



ELSEVIER

Mathematical Social Sciences 38 (1999) 157–170

mathematical
social
sciences

Hierarchies of knowledge: an unbounded stairway[☆]

Aviad Heifetz^a, Dov Samet^{b,*}

^a*School of Economics, Tel Aviv University, Tel Aviv, Israel*

^b*Faculty of Management, Tel Aviv University, Tel Aviv, Israel*

Abstract

The rank of a partition space is the maximal ordinal number of steps in the process, in which events of the space are generated by successively applying knowledge operators, starting with events of nature. It is shown in Heifetz and Samet (1998) that this rank may be an arbitrarily large ordinal [Heifetz, A., Samet, D., 1998. Knowledge spaces with arbitrarily high rank. *Games and Economic Behavior* 22, 260–273]. Here we construct for each ordinal α a canonical partition space U_α , in analogy with the Mertens and Zamir (1985) hierarchical construction for probabilistic beliefs [Mertens, J.F., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. *Int. J. Game Theory* 14, 1–29]. Our main result is that each partition space of rank α is embeddable as a subspace of U_λ , where λ is the first limit ordinal exceeding α . © 1999 Elsevier Science B.V. All rights reserved.

1. Introduction

1.1. Models of knowledge and belief

Both partition spaces and Harsanyi (1967–68)-type spaces are predominant in modeling uncertainty in game theory and economics. The first kind models knowledge in set theoretic terms, the other models belief in probabilistic terms. In both kinds of models, each possible state of the world in the space is associated with a *state of nature* and a *type* for each player. The state of nature is a description of the exogenous parameters that do not depend on the players' uncertainties – the initial endowments, the payoff functions and the like. By type of players we mean their *epistemic state*, that is, an account of their certainties and uncertainties. In Harsanyi type spaces, henceforth

[☆]This is a thoroughly revised version of Heifetz and Samet (1993). We are grateful to Ehud Lehrer, Ron Fagin, Philippe Mongin and Luc Lismont for fruitful discussions and comments.

*Corresponding author. Tel.: +972-3-640-8024; fax: +972-3-640-9560.

E-mail address: samet@vm.tau.ac.il (D. Samet)

probabilistic type spaces, a player's type in a given state (i.e., epistemic state) is a σ -additive probability measure on the space, associated with the state. In partition spaces, a player's type in a state is given by the set of states the player considers possible in that state. This set of possible states is the element of the player's partition which contains that state. In this work we examine the ways in which the general structure of partition spaces differs from and resembles that of probabilistic type spaces.

1.2. *Unfolding epistemic states by α -order events*

The above definition of types enables us to unfold the mutual certainties and uncertainties of the players. In each state and for each player, we can describe explicitly the player's perception of the state of nature, the player's perception of how others perceive the state of nature, and so forth. This explicit description of the epistemic states of the players is reached as follows. First we define epistemic operators on the space. These operators map events to events, where an event is a subset (a measurable one in the probabilistic case) of the state space. In partition spaces each player has one epistemic operator – the knowledge operator. It maps an event E to the event $K(E)$ called 'the player knows E '. This event includes every state for which E is true in all the states considered possible at that state by the player. In probabilistic type spaces each player has, for each p between 0 and 1, a p -belief operator. The event 'the player p -believes E ' consists of all the states in which he ascribes probability at least p to E (see e.g., Monderer and Samet, 1989). Next we define by induction *events of order α* . 0-order events are the natural events – those that contain all the states whose state of nature belongs to some given (measurable, in the probabilistic case) set of states of nature. 1-order events are those describing nature (i.e., 0-order events) and the players' attitude (knowledge or belief, according to the case) towards 0-order events. In general, α -order events describe nature and the epistemic attitude towards lower order events. Thus, the α -order events (for all the ordinals α) to which a given state belongs describe the epistemic state of all the players (as well as the state of nature) in that state.

1.3. *How partition spaces and probabilistic type spaces differ in the construction of α -order events*

How far does this inductive construction of α -order events proceed? Here partition spaces differ from probabilistic type spaces. In the latter, the construction need not go beyond the finite order events. That is, if the event E is in the σ -field generated by all the finite order events, then any belief concerning E already belongs to that σ -field. This property of probabilistic type spaces follows from the main theorem of Mertens and Zamir (1985), which we discuss shortly.

In partition spaces, by contrast, finite order events may not be enough in order to describe the players' scope of knowledge. After constructing all these events we can, in some spaces, construct new events which describe knowledge regarding all the finite order events. And we must not stop at this point either. For any given ordinal α there are spaces in which the construction adds new events for all orders preceding α . In section 3 we describe, for every given ordinal and with only two players, an explicit construction of such a space, taken from Heifetz and Samet (1998).

1.4. *The construction of epistemic states by α -order types – the probabilistic case*

The construction of α -order events helps to unravel in an explicit way the epistemic state of the players in a given state space. There is also another, canonical, way to construct explicitly all the types of the players, starting from a given set of states of nature. This construction was carried out by Armbruster and Böge (1979), Böge and Eisele (1979), and Mertens and Zamir (1985) for the probabilistic case. The construction progresses by inductively defining α -order types. 0-order types are simply the states of nature.¹ A player's 1-order type is a σ -additive probability measure on the 0-order types (i.e., on the set of states of nature). A player's α -order type is a σ -additive probability measure on all the combinations of previously defined types. We define then a state space, called the universal state space, each element of which is a sequence starting with a 0-order type (a state of nature) and specifying the α -order type of each player for all the α 's (taking into account certain consistency restrictions on such sequences). As it turns out, there is no need to go beyond finite-order types. The sequence of all finite-order types of a player (specified in a state) defines a unique σ -additive probability measure on all the sequences of finite order types, that is, it determines its unique ω -order type. Thus the universal state space is a probabilistic type space, where a player's type in a state is his ω -order type, which is determined by the state. Mertens and Zamir (1985) showed that this space is universal not only by virtue of including all possible sequences of finite-order types, but also in light of its relation to the class of probabilistic type spaces: they proved that any non-redundant type space (see Remark 2.4 below) can be embedded in a natural way in the universal type space.

1.5. *The construction of epistemic states by α -order types – the knowledge case*

An analogous construction of α -order types can be carried out for knowledge types. 0-order types are, as before, states of nature. A player's 1-order type is a subset of 0-order types. A player's α -order type is a subset of combinations of all previously defined types (with appropriate consistency restrictions). Again, we can define a state as sequence starting with a 0-order type and specifying the α -order types of the players for each α . But here, as one can expect in light of the construction in section 3, the sequence may develop in a non-trivial way beyond any given ordinal. This result, which we prove in section 4, follows from the construction of section 3, and the theorem we describe in the next paragraph. Consequently, we have a host of canonical spaces – one for each ordinal α . The canonical α -order type space consists of all the sequences of types of orders smaller than α , and is indeed a partition space.

As in the theorem of Mertens and Zamir (1985), we can still relate the class of partition spaces to the hierarchic type spaces. We define the rank of a partition space to be the least ordinal α for which higher-order events appear already as α -order events. In section 4 we prove that a (non-redundant) partition space of rank α can be isomorphical-

¹In the above cited works it is assumed that the set of states of nature is Hausdorff compact. Brandenburger and Dekel (1993) showed that the same results can also be obtained when this set is Polish. Heifetz (1993) relaxed both assumptions and showed it can be done with any Hausdorff space. Various other sufficient topological conditions are discussed in Mertens et al. (1994). In Heifetz and Samet (1996) we showed that these results may fail in the general measure-theoretic case.

ly embedded in the λ -canonical type space, where λ is the least limit ordinal exceeding α . Moreover, the image of the partition space in this λ -canonical type space has the following property, which resembles that of the unique universal space of Mertens and Zamir (1985): each sequence of types of a player in a state of the space can be continued in a unique way beyond λ .

Fagin et al. (1991) were the first to give an example of three players with mutual uncertainties of rank $\omega + 1$. Fagin et al. (1992) presented the inductive construction and found conditions under which a ω hierarchy suffices to characterize the knowledge of the player. In a work parallel to ours, Fagin (1994) proved analogous results, but with a logical, rather than set-theoretical, construction.

The class of partition spaces is sound and complete with respect to the multi-player epistemic logic S5. The transfinite aspects of mutual knowledge explored here arguably call for a logical treatment where infinite conjunctions and disjunctions are allowed in the language. Such a logical language is presented in Fagin (1994), and an axiomatic approach with completeness results is further investigated in Heifetz (1997).

2. Partition spaces and their rank

A *partition space* consists of a set Ω called the space of states of the world. Each state in Ω stands for a situation in which the game or trade may take place. We denote by S the set of *states of nature*. Each element of S is the specification of all the values of the parameters that are relevant to the interaction between the players and that do not depend on their uncertainties. In each state of the world a certain state of nature prevails. The function $\Theta: \Omega \rightarrow S$ maps each state of the world to its state of nature.

We denote by I the set of players. Each player $i \in I$ has a partition Π^i on Ω . The member of i 's partition that contains ω is denoted by $\Pi^i(\omega)$. In state $\omega \in \Omega$ player i considers as possible all the states in $\Pi^i(\omega)$, and he can not tell which of them obtains. To sum up, a partition space is specified by the tuple

$$\langle \Omega, S, \Theta: \Omega \rightarrow S, I, (\Pi^i)_{i \in I} \rangle.$$

Events are subsets of Ω . For a given event E , the event $K^i(E)$ – “player i knows E ” – is: $K^i(E) = \{\omega \in \Omega: \Pi^i(\omega) \subseteq E\}$. That is, i knows E in state ω , if E is true in all the states he considers possible at ω . The operators $K^i: 2^\Omega \rightarrow 2^\Omega$ are called the *knowledge operators*.

Define also $K(E) = \bigcap_{i \in I} K^i(E)$. The event that E is *common knowledge* is $C(E) = \bigcap_{n=1}^\infty (K)^n(E)$ where $(K)^n$ is the operator K applied successively n times.

Using these operators we can reveal the structure of knowledge in the partition space – what players know about the state of nature, what they know about what they know about the state of nature and so on. This is done by constructing for every ordinal α a partition P_α of Ω , by transfinite induction. An arbitrary union of elements of P_α is called an α -order event. The first partition is defined by:

$$P_0(\omega) = P_0(\omega') \quad \text{iff} \quad \Theta(\omega) = \Theta(\omega')$$

For the ordinal α :

$$P_\alpha(\omega) = P_\alpha(\omega') \quad \text{iff}$$

$$\Theta(\omega) = \Theta(\omega'), \text{ and } (\omega \in K^i(E) \Leftrightarrow \omega' \in K^i(E)),$$

for every β -order event E , $\beta < \alpha$ and $i \in I$.

The partition P_0 generates the natural events. The partition P_α generates the mutual-knowledge events, where the chain of mutual references of the players to each other's knowledge is of length at most α . Clearly for $\alpha < \beta$, P_β weakly refines P_α . But the cardinality of a sequence of strictly refining partitions cannot exceed the cardinality of Ω . Therefore, we must have for some α , $P_\alpha = P_{\alpha+1}$, or equivalently, $P_\alpha = P_\beta$ for all $\beta > \alpha$. The sequence of partitions up to P_α gives the full account of the knowledge structure in the space. This leads to the following definition:

Definition 2.1. The *rank* of the partition space Ω , is the minimal ordinal α such that $P_\alpha = P_{\alpha+1}$.

If two states are not separated by P_α , where α is the rank of the space, then they stand for the same natural parameters and uncertainties. Hence the obvious definition of redundancy in partition spaces:

Definition 2.2. The partition space Ω is non-redundant if for α , the rank of the space, P_α is the partition of Ω into single points.

Proposition 2.3. *If the partition space Ω is non-redundant then any two states are either separated by their state of nature or by the partition of some player.*

Proof. Otherwise, there are two states $\omega, \omega' \in \Omega$ with the same state of nature and such that $\Pi^i(\omega) = \Pi^i(\omega') \quad \forall i \in I$. Thus, ω and ω' are not separated by P_0 . Furthermore, if they are not separated by $P_\beta \quad \forall \beta < \alpha$, then they are neither separated by $K^i(E)$, where E is β -order event for some $\beta < \alpha$ and $i \in I$. Hence, ω and ω' are not separated by P_α . QED

Examples.

1. Let $\Omega = \{\omega_1, \omega_2\}$, and suppose that there is only one state of nature. Thus, $P_0 = \Omega$. Suppose that there is only one player, and his partition is the trivial one – $\{\Omega\}$. Then, as $P_0 = P_1$, the rank of Ω is 0. The space is redundant, since ω_1 and ω_2 cannot be separated.

In general, if we replace a given state ω in a partition space by two ‘copies’ ω' and ω'' , leaving the two copies in the same elements of the partitions in which ω was, and associating with them the same state of nature that was associated with ω , then the resulting space is redundant.

2. Consider again the space $\Omega = \{\omega_1, \omega_2\}$, where $P_0 = \Omega$. Assume now that there is a single player with the partition $\{\{\omega_1\}, \{\omega_2\}\}$. The rank of this space is also 0, and the space is redundant.

In general, the union of two isomorphic copies of the same partition space, where the partition on this union is the union of the partitions on each of the copies, is a redundant partition space.

Note that Proposition 2.3 is a necessary condition for non-redundancy, but it is not sufficient, as demonstrated by example (b).

Remark 2.4. Mertens and Zamir [(1985), Def. 2.4] used a condition analogous to the necessary condition in Proposition 2.3 to define non-redundancy. With this definition the union of two isomorphic copies of a non-redundant probabilistic type space is non-redundant, and hence Proposition 2.5 there does not hold. To remedy this, one should define inductively $\mathcal{F}_n = \sigma\{\mathcal{F}_0, \{t^i(\omega)(E) \geq r : r \in [0,1], E \in \mathcal{F}_k, i \in I, k < n\}\}$ (where \mathcal{F}_0 is the σ -field of natural events and $t^i(\omega)$ is the type of player i in ω), and then define a type space to be non-redundant whenever any two of its points are separated by some \mathcal{F}_n .

3. A construction of partition spaces with arbitrary rank

It is possible to define the rank of probabilistic type spaces analogously to the way rank is defined here for partition spaces. However, it follows from the work of Mertens and Zamir (1985) that the rank of a probabilistic type space never exceeds ω . In contrast, for any given set of states of nature S , with at least two elements, there is a partition space on S of rank α , for every ordinal α . We cite here an example from Heifetz and Samet (1998), in which for any α , a non-redundant partition space W_α of rank α is constructed for two players, 1 and 2, and two states of nature. We review the main features of the example, without the proofs.

The states of nature are the two results of tossing a coin, H and T . The space W_α consists of triples $w = (w^0, w^1, w^2)$, where $w^0 \in \{H, T\}$, and w^1 and w^2 are the *consciousness records* (*records* for short) of players 1 and 2, respectively, which we describe in a short while. The first coordinate w^0 is the state of nature at the point w (i.e. $\Theta(w) = w^0$). A record of a player is a sequence of S 's and D 's of length α . That is, $w^1, w^2 \in \{S, D\}^\alpha$ where the ordinal α is, as usual, the set of ordinals smaller than α .

Not all such sequences are qualified as records. The following story motivates the special structure of records. Think of S as standing for Sober and D for Drunk. Our players are hopelessly addicted to alcohol: Whenever they try to avoid it, they finally fail and give up. Formally, for every record and every limit ordinal $\lambda \leq \alpha$, there exists an ordinal $\gamma < \lambda$ such that there are only D 's in the sequence from γ up to (not including)

λ . The record is called λ -even if this γ is even² and λ -odd otherwise. Thus, we have completed the description of the space W_α and the map Θ . Next we describe the partitions Π^i .

If the record of player i starts with Sober (i.e., $w_0^i = S$), that means that player i knows whether the state of nature, w_0 , is H or T . If the player is Drunk in the beginning, he cannot tell what the state of nature is. Thus, if $w_0^i = S$, then $\hat{w} \in \Pi^i(w)$ only if

$$(i) \hat{w}^0 = w^0.$$

If $w_0^i = D$, no restriction is imposed on \hat{w}^0 .

Being Sober at a given non-limit ordinal β means, so goes the story, that the player can tell whether his opponent was Sober or Drunk at the ordinal preceding β . Thus for example, when player i is Sober in the next level ($w_1^i = S$), he can tell whether the other player's record starts with S or with D . This means that he knows whether his opponent knows the state of nature. It makes sense, of course, even if he himself can not tell the true state of nature – in case he is Drunk in the beginning. To summarize, if $w_\beta^i = S$, then $\hat{w} \in \Pi^i(w)$ only if

$$(ii) \hat{w}_{\beta-1}^i = w_{\beta-1}^i.$$

If $w_\beta^i = D$, no restriction is imposed on $\hat{w}_{\beta-1}^i$.

The above interpretation of being sober works only for non-limit ordinals, since they have an immediate predecessor. Being Sober at a limit level λ means that the player knows whether the other player is λ -even or λ -odd. That is, if $w_\lambda^i = S$, then $\hat{w} \in \Pi^i(w)$ only if

$$(iii) \hat{w}_{\beta-1}^i \text{ is of the same parity as } w_{\beta-1}^i.$$

If $w_\lambda^i = D$, no restriction is imposed on the λ parity \hat{w}^i .

The crucial point is that no combination of S 's and D 's up to (not including) a limit ordinal λ ever enables a player to perceive whether the other player is λ -even or λ -odd. This is because in such a combination he himself is always Drunk from some γ on, so he cannot exclude the possibility that the other player stood the temptation longer than he did, and fell Drunk (up to λ) only at some later stage $\gamma' > \gamma$, where γ' may be even as well as odd. Therefore, becoming Sober again in stage λ enables the player to exclude some records of the other player that he can not exclude when he is Drunk there. The state of consciousness of the players at stage λ is, therefore, a meaningful source of uncertainty for the next stage, and so on.

In addition we assume that every player always knows his own record. That is, $\hat{w} \in \Pi^i(w)$ only if

$$(iv) \hat{w}^i = w^i.$$

These restrictions define the partitions Π^i . That is, $\hat{w} \in \Pi^i(w)$, if and only if \hat{w} satisfies conditions (i)–(iv). It is shown in Heifetz and Samet (1998) that the rank of W_α is α .

²i.e., in the unique representation $\gamma = \lambda' + n$, where λ' is a limit ordinal (or 0) and $n \in \mathbb{N}$, n is even.

Remark 3.3. For every ordinal α , the cardinality of W_α is the same as that of α . In particular, W_α is countable whenever α is. Fagin (1994) presented the construction of such partition spaces as an open problem, which was solved by this example.

4. Canonical partition spaces

In section 2 we derived the mutual knowledge of the players, in a state of a partition space, by constructing α -order events. In this section we describe the players' knowledge explicitly, in analogy with the construction of the universal space for probabilistic beliefs of Mertens and Zamir (1985). We construct for each ordinal α a canonical space U_α of rank α . Each point of U_α is a hierarchy of length α – it describes the state of nature and the types of the players of order smaller than α – i.e., what players know about lower order types. In the case of the universal space for probabilistic beliefs, only finite order types are required. Once these types are specified, further, infinite-order types are determined uniquely. The partition spaces constructed in section 3 suggest that in the case of knowledge, for any ordinal α , α -order types may not be determined by lower order types. In order to show it we have to establish the relation between the class of partition spaces and the canonical α -rank type spaces U_α . The relation of probabilistic type spaces and the universal type space, established by Mertens and Zamir (1985) is relatively simple. Each non-redundant probabilistic type space can be isomorphically embedded in the universal space. In the knowledge case such a simple relation cannot exist, since there is no universal space. We show though, in Theorem 4.3, that every non-redundant partition space can be embedded in some U_α , for big enough α which depends on the rank of the partition space.

We start with the construction, by induction, of the canonical type spaces U_α of rank α . The same construction, using different notations, is carried out in Fagin et al. (1992)

Let S be the set of states of nature and I the set of players. Define U_α by transfinite induction:

$$\begin{aligned}
 U_0 &= S \\
 U_\alpha &= \{(s, (t_\beta^j)_{j \in I}^{\beta < \alpha}) \in S \times \prod_{\beta < \alpha} (2^{U_\beta})^I : \forall \beta < \alpha \quad \forall i \in I \\
 \text{(I)} \quad &(s, (t_\gamma^j)_{j \in I}^{\gamma < \beta}) \in t_\beta^i \\
 \text{(II)} \quad &(\bar{s}, (\bar{t}_\gamma^j)_{j \in I}^{\gamma < \beta}) \in t_\beta^i \Rightarrow \bar{t}_\gamma^i = t_\gamma^i \quad \forall \gamma < \beta \\
 \text{(III)} \quad &\forall \gamma < \beta \text{ the projection of } t_\beta^i \text{ on } U_\gamma \text{ is } t_\gamma^i \}
 \end{aligned}$$

Let $\Phi_{\alpha\beta}$ denote the projection from U_α to U_β for $\beta \leq \alpha$.

For a given $u = (s, (t_\beta^j)_{j \in I}^{\beta < \alpha})$ in U_α , t_β^i is i 's type of order β in u . It is a set of β -order states (i.e., a subset of the space U_β) that i considers possible in state u . The point u specifies a state of nature, and for each player i , a sequence $(t_\beta^i)_{\beta < \alpha}$ of his types for all orders smaller than α . The point u should satisfy for each player i :

I *Correctness*: i 's uncertainties regarding the β -order state of the world, for $\beta < \alpha$, should allow for the actual state, that is, the set t_β^i should contain the β -initial segment of u .

- II *Introspection*: in all β -order states that player i considers possible, for $\beta < \alpha$, (i.e., in all the points of t_β^i) his β -order type is the same as his actual one in u .
- III *Coherence*: the set of γ -order states that i considers possible should be the same at all higher levels. That is, for $\gamma < \beta < \alpha$, t_γ^i should coincide with the set of γ -initial segments of the points in t_β^i .

Note that by condition (III), for each β and i , t_β^i determines t_γ^i for all $\gamma < \beta$ and therefore t_β^i can be identified with the sequence $(t_\gamma^i)_{\gamma < \beta}$.

Conditions (II) and (III) are similar to the ones used in the construction of the universal space of probabilistic beliefs. Condition (I) does not have a corresponding one in this construction. This condition makes the knowledge of different players and state of nature correlated (through the requirement to consider as possible the actual lower-order types of each other). In the Mertens and Zamir (1985) construction there need not be any correlation between the beliefs of different players in a given state.

The canonical α -rank type spaces U_α can be made into a partition space very naturally. Its map to the state of nature is given by the projection on the first coordinate. Player i 's partition, P_α^i , of U_α , partitions it to i 's α -order types. That is, for every $u = (t_\beta^j)_{j \in I} \in U_\alpha$

$$P_\alpha^i(u) = \{(\bar{s}, (\bar{t}_\beta^j)_{j \in I}^{\beta < \alpha}) : \bar{t}_\beta^i = t_\beta^i \ \forall \beta < \alpha\}.$$

Lemma 4.0. *For every $u = (s, (t_\beta^i)_{i \in I}^{\beta < \alpha}) \in U_\alpha$ we have*

$$(u, (P_\alpha^i(u))_{i \in I}) \in U_{\alpha+1}.$$

Proof. We have to prove that $(u, (P_\alpha^i(u))_{i \in I})$ satisfies conditions (I), (II) and (III) for $\beta = \alpha$ in the definition of $U_{\alpha+1}$ (for $\beta < \alpha$, these conditions obtain since $u \in U_\alpha$). (I) and (II) are immediate from the definition of P_α^i . For (III), we have to show that $\Phi_{\alpha\gamma}(P_\alpha^i(u)) = t_\gamma^i$ for every $\gamma < \alpha$.

Condition (I) in the definition of U_α says that for every $\gamma < \alpha$, every $\bar{u}_\gamma \in \Phi_{\alpha\gamma}(P_\alpha^i(u))$ satisfies $\bar{u}_\gamma \in t_\gamma^i$, and hence that $\Phi_{\alpha\gamma}(P_\alpha^i(u)) \subseteq t_\gamma^i$; and condition (II) in the definition of U_α says that for every $\gamma < \alpha$, every $\bar{u}_\gamma \in t_\gamma^i$ satisfies $\bar{u}_\gamma \in \Phi_{\alpha\gamma}(P_\alpha^i(u))$, and hence that $t_\gamma^i \subseteq \Phi_{\alpha\gamma}(P_\alpha^i(u))$. Together this means that $\Phi_{\alpha\gamma}(P_\alpha^i(u)) = t_\gamma^i$, as required. *QED*

Now, knowledge is defined on U_α in two ways. First, the internal structure of each state defines the players' knowledge. Second, the partitions defined on U_α also define the players' knowledge, say by constructing the sequence of β -order events as in section 2. There is a discrepancy, though, between the two. The partition structure determines knowledge completely. That is, we can construct by it β -order events for any ordinal β (although it is unnecessary for β bigger than the rank of the space). The internal structure of a state does not enable such a complete description of knowledge. States in U_α determine all β -order types for $\beta \leq \alpha$ but they do not necessarily determine types of order higher than α . Thus, by Lemma 4.0, $(u, (P_\alpha^i(u))_{i \in I}) \in U_{\alpha+1}$ for each $u \in U_\alpha$. For some u this may be the only possible extension of u to a point in $U_{\alpha+1}$. But for some u in U_α , there may be other extensions to points in $U_{\alpha+1}$, which means that u does not determine the $\alpha + 1$ types of the players.

This motivates the following definition of *subspaces* in U_α as those subsets of U_α where the above mentioned discrepancy does not exist.

Definition 4.1. U is a *subspace* of U_α if $\forall u \in U$:

1. $P_\alpha^i(u) \subseteq U \quad \forall i \in I$
2. $(u, (P_\alpha^i(u))_{i \in I})$ is the only possible extension of u to a point in $U_{\alpha+1}$.

Remark.

1. It is easy to verify by induction, that when U is a subspace of U_α , every $u \in U$ has a unique extension in every U_γ , $\gamma > \alpha$.
2. Note that property (1) in Definition 4.1 guarantees that U is common knowledge in U_α , i.e. $U = C_\alpha(U)$ where C_α is the common knowledge operator on $\langle U_\alpha, (P_\alpha^i)_{i \in I} \rangle$. Property (2) implies further that for the set $\hat{U} = \{(u, (P_\alpha^i(u))_{i \in I}) : u \in U\} \subseteq U_{\alpha+1}$ we have $\hat{U} = C_{\alpha+1}(\hat{U})$.

We define now a notion of a morphism between partition spaces (which have the same set of states of nature) that preserves the natural and knowledge characteristics of states.

Definition 4.2. A partition space $\langle \Omega, S, \Theta: \Omega \rightarrow S, I, (\Pi^i)_{i \in I} \rangle$ is *epimorphic* to the partition space $\langle \Omega, S, \bar{\Theta}: \Omega \rightarrow S, \bar{I}, (\bar{\Pi}^i)_{i \in \bar{I}} \rangle$ if there is an onto function $H: \Omega \rightarrow \bar{\Omega}$ such that for all $\omega \in \Omega$ and $i \in I$, $\Theta(H(\omega)) = \bar{\Theta}(\omega)$, and $\bar{\Pi}^i(H(\omega)) = H(\Pi^i(\omega))$. If H is one-to-one, the two partition spaces are called *isomorphic*.

We now come to the main theorem of this section.

Theorem 4.3. Let $\langle \Omega, S, \Theta: \Omega \rightarrow S, I, (\Pi^i)_{i \in I} \rangle$ be a partition space of rank α , and λ the least limit ordinal greater than α . Then Ω is epimorphic to a subspace U of the canonical λ -order type space U_λ . Moreover, $\Phi_{\lambda\alpha}: U \rightarrow U_\alpha$ is one-to one, and α is the least ordinal with this property. If Ω is non-redundant then the epimorphism of Ω onto U is an isomorphism.

Proof. Define inductively mappings $H_\beta: \Omega \rightarrow U_\beta$:

$$H_0(\omega) = \Theta(\omega)$$

$$H_\beta(\omega) = (H_0(\omega), (H_\gamma(\Pi^i(\omega)))_{\substack{\gamma \leq \beta \\ i \in I}})$$

It is straightforward to verify, by transfinite induction, that indeed $H_\beta(\omega) \in U_\beta$.

Claim 1. Two points ω and ω' in Ω are separated by H_β iff they are separated by P_β .³

Proof of Claim 1. Note that the claim is equivalent to saying that an event in Ω is a β -order event iff it is the inverse image, by H_β , of some event in U_β .

³ P_β was defined in section 2.

We prove the claim by induction. H_0 maps each point to its state of nature, so ω and ω' are separated by H_0 exactly when they are separated by P_0 .

Suppose the claim holds $\forall \gamma < \beta$. Then H_β separates ω and ω' either when they differ by their state of nature, which is exactly the case when they are separated by P_0 and hence also by P_β , or when $H_\gamma(\Pi^i(\omega)) \neq H_\gamma(\Pi^i(\omega'))$ for some $\gamma < \beta$ and $i \in I$. This holds exactly when $\Pi^i(\omega')$ is not a subset of the event $H_\gamma^{-1}(H_\gamma(\Pi^i(\omega)))$. By the induction hypothesis this event is γ -order event, and player i does not know it in ω' while he does in ω . This happens when and only when ω and ω' are separated by $P_{\gamma+1}$ and hence also by P_β . Thus claim 1 is proved.

If two states in Ω are separated by some P_β then they are separated by P_β for every $\beta \geq \alpha$ where α is the rank of Ω . By claim 1 they are therefore mapped by H_β to different states in U_β for every $\beta \geq \alpha$. If, however, two states are not separated by any P_β , then by claim 1 they have the same image by all H_β . We conclude that $\forall \gamma > \beta \geq \alpha$, the projection $\Phi_{\gamma\beta}$ of $H_\gamma(\Omega)$ to U_β is one-to-one.

Let λ be the least limit ordinal greater than α . We now head to show that $H_\lambda(\Omega)$ is a subspace isomorphic to Ω .

Claim 2. For every $\omega \in \Omega$ $(H_\lambda(\omega), H_\lambda(\Pi^i(\omega))_{i \in I})$ is the unique extension of $H_\lambda(\omega)$ to a state in $U_{\lambda+1}$.

Proof of Claim 2. Suppose

$$(H_\lambda(\omega), (t_\lambda^i)_{i \in I}) \in U_{\lambda+1}.$$

Then by property (III) in the definition of $U_{\lambda+1}$, for every $i \in I$ and $\beta < \lambda$ the projection of t_λ^i on U_β is $H_\beta(\Pi^i(\omega))$. This means that for every $u_\lambda^* \in t_\lambda^i$ there is a $u_\lambda^\beta \in H_\lambda(\Pi^i(\omega))$ such that

$$\Phi_{\lambda\beta}(u_\lambda^\beta) = \Phi_{\lambda\beta}(u_\lambda^*).$$

By the conclusion that followed claim 1, for $\alpha \leq \beta < \lambda$ the projection $\Phi_{\lambda\beta}$ of $H_\lambda(\Omega)$ on U_β is one to one. So for every $u_\lambda^* \in t_\lambda^i$ there is a unique $u_\lambda \in H_\lambda(\Pi^i(\omega))$ such that

$$\Phi_{\lambda\beta}(u_\lambda) = \Phi_{\lambda\beta}(u_\lambda^*) \quad \forall \alpha \leq \beta < \lambda.$$

Since λ is a limit ordinal, this simply means that $u_\lambda = u_\lambda^*$. Hence $t_\lambda^i \subseteq H_\lambda(\Pi^i(\omega))$.

On the other hand, we actually have $t_\lambda^i = H_\lambda(\Pi^i(\omega))$, because by property (III) in the definition of $U_{\lambda+1}$,

$$\Phi_{\lambda\alpha}(t_\lambda^i) = H_\alpha(\Pi^i(\omega)) = (\Phi_{\lambda\alpha}(H_\lambda(\Pi^i(\omega))), \tag{1}$$

and for every u_α in (1) there is a unique $u_\lambda \in H_\lambda(\Pi^i(\omega))$ such that

$$u_\alpha = \Phi_{\lambda\alpha}(u_\lambda).$$

Thus claim 2 is proved.

Claim 2 implies that $H_\lambda(\Omega)$ satisfies condition (2) in the definition of a subspace (Definition 4.1). Furthermore, since by Lemma 4.0

$$(H_\lambda(\omega), P_\lambda^i(H_\lambda(\omega))_{i \in I}) \in U_{\lambda+1},$$

we have

$$P_\lambda^i(H_\lambda(\omega)) = H_\lambda(\Pi^i(\omega)) \tag{2}$$

for all $\omega \in \Omega$, which proves condition (1) of that definition. Thus $H_\lambda(\Omega)$ is a subspace of U_λ .

Now note that by definition $\Phi_{\lambda 0}(H_\lambda(\omega)) = \Theta(\omega)$, where $\Phi_{\lambda 0}$ maps U_λ to S . This together with (2) shows that H_λ is an epimorphism of Ω onto $H_\lambda(\Omega)$. If Ω is non-redundant, then by claim 1, H_λ is one-to-one on Ω , and hence an isomorphism. This completes the proof of the theorem. *QED*

Note that although $H_\beta(\Omega)$ for $\alpha \leq \beta < \lambda$ is a one-to-one image of Ω in U_β , it fails in general to be a subspace of U_β . Thus by part (2) of the remark after Definition 4.1, $H_{\beta+1}(\Omega)$ is not common knowledge in $U_{\beta+1}$. λ is the minimal ordinal γ for which it is guaranteed that $H_{\gamma+1}(\Omega)$ is common knowledge in $U_{\gamma+1}$.

The construction of the α -rank spaces W_α in section 3 enables us to show that for $\alpha_1 < \alpha_2$ the projection $\Phi_{\alpha_2 \alpha_1}$ from U_{α_2} to U_{α_1} is many-to-one. Indeed, let λ_1 and λ_2 be the limit ordinals following α_1 and α_2 , respectively. Then by theorem 4.3 for $k = 1, 2$, W_{α_k} can be embedded as a subspace in U_{λ_k} , and the projection of this subspace to U_{α_k} is one-to-one. But the projection from W_{α_2} to W_{α_1} is many-to-one, so $\Phi_{\alpha_2 \alpha_1}$ is many-to-one, at least on the image of W_{α_2} in U_{α_2} .

If the rank of the partition space Ω is a limit ordinal λ , it may sometimes be possible to embed Ω already in U_λ (and not only in $U_{\lambda+\omega}$, as the theorem guarantees). The ‘coordinated attack’ is such an example:

Example 4.4. The day before the campaign two generals agree to attack together either at dawn or at twilight. General A is supposed to send general B a personal messenger with a message about the time. Then B is supposed to send the messenger back with a confirmation, A is supposed to send him once more to confirm the confirmation, and so on. For each decision of general A and for each natural number $n \in \mathbb{N}$, there would be a state for the state of mind of the generals when the messenger made his way successfully exactly n times. We denote the partition of general A by round brackets, that of general B with curly brackets:

$$\begin{array}{l} \text{dawn} \quad \left(\overset{\frown}{\bullet} \right) \quad \{ \bullet \} \quad (\bullet) \quad \{ \bullet \} \quad \dots \\ \text{twilight} \quad \left(\underset{\smile}{\bullet} \right) \quad \{ \bullet \} \quad (\bullet) \quad \{ \bullet \} \quad \dots \\ \quad \quad \quad 0 \quad 1 \quad 2 \quad 3 \quad \dots \end{array}$$

When this partition space is mapped into the canonical spaces with states of nature {dawn, twilight}, the states that correspond to the success of the n th mission of the messenger are first separated in U_{n+1} from those that correspond to the success of the earlier missions. Hence all the states are separated only in U_ω , so ω is the rank of this partition space.

In particular, a partition member of the form $(n - 1, n)$ is mapped in a one-to-one fashion to U_{n+1} and on. This means that the projection of this two-point set from U_ω downwards is one-to-one from the projection to U_{n+1} and on. By the same argument as in the proof of Claim 2 in the theorem, this is the unique belief of the general on U_ω which is coherent with his lower-level beliefs there. Hence the image of the partition space in U_ω is a subspace. In this example the partition space is therefore of rank ω , but embeddable already in U_ω (and not only in $U_{\omega+\omega}$, as the theorem guarantees).

On the other hand, the canonical space U_ω itself is of rank ω , and is only embeddable in $U_{\omega+\omega}$ (because many states in U_ω have multiple extensions to $U_{\omega+1}$, and so on for $\omega + 2, \omega + 3 \dots$). Thus, embedding in the canonical space of the next limit ordinal is the best general result we can attain.

5. The canonical class partition space

When \hat{U}_λ is a subspace of U_λ and $\alpha \geq \lambda$, the function

$$\Psi_{\lambda\alpha} : \hat{U}_\lambda \rightarrow U_\alpha$$

defined inductively by

$$\Psi_{\lambda\alpha}(\hat{u}_\lambda) = (\hat{u}_\lambda, (P^i_\beta(\Psi_{\lambda\beta}(\hat{u}_\lambda)))_{\lambda \leq \beta < \alpha} \quad i \in I)$$

maps each $\hat{u}_\lambda \in \hat{U}_\lambda$ to its unique extension in U_α . Let On be the class of all the ordinals (for an axiomatic treatment of classes see, for instance, Devlin, 1993). For each $\hat{u}_\lambda \in \hat{U}_\lambda$ define the point $\Psi_{\lambda, On}(\hat{u}_\lambda) = (\hat{u}_\lambda, (P^i_\beta(\Psi_{\lambda\beta}(\hat{u}_\lambda)))_{\substack{\lambda \leq \beta \in On \\ i \in I}}$. Using this definition, we can now define the class⁴

$$U_{On} = \{\Psi_{\lambda, On}(\hat{u}_\lambda) : \hat{u}_\lambda \in \hat{U}_\lambda, \hat{U}_\lambda \text{ is a subspace of } U_\lambda, \lambda \in On\}.$$

For all $i \in I$ and for all $u_{On} = (s, (t^j_\beta)_{\substack{\beta \in On \\ j \in I}}) \in U_{On}$ define also the class

$$P^i_{On}(u_{On}) = \{(\bar{s}, (\bar{t}^j_\beta)_{\substack{\beta \in On \\ j \in I}}) \in U_{On} : \bar{t}^i_\beta = t^i_\beta \quad \forall \beta \in On\}. \tag{3}$$

The same argument as in claim 2 in Theorem 4.3 implies that

$$P^i_{On}(\Psi_{\lambda, On}(\hat{u}_\lambda)) = \Psi_{\lambda, On}(P^i_\lambda(\hat{u}_\lambda))$$

(and in particular that (3) is actually a set.) In this sense, \hat{U}_λ can be embedded as a subspace of U_{On} . Therefore, by that theorem, any partition space *which is a set* can be

⁴As written, U_{On} seems to be a family of classes, since each point $\Psi_{\lambda, On}(\hat{u}_\lambda)$ is actually a class. However,

$$D_\lambda = \{\hat{u}_\lambda \in U_\lambda : \hat{u}_\lambda \text{ belongs to some subspace of } U_\lambda\}$$

is a set, and $D = \bigcup_{\lambda \in On} D_\lambda$ is a class. Denote by V the class of all sets. Then U_{On} is the family of classes $(\Psi_{\lambda, On}(\hat{u}_\lambda))_{\hat{u}_\lambda \in D}$, which can be represented as a relation $R \subseteq V \times D$, defined by $(v, \hat{u}_\lambda) \in R$ iff $v = \Psi_{\lambda\beta}(\hat{u}_\lambda)$ for some $\beta \in On$.

Such a relation in the class $V \times D$ is a class by definition.

embedded as a subspace of U_{On} . Hence, it may be convenient to treat U_{On} as the canonical hierarchic construction for (set) partition spaces.

However, if we allow partition spaces to be classes, we shall have to continue the construction further. For instance, if we extend the class of examples in section 3 and allow for a space with class-long sobriety records that have an entry S or D for every ordinal, a state where the players are always Drunk will not be in U_{On} , and will have more than one potential ‘extension’. We will avoid here the technical details, which are, however, completely analogous to those of section 3.

To sum up, U_{On} has indeed the desired canonical properties for set partition spaces. On the other hand it is a class, which can not embed every class partition space. We can not achieve with it ‘self sufficiency’, exactly in the same way we could not with all the previous U_α ’s. This is in sharp contrast with the Mertens–Zamir construction for σ -additive Bayesian types, which achieves this ‘self sufficiency’ in just ω steps.

References

- Armbruster, W., Böge, W., 1979. Bayesian game theory. In: Moeschlin, O., Pallaschke, D. (Eds.), *Game Theory and Related Topics*, North Holland, Amsterdam.
- Böge, W., Eisele, T., 1979. On solutions of Bayesian games. *Int. Jour. Game Theory* 8, 193–215.
- Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *J. Econ. Theory* 59, 189–198.
- Devlin, K., 1993. *The Joy of Sets*, Springer–Verlag.
- Fagin, R., Halpern, J.Y., Vardi, M.Y., 1991. A model-theoretic analysis of knowledge. *J. of the ACM* 91 (2), 382–428.
- Fagin, R., 1994. A quantitative analysis of modal logic. *The Journal of Symbolic Logic* 59, 209–252.
- Fagin, R., Geanakoplos, J., Halpern, J., Vardi, M.Y., 1992. The expressive power of the hierarchical approach to modeling knowledge and common knowledge. In: *Proc. of the 4th Conference on Theoretical Aspects of Reasoning about Knowledge*, Morgan Kaufmann, California, pp. 229–244.
- Harsanyi, J.C. (1967–68), *Games with Incomplete Information Played by Bayesian Players*, parts I, II, & III, *Man. Sc.* 14, pp. 159–182, 320–334, 486–502.
- Heifetz, A., 1993. The Bayesian formulation of incomplete information – The non-compact case. *International Journal of Game Theory* 21, 329–338.
- Heifetz, A., Samet, D., 1998. Knowledge spaces with arbitrarily high rank. *Games and Economic Behavior* 22, 260–273.
- Heifetz, A., 1997. Infinitary $S5$ epistemic logic. *Mathematical Logic Quarterly* 43, 333–342, A preliminary version appeared in the proc. of the 5th Conference on Theoretical Aspects of Reasoning about Knowledge (1994), Morgan Kaufmann, California, pp. 95–107.
- Heifetz, A., Samet, D., 1993. Universal partition structures, IIBR working paper, Vol. 26/93, Faculty of Management, Tel Aviv University.
- Heifetz, A., Samet, D., 1996. Coherent beliefs are not always types, Tel Aviv University, mimeo.
- Mertens, J.F., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. *Int. Jour. Game Theory* 14, 1–29.
- Mertens, J.F., Sorin, S., Zamir, S., 1994. Repeated Games. Part A: Background Material. In: *CORE Discussion Paper*, Vol. 9420, Université Catholique de Louvain.
- Monderer, D., Samet, D., 1989. Approximating common knowledge with common beliefs. *Games and Economic Behavior* 1, 170–190.