \setminus K_i(\varphi^{\underline{M}})$ and therefore $m \in K_i(M \setminus K_i(\varphi^{\underline{M}}))$. By the definitions, we have $D(m) \in [k_i(\neg k_i(\varphi))]$. It follows that $D(m) \in K_i^*(\Omega \setminus K_i^*([\varphi]))$.
- (4) Let $D(m) \in \bigcap_{n \in \mathbb{N}} K_i^*([\chi_n])$. It follows that $m \in \bigcap_{n \in \mathbb{N}} K_i(\chi_n^{\underline{M}})$ and therefore $m \in K_i(\bigcap_{n \in \mathbb{N}} \chi_n^{\underline{M}})$. By the definitions, we have $D(m) \in [k_i(\bigwedge_{n \in \mathbb{N}} \chi_n)]$. It follows that $D(m) \in K_i^*(\bigcap_{n \in \mathbb{N}} [\chi_n])$. $\square$

Proof of Proposition 4. It remains to show

- (1) Ω is non-empty.
- (2) For every $i \in I$: T_i^* is measurable as a function from Ω to $\Delta(\Omega, \Sigma_\Omega)$.
- (3) For every $i \in I$, $\omega \in \Omega$ and $\varphi \in \Phi$: $\omega \in K_i^*([\varphi])$ implies $T_i^*(\omega)([\varphi]) = 1$.
- (4) For every $i \in I$, $\omega \in \Omega$ and $\varphi \in \Phi$

$$\{\omega' \in \Omega \mid T_i^*(\omega') = T_i^*(\omega)\} \subseteq [\varphi],$$

implies $\omega \in K_i^*([\varphi])$.

To:

- (1) Let $M := \{m\}$ and choose $s \in S$. Set $\Sigma := \text{Pow}(M)$, $K_i(E) := E$, for $E \subseteq M$ and $i \in I$, $T(m) := \delta_m$, for $i \in I$, and $\theta(m) := s$. Then

$$\langle M, \Sigma, (K_i)_{i \in I}, (T_i)_{i \in I}, \theta \rangle$$

is a knowledge-belief space on S for player set I .

- (2) Since inverse images commute with unions, intersections and complements, it is enough to show that $T_i^{*-1}(\beta^p(E)) \in \Sigma_\Omega$, for

$$\beta^p(E) := \{\mu \in \Delta(\Omega, \Sigma_\Omega) \mid \mu(E) \geq p\},$$

where $E \in \Sigma_\Omega$ and $p \in [0, 1]$. We have

$$T_i^{*-1}(\{\mu \in \Delta(\Omega, \Sigma_\Omega) \mid \mu(E) \geq p\}) = \{\omega \in \Omega \mid T_i^*(\omega)(E) \geq p\}.$$

Since $E \in \Sigma_\Omega$, there is a kb-expression φ such that $E = [\varphi]$. Note that if $p \in [0, 1]$ and $p = \sup\{q \mid b_i^q(\varphi) \in \omega\}$, then $b_i^p(\varphi) \in \omega$. This implies that

$$\omega \in T_i^{*-1}(\{\mu \in \Delta(\Omega, \Sigma_\Omega) \mid \mu(E) \geq p\}) \quad \text{iff} \quad b_i^p(\varphi) \in \omega.$$

It follows that $T_i^{*-1}(\{\mu \in \Delta(\Omega, \Sigma_\Omega) \mid \mu(E) \geq p\}) = [b_i^p(\varphi)]$.

- (3) Chose a knowledge–belief space $\underline{M}$ on S for player set I and $m \in M$ such that $D(m) = \omega$. Since $\omega \in K_i^*([\varphi])$, we have by Proposition 2 and Lemma 1 $m \in K_i(\varphi^M)$ and therefore $1 = T_i(m)(\varphi^M) = T_i(m)(D^{-1}([\varphi])) = T_i^*(D(m))([\varphi])$.
- (4) Let φ be a kb-expression and

$$\{\omega' \in \Omega \mid T_i^*(\omega') = T_i^*(\omega)\} \subseteq [\varphi].$$

Chose a knowledge–belief space $\underline{M}$ on S for player set I and $m \in M$ such that $D(m) = \omega$. Let $m' \in M$. If $T_i^*(D(m')) \neq T_i^*(D(m))$, then there is a kb-expression ψ such that $T_i(m')(\psi^M) \neq T_i(m)(\psi^M)$. It follows that

$$D(\{m' \in M \mid T_i(m') = T_i(m)\}) \subseteq \{\omega' \in \Omega \mid T_i^*(\omega') = T_i^*(\omega)\},$$

which implies

$$\{m' \in M \mid T_i(m') = T_i(m)\} \subseteq D^{-1}([\varphi]) = \varphi^M.$$

So we have

$$m \in K_i(\varphi^M) = (k_i(\varphi))^M.$$

And hence, $\omega \in [k_i(\varphi)] = K_i^*([\varphi])$. $\square$

Proof of Lemma 4. For $\omega \in \Omega$, we have

$$\omega = \{\varphi \in \Phi \mid \omega \in [\varphi]\}.$$

We have to show that for every kb-expression φ and every $\omega \in \Omega$: $\omega \in \varphi^\Omega$ iff $\omega \in [\varphi]$. We know this already if $\varphi = E$, where $E \in \Sigma_S$. It is obvious that $\Omega \setminus [\varphi] = [\neg\varphi]$, and that if Ψ is a non-empty set of kb-expressions of cardinality $\leq \aleph_0$, then

$$\bigcap_{\varphi \in \Psi} [\varphi] = \left[\bigwedge_{\varphi \in \Psi} \varphi \right].$$

By the definition, we have $K_i^*([\varphi]) = [k_i(\varphi)]$. So it remains to show that $[\varphi] = \varphi^\Omega$ implies $[b_i^p(\varphi)] = B_i^p([\varphi])$. For $\omega \in \Omega$, choose a knowledge–belief space $\underline{M}$ on S for player set I and $m \in M$ such that $D(m) = \omega$. We have

$$D(m) \in [b_i^p(\varphi)] \quad \text{iff} \quad b_i^p(\varphi) \in D(m) \quad \text{iff} \quad T_i(m)(\varphi^M) \geq p.$$

But we have

$$T_i^*(\omega)([\varphi]) = T_i(m)(D^{-1}([\varphi])) = T_i(m)(\varphi^M).$$

This implies that $[b_i^p(\varphi)] = B_i^p([\varphi])$. $\square$

Acknowledgments

I thank Adam Brandenburger, Joe Halpern, several anonymous referees and an anonymous Associate Editor for very helpful comments. This work was supported by the European Union via a Marie-Curie Individual Fellowship (CORE), and later by the Spanish Ministerio de Educación y Ciencia via a Ramon y Cajal Fellowship (IAE-CSIC) and a Research Grant (SEJ 2004-07861), as well as the Fundación Ramón Areces and Barcelona Economics (XREA). Parts of the material appeared without proofs in my paper “Harsanyi Type Spaces with Knowledge Operators” prepared for the conference TARK 2005.

References

- Aumann, R.J., 1976. Agreeing to disagree. *Ann. Statist.* 4, 1236–1239.
- Aumann, R.J., 1999a. Interactive epistemology I: Knowledge. *Int. J. Game Theory* 28, 263–300.
- Aumann, R.J., 1999b. Interactive epistemology II: Probability. *Int. J. Game Theory* 28, 301–314.
- Battigalli, P., Siniscalchi, M., 1999. Hierarchies of conditional beliefs and interactive epistemology in dynamic games. *J. Econ. Theory* 88, 188–230.
- Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *J. Econ. Theory* 59, 189–198.
- Brandenburger, A., Keisler, H.J., 2006. An impossibility theorem on beliefs in games. *Studia Logica* 84, 211–240.
- Fagin, R., Halpern, J., 1994. Reasoning about knowledge and probability. *J. ACM* 43, 340–367.
- Harsanyi, J.C., 1967. Games with incomplete information played by Bayesian players, part I. *Manage. Sci.* 14, 159–182.
- Harsanyi, J.C., 1968a. Games with incomplete information played by Bayesian players, part II. *Manage. Sci.* 14, 320–334.
- Harsanyi, J.C., 1968b. Games with incomplete information played by Bayesian players, part III. *Manage. Sci.* 14, 486–502.
- Heifetz, A., 1993. The Bayesian formulation of incomplete information—The non-compact case. *Int. J. Game Theory* 21, 329–338.
- Heifetz, A., 1997. Infinitary $S5$ -epistemic logic. *Math. Log. Quart.* 43, 333–342.
- Heifetz, A., Mongin, P., 2001. Probability logic for type spaces. *Games Econ. Behav.* 35, 31–53.
- Heifetz, A., Samet, D., 1998a. Knowledge spaces with arbitrarily high rank. *Games Econ. Behav.* 22, 260–273.
- Heifetz, A., Samet, D., 1998b. Topology-free typology of beliefs. *J. Econ. Theory* 82, 324–341.
- Mariotti, T., Meier, M., Piccione, M., 2005. Hierarchies of beliefs for compact possibility models. *J. Math. Econ.* 41, 303–324.
- Meier, M., 2005. On the nonexistence of universal information structures. *J. Econ. Theory* 122, 132–139.
- Meier, M., 2006. Finitely additive beliefs and universal type spaces. *Ann. Probab.* 34, 386–422.
- Mertens, J.F., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. *Int. J. Game Theory* 14, 1–29.
- Mertens, J.F., Sorin, S., Zamir, S., 1994. Repeated games: Part A: Background material. CORE Discussion Paper Nr. 9420, Université Catholique de Louvain.