

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

**BAYESIAN GAMES:
GAMES WITH INCOMPLETE INFORMATION**

by

SHMUEL ZAMIR

Discussion Paper # 486

June 2008

מרכז לחקר הרציונליות

**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel

PHONE: [972]-2-6584135 FAX: [972]-2-6513681

E-MAIL: ratio@math.huji.ac.il

URL: <http://www.ratio.huji.ac.il/>

Bayesian games: Games with incomplete information¹

Shmuel Zamir

Center for the Study of Rationality
Hebrew University, Israel²
zamir@math.huji.ac.il

June 4, 2008

Article Outline

Glossary

1. Definition of the subject and its importance
2. Introduction: Modeling incomplete information
3. Harsanyi's model: The notion of *type*
4. Aumann's model
5. Harsanyi's model and the hierarchies of beliefs
6. The Universal Belief Space
7. Belief subspaces
8. Consistent beliefs and Common prior
9. Bayesian games and Bayesian equilibrium
10. Bayesian equilibrium and Correlated equilibrium
11. Concluding remarks and future directions
12. Bibliography

Glossary

Bayesian game: An interactive decision situation involving several decision makers (players) in which each player has beliefs about (i.e. assigns probability distribution to) the payoff relevant parameters and the beliefs of the other players.

State of nature: Payoff relevant data of the game such as payoff functions, value of a random variable etc. It is convenient to think of a state of nature as a full description of a 'game-form' (actions and payoff functions).

¹Forthcoming in: *Encyclopedia of Complexity and System Science*, Bob Meyers Editor, Springer 2008

²I am grateful to two anonymous reviewers for their helpful comments.

Type: Also known as *state of mind*, is a full description of player's beliefs (on the state of nature), beliefs on beliefs of the other players, beliefs on the beliefs on his beliefs etc. ad infinitum.

State of the world: A specification of state of nature (payoff relevant parameters) and the players' types (belief of all levels). That is, a state of the world is a state of nature and a list of the states of mind of all player.

Common prior and consistent beliefs: The beliefs of players in a game with incomplete information are said to be consistent if they are derived from the same probability distribution (the common prior) by conditioning on each player's private information. In other words, if the beliefs are consistent, the only source of differences in beliefs is difference in information.

Bayesian equilibrium: A Nash equilibrium of a Bayesian game: A list of behavior and beliefs such that each player is doing his best to maximize his payoff, according to his beliefs about the behavior of the other players.

Correlated equilibrium: A Nash equilibrium in an extension of the game in which there is a chance move, that each player has only partial information about its outcome.

1 Definition of the subject and its importance

Bayesian games (also known as *Games with Incomplete Information*) are models of interactive decision situations in which the decision makers (players) have only partial information about the data of the game and about the other players. Clearly this is typically the situation we are facing and hence the importance of the subject: The basic underlying assumption of classical game theory according to which the data of the game is *common knowledge* (CK) among the players, is too strong and often implausible in real situations. The importance of Bayesian games is in providing the tools and methodology to relax this implausible assumption, to enable modeling the overwhelming majority of real life situations in which players have only partial information about the payoff relevant data. As a result of the interactive nature of the situation, this methodology turned out to be rather deep and sophisticated, both conceptually and mathematically: Adopting the classical Bayesian approach of Statistics, we encounter the need to deal *infinite hierarchy of beliefs*: what each player believes that the other player believes about what he believes....is the actual payoff associated with a certain outcome.? It is not surprising that this methodological difficulty was a major obstacle in the development of the theory, and this article is largely devoted to explain and resolve this methodological difficulty.

2 Introduction: Modeling incomplete information

A game is a mathematical model for an interactive decision situation involving several decision makers (players) whose decisions affect each other. A basic, often implicit, assumption is that the data of the game, which we call the *state of nature*, are *common knowledge* (CK) among the players (see article 22 on *Knowledge*). In particular the actions available

to the players and the payoff functions are CK. This is a rather strong assumption that says that every player knows all actions and payoff functions of all players, every player knows that all other players know all actions and payoff functions, every player knows that every player knows that every player knows.... etc. ad infinitum. *Bayesian games* (also known as *Games with incomplete information*), which are the subject of this article, are models of interactive decision situations in which each player has only partial information about the payoff relevant parameters of the given situation.

Adopting the Bayesian approach, we assume that a player who has only partial knowledge about the state of nature has some *beliefs*, namely prior distribution, on the parameters which he does not know or he is uncertain about. However, unlike in a statistical problem which involves a single decision maker, this is not enough in an interactive situation: As the decisions of other players are relevant, so are their beliefs, since they affect their decisions. Thus a player must have beliefs on the beliefs of other players. By the same reason, a player needs beliefs on the the beliefs of other players on his beliefs and so on. This interactive reasoning about beliefs leads unavoidably to *infinite hierarchies of beliefs* which looks rather intractable. The natural emergence of hierarchies of beliefs is illustrated in the following example:

Example 1. Two players, P1 and P2, play a 2×2 game whose payoffs depend on an unknown state of nature $s \in \{1, 2\}$. Player P1 actions are $\{T, B\}$, player P2 actions are $\{L, R\}$ and the payoffs are given in the following matrices.

		P2	
		L	R
P1	T	0, 1	1, 0
	B	1, 0	0, 1

Payoffs when $s = 1$

		P2	
		L	R
P1	T	1, 0	0, 1
	B	0, 1	1, 0

Payoffs when $s = 2$

Assume that the belief (prior) of P1 on the event $\{s = 1\}$ is p and the belief of P2 on the same event is q . The best action of P1 depends both on his prior and on the action of P2, and similarly for the best action of P2. This is given in the following tables:

		P2's action	
		L	R
$p < 0.5$	T	B	
$p > 0.5$	B	T	

Best reply of P1

		$q < 0.5$	$q > 0.5$
P1's action	T	R	L
	B	L	R

Best reply of P2

Now, since the optimal action of P1 depends not only on his belief p but also on the, unknown to him, action of P2, which depends on his belief q , player P1 must therefore have beliefs on q . These are his *second level beliefs* namely beliefs on beliefs. But then, since this is relevant and unknown to P2, he must have beliefs on that which will be *third level beliefs* of P2, and so on. The whole infinite hierarchies of beliefs of the two players pop out naturally in the analysis of this simple two person game of incomplete information.

The objective of this article is to model this kind of situation. Most of the effort will be

devoted to the modeling of the mutual beliefs structure and only then we add the underlying game which, together with the the beliefs structure, defines a *Bayesian game* for which we define the notion of *Bayesian equilibrium*.

3 Harsanyi's model: The notion of *type*.

As suggested by our introductory example, the straightforward way to describe the mutual belief structure in a situation of incomplete information is to specify explicitly the whole hierarchies of beliefs of the players that is: the beliefs of each player on the unknown parameters of the game, his beliefs on these parameters and on the beliefs of the other players, his beliefs on the parameters, the beliefs of the other players on the parameters and on his own beliefs and so on ad infinitum. This may be called the *explicit* approach and is in fact feasible and was explored and developed at a later stage of the theory (see Mertens and Zamir [13], Aumann [5],[6], Aumann Heifetz [7]). We will come back to it in section 6 when we discuss the *Universal Belief Space*. However, for obvious reasons, the explicit approach is mathematically rather cumbersome and hardly manageable. Indeed this was a major obstacle to the development of the theory of games with incomplete information at its early stages. The breakthrough was provided by John Harsanyi (Harsanyi [10]) in a seminal work that awarded him the Nobel Prize some thirty years later. While Harsanyi actually formulated the problem verbally, in an explicit way, he suggested a solution that 'avoided' the difficulty of having to deal with infinite hierarchies of beliefs, by providing a much more workable *implicit*, encapsulated model which we present now.

The key notion in Harsanyi's model is that of *type*. Each player can be of several types where a type is to be thought of as a full description of the player's beliefs on the *state of nature* (the data of the game) , beliefs on the beliefs of other players on the state of nature and on his own beliefs etc. One may think of a player's type as his/her *state of mind*; a specific configuration of his brain that contains an answer to any question regarding beliefs on the state of nature and on the types of the other players. Note that this implies self reference (of a type to itself through the types of other players) which is unavoidable in an interactive decision situation. A Harsanyi game of incomplete information consists of the following ingredients: (To simplify notations, assume all sets to be finite.)

- I – Players set.
- S – The set of states of nature.
- T_i – The type set of player $i \in I$.

Let $T = \times_{i \in I} T_i$ – denote the *type set* that is, the set type profiles.

- $Y \subset S \times T$, a set of *states of the world*.
- $p \in \Delta(Y)$ – probability distribution on Y , called the *common prior*.

(For a set A , we denote the set of probability distributions on A by $\Delta(A)$.)

Remarks: A state of the world ω consists thus of a state of nature and a list of the types of the players. we denote it as

$$\omega = (s(\omega); t_1(\omega), \dots, t_n(\omega))$$

We think of the state of nature as a full description of the game which we call a *game-form*. So, if it is a game in strategic form, we write the state of nature at state of the world ω as:

$$s(\omega) = (I, (A_i(\omega))_{i \in I}, (u_i(\cdot; \omega))_{i \in I})$$

The payoff functions u_i depend only on the state of nature and not on the types. That is, for all $i \in I$:

$$s(\omega) = s(\omega') \Rightarrow u_i(\cdot; \omega) = u_i(\cdot; \omega').$$

The game with incomplete information is played as follows:

(1) A chance move chooses $\omega = (s(\omega); t_1(\omega), \dots, t_n(\omega)) \in Y$ using the probability distribution p .

(2) Each player is told his chosen type $t_i(\omega)$ (but not the the chosen state of nature $s(\omega)$ and not the other platers' types $t_{-i}(\omega) = (t_j(\omega))_{j \neq i}$.)

(3) The players choose simultaneously an action: player i chooses $a_i \in A_i(\omega)$ and receives a payoff $u_i(a; \omega)$ where $a = (a_1, \dots, a_n)$ is the vector of chosen actions and ω is the state of the world chosen by the chance move.

Remark: The set $A_i(\omega)$, of actions available to player i in state of the world ω must be known to him. Since his only information is his type $t_i(\omega)$, we must impose that $A_i(\omega)$ is T_i -measurable i.e.

$$t_i(\omega) = t_i(\omega') \Rightarrow A_i(\omega) = A_i(\omega').$$

Note that if $s(\omega)$ was commonly known among the players, it would be a regular game in strategic form. We use the term 'game-form' to indicate that the players have only partial information about $s(\omega)$. The players do not know which $s(\omega)$ is being played. In other words, in the extensive form game of Harsanyi, the game-forms $(s(\omega))_{\omega \in Y}$ are not a subgames since they are interconnected by information sets: Player i does not know which $s(\omega)$ is being played since he does not know ω ; he knows only his own type $t_i(\omega)$.

An important application of Harsanyi's model is made in *auction theory*, as auction is a clear situation of incomplete information. For example in a closed private value auction of a single indivisible object, the type of a player is his private value for the object, which is typically known to him and not to other players. We come back to this in section 9.1.1.

4 Aumann's model

A frequently used model of incomplete information was given by Aumann (Aumann [1]).

Definition 2. An Aumann model of incomplete information is $(I, Y, (\pi_i)_{i \in I}, P)$ where:

- I – is the players set.
- Y is a (finite) set whose elements are called *states of the world*.
- For $i \in I$, π_i is a partition of Y .
- P is a probability distribution on Y , also called the *common prior*.

In this model a state of the world $\omega \in Y$ is chosen according to the probability distribution P , and each player i is informed of $\pi_i(\omega)$, the element of his partition that contains the chosen state of the world ω . This is the informational structure which becomes a game with incomplete information if we add a mapping $s : Y \rightarrow S$. The state of nature $s(\omega)$ is the game-form corresponding to the state of the world ω (with the requirement that the action sets $A_i(\omega)$ are π_i -measurable).

It is readily seen that an Aumann's model is a Harsanyi model in which the type set T_i of player i is the set of his partition elements i.e., $T_i = \{\pi_i(\omega) | \omega \in Y\}$, and the common prior on Y is P . Conversely, any Harsanyi model is an Aumann model in which the partitions are those defined by the types i.e., $\pi_i(\omega) = \{\omega' \in Y | t_i(\omega') = t_i(\omega)\}$.

5 Harsanyi's model and hierarchies of beliefs

As our starting point in modeling incomplete information situations was the appearance of hierarchies of beliefs, one may ask how is Harsanyi (or Aumann) model related to hierarchies of beliefs and how does it capture this unavoidable feature of incomplete information situations? The main observation towards answering this question is the following:

Proposition 3. *Any state of the world in Aumann's model or any type profile $t \in T$ in Harsanyi's model defines (uniquely) a hierarchy of mutual beliefs among the players.*

Let us illustrate the idea of the proof by the following example:

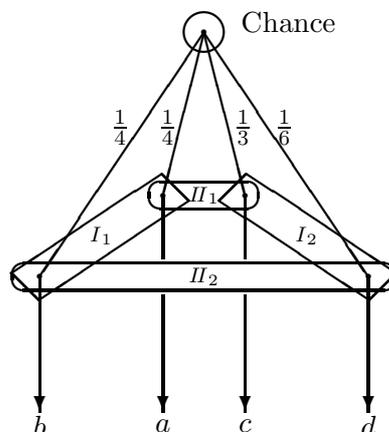
Example

Consider a Harsanyi model with two players, I and II , each of which can be of two types: $T_I = \{I_1, I_2\}$, $T_{II} = \{II_1, II_2\}$ and thus: $T = \{(I_1, II_1), (I_1, II_2), (I_2, II_1), (I_2, II_2)\}$. The probability p on types is given by:

	II_1	II_2
I_1	$\frac{1}{4}$	$\frac{1}{4}$
I_2	$\frac{1}{3}$	$\frac{1}{6}$

Denote the corresponding states of nature by $a = s(I_1 II_1)$, $b = s(I_1 II_2)$, $c = s(I_2 II_1)$ and $d = s(I_2 II_2)$. These are the *states of nature* about which there is incomplete information.

The game in extensive form:



Assume that the state of nature is a ; What are the belief hierarchies of the players?

	Π_1	Π_2
I_1	a $1/4$	b $1/4$
I_2	c $1/3$	d $1/6$

First level belief is obtained by each player from p , by conditioning on his type:

- I_1 : With probability $\frac{1}{2}$ the state is a and with probability $\frac{1}{2}$ the state is b .
- I_2 : With probability $\frac{2}{3}$ the state is c and with probability $\frac{1}{3}$ the state is d .
- Π_1 : With probability $\frac{3}{7}$ the state is a and with probability $\frac{4}{7}$ the state is c .
- Π_2 : With probability $\frac{3}{5}$ the state is b and with probability $\frac{2}{5}$ the state is d .

Second level beliefs (using short notation for the above beliefs: $\left(\frac{1}{2}a + \frac{1}{2}b\right)$ etc.):

- I_1 : With probability $\frac{1}{2}$, player Π believes $\left(\frac{3}{7}a + \frac{4}{7}c\right)$, and with probability $\frac{1}{2}$, player Π believes $\left(\frac{3}{5}b + \frac{2}{5}d\right)$.
- I_2 : With probability $\frac{2}{3}$, player Π believes $\left(\frac{3}{7}a + \frac{4}{7}c\right)$, and with probability $\frac{1}{3}$, player Π believes $\left(\frac{3}{5}b + \frac{2}{5}d\right)$.
- Π_1 : With probability $\frac{3}{7}$, player I believes $\left(\frac{1}{2}a + \frac{1}{2}b\right)$, and with probability $\frac{4}{7}$, player I believes $\left(\frac{2}{3}c + \frac{1}{3}d\right)$.
- Π_2 : With probability $\frac{3}{5}$, player I believes $\left(\frac{1}{2}a + \frac{1}{2}b\right)$, and with probability $\frac{2}{5}$, player I believes $\left(\frac{2}{3}c + \frac{1}{3}d\right)$.

Third level beliefs:

- I_1 : With probability $\frac{1}{2}$, player II believes that: “With probability $\frac{3}{7}$, player I believes $\left(\frac{1}{2}a + \frac{1}{2}b\right)$ and with probability $\frac{4}{7}$, player I believes $\left(\frac{2}{3}c + \frac{1}{3}d\right)$.”
and with probability $\frac{1}{2}$, player II believes that: “With probability $\frac{3}{5}$, player I believes $\left(\frac{1}{2}a + \frac{1}{2}b\right)$ and with probability $\frac{2}{5}$, player I believes $\left(\frac{2}{3}c + \frac{1}{3}d\right)$,”

and so on and so on. The idea is very simple and powerful; Since each player of a given type has a probability distribution (beliefs) both on the types of the other players and on the set S of states of nature, the hierarchies of beliefs are constructed inductively: If the k -th level beliefs (on S) are defined for each type, then the beliefs on types generates the $(k + 1)$ -th level of beliefs.

Thus the compact model of Harsanyi *does* capture the whole hierarchies of beliefs and it is rather tractable. the natural question is whether this model can be used for *all* hierarchy of beliefs? In other words, given any hierarchy of mutual beliefs of a set of players I on a set S of states of nature, can it be represented by a Harsanyi game? This was answered by Mertens and Zamir [18], who constructed the *Universal Belief Space* that is, given a set S of states of nature and a finite set I of players, they looked for the space Ω of all possible hierarchies of mutual beliefs on S among the players in I . This construction is outlined in the next section.

6 The Universal belief space

Given a finite set of players $I = \{1, \dots, n\}$ and a set S of states of nature, which assumed to be compact, we first identify the mathematical spaces in which lie the hierarchies of beliefs. Recall that $\Delta(A)$ denotes the set of probability distributions on A and define inductively the sequence of spaces $(X_k)_{k=1}^{\infty}$ by

$$X_1 = \Delta(S) \tag{1}$$

$$X_{k+1} = X_k \times \Delta(S \times X_k^{n-1}), \text{ for } k = 1, 2, \dots, \tag{2}$$

Any probability distribution on S can be a first level belief and is thus in X_1 . A second level belief is a joint probability distribution on S and the first level beliefs of the other $(n - 1)$ players. this is an element in $\Delta(S \times X_1^{n-1})$ and therefore a two level hierarchy is an element of the product space $X_1 \times \Delta(S \times X_1^{n-1})$, and so on for any level. Note that at each level belief is a joint probability distribution on S and the previous level beliefs, allowing for correlation between the two. In dealing with these spaces we need to have some mathematical structure on these probability spaces. More specifically we make use of the *weak** topology:

Definition 4. A sequence $(F_n)_{n=1}^{\infty}$ of probability measures (on Ω) converges in the *weak** topology to the probability F if and only if $\lim_{n \rightarrow \infty} \int_{\Omega} g(\omega) dF_n = \int_{\Omega} g(\omega) dF$ for all bounded and continuous functions $g : \Omega \rightarrow \mathbb{R}$.

It follows from the compactness of S that all spaces defined by (1)–(2) are compact in the *weak** topology. However, for $k > 1$, not every element of X_k represents a coherent

hierarchy of beliefs of level k . For example, if $(\mu_1, \mu_2) \in X_2$ where $\mu_1 \in \Delta(S) = X_1$ and $\mu_2 \in \Delta(S \times X_1^{n-1})$, then for this to describe meaningful beliefs of a player, the marginal distribution of μ_2 on S must coincide with μ_1 . More generally, any event A in the space of k level beliefs, has to have the same (marginal) probability in any higher level beliefs. Furthermore, not only that each player's beliefs are coherent but he also considers only coherent beliefs of the other players (only those are in the support of his beliefs). Expressing formally this coherency condition yields a selection $T_k \subseteq X_k$ such that $T_1 = X_1 = \Delta(S)$. It is proved that the the projection of T_{k+1} on X_k is T_k (that is, any coherent k level hierarchy can be extended to a coherent $k + 1$ level hierarchy) and that all the sets T_k are compact. Therefore, the *projective limit*, $T = \lim_{\infty \leftarrow k} T_k$ is well defined and nonempty.³

Definition 5. The *universal type set* T is the projective limit of the spaces $(T_k)_{k=1}^{\infty}$.

That is, T is the set of *all coherent infinite hierarchies of beliefs* regarding S , of a player in I . It does not depend on i since, by construction it contains *all possible* hierarchies of beliefs regarding S , and it is therefore the same for all players. It is determined only by S and the number of players n .

Proposition 6. *The universal type space T is compact and satisfies*

$$T \approx \Delta(S \times T^{n-1}) \quad (3)$$

The \approx sign in (3) is to be read as an *isomorphism* and Proposition 6 says that a type of a player can be identified with a joint probability distribution on the state of nature and the types of the other players. The implicit equation (3) reflects the self reference and circularity of the notion of type: The type of a player is his beliefs about the state of nature and all beliefs of other players, in particular, their beliefs on his own beliefs.

Definition 7. The *Universal Belief Space* (UBS) is the space Ω defined by:

$$\Omega = S \times T^n. \quad (4)$$

An element of Ω is called a *state of the world*.

Thus a state of the world is $\omega = (s(\omega); t_1(\omega), t_2(\omega), \dots, t_n(\omega))$ with $s(\omega) \in S$ and $t_i(\omega) \in T$ for all i in I . This is the specifications of the states of nature and the types of all players. The universal belief space Ω is what we looked for: the set of *all* incomplete information and mutual belief configurations of a set of n players regarding the state of nature. In particular, as we will see later, all Harsanyi and Aumann models are embedded in Ω , but it includes also belief configurations that cannot be modeled as Harsanyi games. As we noted before, the UBS is determined only by the set of states of nature S and the set of players I , so it should be denoted as $\Omega(S, I)$. For the sake of simplicity we shall omit the arguments and write Ω unless we wish to emphasize the underlying sets S and I .

The execution of the construction of the UBS according the outline above, involves some non trivial mathematics as can be seen in Mertens and Zamir [18]. The reason is that even

³The projective limit, (also known as the *inverse limit*), of the sequence $(T_k)_{k=1}^{\infty}$ is the space T of all sequences $(\mu_1, \mu_2, \dots) \in \times_{k=1}^{\infty} T_k$ which satisfy: For any $k \in \mathbb{N}$, there is a probability distribution $\nu_k \in \Delta(S \times T_k^{n-1})$ such that $\mu_{k+1} = (\mu_k, \nu_k)$.

with a finite number of states of nature, the space of first level beliefs is a continuum, the second level is the space of probability distributions on a continuum and the third level is the space of probability distributions on the space of probability distributions on a continuum. This requires some structure to these spaces: For a (Borel) measurable event E let $B_i^p(E)$ be the event “player i of type t_i believes that the probability of E is at least p ” that is,

$$B_i^p(E) = \{\omega \in \Omega | t_i(E) \geq p\}$$

Since this is the object of beliefs of players other than i (beliefs of $j \neq i$ on the beliefs of i), this set must also be measurable. Mertens and Zamir used the *weak** topology which is the minimal topology with which the event $B_i^p(E)$ is (Borel) measurable for any (Borel) measurable event E . In this topology, if A is a compact set then $\Delta(A)$, the space of all probability distributions on A , is also compact. However, the hierarchic construction can also be made with stronger topologies on $\Delta(A)$ (see Brandenburger and Dekel [9], Heifetz [12], Mertens, Sorin and Zamir [17]). Heifetz and Samet [14], worked out the construction of the universal belief space without topology, using only a measurable structure (which is implied by the assumption that the beliefs of the players are measurable). All these explicit construction of the belief space are within what is called the *semantic* approach. Aumann [6], provided another construction of a belief system using the *syntactic* approach based on *sentences and logical formulas* specifying explicitly what each player believes about the state of nature, about the beliefs of the other players on the state of nature and so on. For detailed construction see Aumann [6], Heifetz and Mongin [13], and Meier [16]. For comparison of the syntactic with semantic approach see Aumann and Heifetz [7].

7 Belief subspaces

In constructing the universal belief space we implicitly assumed that each player knows his own type since we specified his beliefs only on the state of nature and the beliefs of the *other* players. In view of that and since by (3) a type of player i is a probability distribution on $S \times T^{I \setminus \{i\}}$, we can view a type t_i also as a probability distribution on $\Omega = S \times T^I$ in which the marginal distribution on T_i is a degenerate delta function at t_i that is, if $\omega = (s(\omega); t_1(\omega), t_2(\omega), \dots, t_n(\omega))$ then for all i in I ,

$$t_i(\omega) \in \Delta(\Omega) \quad \text{and} \quad t_i(\omega)[t_i = t_i(\omega)] = 1. \quad (5)$$

In particular it follow that if $Supp(t_i)$ denotes the support of t_i then,

$$\omega' \in Supp(t_i(\omega)) \Rightarrow t_i(\omega') = t_i(\omega). \quad (6)$$

Let $P_i(\omega) = Supp(t_i(\omega)) \subseteq \Omega$. This defines a *possibility correspondence*; at state of the world ω , player i does not consider as possible any point not in $P_i(\omega)$. By (6),

$$P_i(\omega) \cap P_i(\omega') \neq \emptyset \Rightarrow P_i(\omega) = P_i(\omega'),$$

However, unlike in Aumann’s model, P_i does not define a partition of Ω since it is possible that $\omega \notin P_i(\omega)$, and hence the union $\cup_{\omega \in \Omega} P_i(\omega)$ may be strictly smaller than Ω (see Example

7 below on page 16). If $\omega \in P_i(\omega) \subseteq Y$ holds for all ω in some subspace $Y \subset \Omega$, then $(P_i(\omega))_{\omega \in Y}$ is a partition of Y .

As we said, the universal belief space includes all possible beliefs and mutual belief structures over the state of nature. However, in a specific situation of incomplete information, it may well be that only part of Ω is relevant for describing the situation. If the state of the world is ω then clearly all states of the world in $\cup_{i \in I} P_i(\omega)$ are relevant, but this is not all because if $\omega' \in P_i(\omega)$ then all states in $P_j(\omega')$, for $j \neq i$ are also relevant in the considerations of player i . This observation motivates the following definition.

Definition 8. A *belief subspace* (*BL*-subspace) is a closed subset Y of Ω which satisfies:

$$P_i(\omega) \subseteq Y \quad \forall i \in I \quad \text{and} \quad \forall \omega \in Y. \quad (7)$$

A belief subspace is *minimal* if it has no proper subset which is also a belief subspace. Given $\omega \in \Omega$, the belief subspace at ω , denoted by $Y(\omega)$, is the minimal subspace containing ω .

Since Ω is a *BL*-subspace, $Y(\omega)$ is well defined for all $\omega \in \Omega$. A *BL*-subspace is a closed subset of Ω which is also *closed under beliefs* of the players. In any $\omega \in Y$, it contains all states of the world which are relevant to the situation: If $\omega' \notin Y$, then no player believes that ω' is possible, no player believes that any other player believes that ω' is possible, no player believes that any player believes that any player believes.... etc.

Remark 9. The subspace $Y(\omega)$ is meant to be the minimal subspace which is belief closed by all player *at the state* ω . Thus, a natural definition would be: $\tilde{Y}(\omega)$ is the minimal *BL*-subspace containing $P_i(\omega)$ for all i in I . However, if for every player the state ω is not in $P_i(\omega)$ then $\omega \notin \tilde{Y}(\omega)$. Yet, even if it is not in the belief closure of the players, the real state ω is still relevant (at least for the analyst) because it determines the true state of nature that is, it determines the true payoffs of the game. This is the reason for adding the true state of the world ω , even though ‘it may not be in the mind of the players’.

It follows from (5),(6) and (7) that a *BL*-subspace Y has the following structure:

Proposition 10. A closed subset Y of the universal belief space Ω is a *BL*-subspace if and only if it satisfy the following conditions:

1. For any $\omega = (s(\omega); t_1(\omega), t_2(\omega), \dots, t_n(\omega)) \in Y$, and for all i , the type $t_i(\omega)$ is a probability distribution on Y .

2. For any ω and ω' in Y ,

$$\omega' \in \text{Supp}(t_i(\omega)) \Rightarrow t_i(\omega') = t_i(\omega).$$

In fact condition 1. follows directly from Definition 8 while condition 2. follows from the general property of the UBS expressed in (6).

Given a *BL*-subspace Y in $\Omega(S, I)$ we denote by T_i the type set of player i ;

$$T_i = \{t_i(\omega) | \omega \in Y\},$$

and note that unlike in the UBS, in a specific model Y , the type sets are typically *not the same* for all i , and the analogue of (4) is

$$Y \subseteq S \times T_1 \times \dots \times T_n.$$

A BL -subspace is a model of incomplete information about the state of nature. As we saw in Harsanyi's model, in any model of incomplete information about a fixed set S of states of nature, involving the same set of players I , a state of the world ω defines (encapsulates) an infinite hierarchy of mutual beliefs of the players I on S . By the universality of the belief space $\Omega(S, I)$, there is $\omega' \in \Omega(S, I)$ with the same hierarchy of beliefs as that of ω . The mapping of each ω to its corresponding ω' in $\Omega(S, I)$ is called a *belief morphism* as it preserves the belief structure. Mertens and Zamir [18] proved the space $\Omega(S, I)$ is universal in the sense that *any* model Y of incomplete information of the set of players I about the state of nature $s \in S$, can be embedded in $\Omega(S, I)$ via belief morphism $\varphi : Y \rightarrow \Omega(S, I)$ so that $\varphi(Y)$ is a belief subspace in $\Omega(S, I)$. In the following examples we give the BL -subspaces representing some known models.

Examples of belief subspaces.

Example 1. A game with complete information.

If the state of nature is $s_0 \in S$ then in the universal belief space $\Omega(S, I)$, the game is described by a BL -subspace Y consisting of a single state of the world:

$$Y = \{\omega\} \text{ where } \omega = (s_0; [1\omega], \dots, [1\omega]).$$

Here $[1\omega]$ is the only possible probability distribution on Y namely the trivial distribution supported by ω . In particular, the state of nature s_0 (i.e. the data of the game) is commonly known.

Example 2. Commonly known uncertainty about the state of nature.

Assume that the players set is $I = \{1, \dots, n\}$ and there are k states of nature representing, say k possible n -dimensional payoff matrices G_1, \dots, G_k . At the beginning of the game, the payoff matrix is chosen by a chance move according to the probability vector $p = (p_1, \dots, p_k)$ which is commonly known by the players but no player receives any information about the outcome of the chance move. The set of states of nature is $S = \{G_1, \dots, G_k\}$. The situation described above is embedded in the UBS is $\Omega(S, I)$ as the following BL -subspace Y consisting of k states of the world (denoting $p \in \Delta(Y)$ by $[p_1\omega_1, \dots, p_k\omega_k]$):

- $Y = \{\omega_1, \dots, \omega_k\}$
- $\omega_1 = (G_1; [p_1\omega_1, \dots, p_k\omega_k], \dots, [p_1\omega_1, \dots, p_k\omega_k])$
- $\omega_2 = (G_2; [p_1\omega_1, \dots, p_k\omega_k], \dots, [p_1\omega_1, \dots, p_k\omega_k])$
- \dots
- $\omega_k = (G_k; [p_1\omega_1, \dots, p_k\omega_k], \dots, [p_1\omega_1, \dots, p_k\omega_k])$

There is a single type, $[p_1\omega_1, \dots, p_k\omega_k]$, which is the same for all players. It should be emphasized that the type is a distribution on Y (and not just on the states of nature),

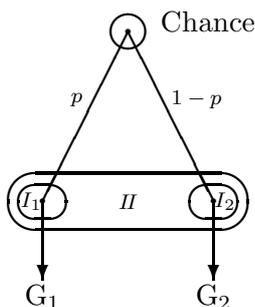
which implies that the beliefs $[p_1G_1, \dots, p_kG_k]$ on the state of nature are commonly known among players.

Example 3. Two players with incomplete information on one side.

There are two players, $I = \{I, II\}$ and two possible payoff matrices, $S = \{G_1, G_2\}$. The payoff matrix is chosen at random with $P(s = G_1) = p$, known to both players. The outcome of this chance move is known only to player I . Aumann and Maschler have studied such situations in which the chosen matrix is played repeatedly and the issue is how does the informed player strategically uses his information (see Aumann and Maschler [8] and its references). This situation is presented in the UBS by the following BL -subspace:

- $Y = \{\omega_1, \omega_2\}$
- $\omega_1 = (G_1; [1\omega_1], [p\omega_1, (1-p)\omega_2])$
- $\omega_2 = (G_2; [1\omega_2], [p\omega_1, (1-p)\omega_2])$

Player I has two possible types, $I_1 = [1\omega_1]$ when he is informed of G_1 and $I_2 = [1\omega_2]$ when he is informed of G_2 . Player II has only one type, $II = [p\omega_1, (1-p)\omega_2]$. We describe this situation in the following ‘extensive form like’ figure in which the oval forms describe the types of the players in the various vertices.



Example 4. Incomplete information about the information others.

In the next example, taken from Sorin and Zamir [23], one of two players always knows the state of nature but he may be uncertain whether the other player knows it. There are two players, $I = \{I, II\}$ and two possible payoff matrices, $S = \{G_1, G_2\}$. It is commonly known to both players that the payoff matrix is chosen at random by a toss of a fair coin; $P(s = G_1) = 1/2$. The outcome of this chance move is told to player I . In addition, if (and only if) the matrix G_1 was chosen, another fair coin toss determines whether to inform player II which payoff matrix was chosen. In any case player I is *not told* the result of the second coin toss. This situation is described by the following belief space with three states of the world.

- $Y = \{\omega_1, \omega_2, \omega_3\}$
- $\omega_1 = (G_1; [\frac{1}{2}\omega_1, \frac{1}{2}\omega_2], [1\omega_1])$
- $\omega_2 = (G_1; [\frac{1}{2}\omega_1, \frac{1}{2}\omega_2], [\frac{1}{3}\omega_2, \frac{2}{3}\omega_3])$

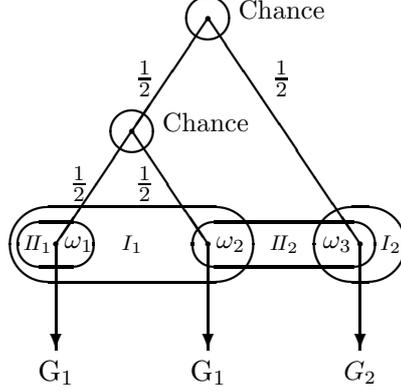
- $\omega_3 = (G_2; [1\omega_3], [\frac{1}{2}\omega_2, \frac{1}{2}\omega_3])$

Each player has two types and the type sets are:

$$T_I = \{I_1, I_2\} = \{[\frac{1}{2}\omega_1, \frac{1}{2}\omega_2], [1\omega_3]\}$$

$$T_{II} = \{II_1, II_2\} = \{[1\omega_1], [\frac{1}{3}\omega_2, \frac{2}{3}\omega_3]\}.$$

Note that in all our examples of belief subspaces, condition (6) is satisfied; the support of a player's type contains only states of the world in which he has that type. The game with incomplete information is described in the following figure.



Example 5. Incomplete information on two sides: A Harsanyi game.

In this example, the set of players is again $I = \{I, II\}$ and the set of states of nature is $S = \{s_{11}, s_{12}, s_{21}, s_{22}\}$. In the universal belief space $\Omega(S, I)$ consider the following *BL*-subspace consisting of four states of the world.

- $Y = \{\omega_{11}, \omega_{12}, \omega_{21}, \omega_{22}\}$
- $\omega_{11} = (s_{11}; [\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12}], [\frac{3}{5}\omega_{11}, \frac{2}{5}\omega_{21}])$
- $\omega_{12} = (s_{12}; [\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12}], [\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22}])$
- $\omega_{21} = (s_{21}; [\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22}], [\frac{3}{5}\omega_{11}, \frac{2}{5}\omega_{21}])$
- $\omega_{22} = (s_{22}; [\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22}], [\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22}])$

Again, each player has two types and the type sets are:

$$T_I = \{I_1, I_2\} = \{[\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12}], [\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22}]\}$$

$$T_{II} = \{II_1, II_2\} = \{[\frac{3}{5}\omega_{11}, \frac{2}{5}\omega_{21}], [\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22}]\}$$

The type of a player determines his beliefs about the type of the other player. For example player I of type I_1 assigns probability $3/7$ to the state of the world ω_{11} in which player II is of type II_1 , and probability $4/7$ to the state of the world ω_{12} in which player II is of type II_2 . Therefore, the beliefs of type I_1 on the types of player II are: $P(II_1) = 3/7$, $P(II_2) = 4/7$. The mutual beliefs on each other type are given in the following tables.

	II_1	II_2
I_1	$3/7$	$4/7$
I_2	$2/3$	$1/3$

	II_1	II_2
I_1	$3/5$	$4/5$
I_2	$2/5$	$1/5$

Beliefs of player I

Beliefs of player II

These are precisely the beliefs of Bayesian players if the pair of types (t_I, t_{II}) in $T = T_I \times T_{II}$ is chosen according to the prior probability distribution p below, and each player is then informed of his own type:

	II_1	II_2
I_1	0.3	0.4
I_2	0.2	0.1

Prior distribution p on T .

In other words, this BL -subspace is a Harsanyi game with type sets T_I, T_{II} and the prior probability distribution p on the types. Actually, as there is one to one mapping between the type set T and the set S of states of nature, the situation is generated by a chance move choosing the state of nature $s_{ij} \in S$ according to the distribution p (that is $P(s_{ij}) = P(I_i, II_j)$ for i and j in $\{1, 2\}$) and then, player I is informed of i and player II is informed of j . As a matter of fact, *all* the BL -subspaces in the previous examples can also be written as Harsanyi games, mostly in a trivial way.

Example 6. Inconsistent beliefs.

In the same universal belief space, $\Omega(S, I)$ of the previous example, consider now another BL -subspace \tilde{Y} which differs from Y only by changing the type II_1 of player II from $[\frac{3}{5}\omega_{11}, \frac{2}{5}\omega_{21}]$ to $[\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}]$ that is,

- $\tilde{Y} = \{\omega_{11}, \omega_{12}, \omega_{21}, \omega_{22}\}$
- $\omega_{11} = (s_{11}; [\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12}], [\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}])$
- $\omega_{12} = (s_{12}; [\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12}], [\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22}])$
- $\omega_{21} = (s_{21}; [\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22}], [\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}])$
- $\omega_{22} = (s_{22}; [\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22}], [\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22}])$

with type sets:

$$T_I = \{I_1, I_2\} = \{[\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12}], [\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22}]\}$$

$$T_{II} = \{II_1, II_2\} = \{[\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}], [\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22}]\}.$$

Now, the mutual beliefs on each other's type are

	II_1	II_2
I_1	3/7	4/7
I_2	2/3	1/3

Beliefs of player I

	II_1	II_2
I_1	1/2	4/5
I_2	1/2	1/5

Beliefs of player II

Unlike in the previous example, these beliefs *cannot be derived from a prior distribution* p . According to Harsanyi, these are *inconsistent beliefs*. A BL -subspace with inconsistent beliefs, *cannot be described* as a Harsanyi or Aumann model; It cannot be described as a game in extensive form.

Example 7. ‘Highly inconsistent’ beliefs.

In the previous example, even though the beliefs of the players were inconsistent, in all states of the world, the true state was considered possible by all players (for example in the state ω_{12} player I assigns to this state probability $4/7$ and player II assigns to it probability $4/5$). As was emphasized before, the UBS contains *all* belief configurations including ‘highly inconsistent’ or ‘wrong’ beliefs as the following example shows. The belief subspace of the two players I and II concerning the state of nature which can be s_1 or s_2 is given by:

- $Y = \{\omega_1, \omega_2\}$
- $\omega_1 = (s_1; [\frac{1}{2}\omega_1, \frac{1}{2}\omega_2], [1\omega_2])$
- $\omega_2 = (s_2; [\frac{1}{2}\omega_1, \frac{1}{2}\omega_2], [1\omega_2])$

In the state of the world ω_1 , the state of nature is s_1 , player I assigns equal probabilities to s_1 and s_2 but player II assigns probability 1 to s_2 . In other words, he does not consider as possible the true state of the world (and also the true state of nature): $\omega_1 \notin P_{II}(\omega_1)$ and consequently $\cup_{\omega \in Y} P_{II}(\omega) = \{\omega_2\}$ which is strictly smaller than Y . By the definition of belief subspace and condition (6), this also implies that $\cup_{\omega \in \Omega} P_{II}(\omega)$ is strictly smaller than Ω (as it does not contain ω_1).

8 Consistent beliefs and common priors

A BL -subspace Y is a semantic belief system presenting, via the notion of types, the hierarchies of belief of a set of players having incomplete information about the state of nature. A state of the world captures the situation at what is called *the interim stage*: Each player knows his own type and has beliefs about the state of nature and the types of the other players. The question ‘what is the *real state of the world* ω ’? is not addressed. In a BL -subspace, there is no chance move with explicit probability distribution that chooses the state of the world, while such a probability distribution is part of a Harsanyi or an Aumann model. Yet, in the belief space Y of Example 5 in the previous section, such a *prior distribution* p emerged *endogenously* from the structure of Y . More specifically, If the state $\omega \in Y$ is chosen by a chance move according to the probability distribution p and each player i is told his type $t_i(\omega)$, then his beliefs are *precisely* those described by $t_i(\omega)$. This is a property of the BL -subspace that we call *consistency* (that does not hold, for instance, for the BL -subspace \tilde{Y} in Example 6) and which we define now: Let $Y \subseteq \Omega$ be a BL -subspace.

Definition 11. (i) A probability distribution $p \in \Delta(Y)$ is said to be *consistent* if for any player $i \in I$,

$$p = \int_Y t_i(\omega) dp. \tag{8}$$

(ii) A BL -subspace Y is said to be consistent if there is a consistent probability distribution p with $Supp(p) = Y$. A consistent BL -subspace will be called a C -subspace. A state of the world $\omega \in \Omega$ is said to be consistent if it is a point in a C -subspace.

The interpretation of equation (8) is that the probability distribution p is ‘the average’ of the types $t_i(\omega)$ of player i (which are also a probability distributions on Y), when the

average is taken on Y according to p . This definition is not transparent; it is not clear how it captures the consistency property we have just explained, in terms of a chance move choosing $\omega \in Y$ according to p . However, it turns out to be equivalent.

For $\omega \in Y$ denote $\pi_i(\omega) = \{\omega' \in Y | t_i(\omega') = t_i(\omega)\}$, then we have:

Proposition 12. *A probability distribution $p \in \Delta(Y)$ is consistent if and only if,*

$$t_i(\omega)(A) = p(A|\pi_i(\omega)), \quad (9)$$

holds for all $i \in I$ and for any measurable set $A \subseteq Y$.

In particular, a Harsanyi or an Aumann model is represented by a consistent BL -subspace since, by construction, the beliefs are derived from a common prior distribution which is part of the data of the model. The role of the prior distribution p in these models is actually not that of an additional parameter of the model but rather that of an additional assumption on the belief system namely, the *consistency assumption*. In fact, if a minimal belief subspace is consistent then, the common prior p is uniquely determined by the beliefs as we saw in example 5; *there is no need to specify p as an additional data of the system.*

Proposition 13. *If $\omega \in \Omega$ is a consistent state of the world, and if $Y(\omega)$ is the smallest consistent BL -subspace containing ω , then the consistent probability distribution p on $Y(\omega)$ is uniquely determined.*

(The formulation of this proposition requires some technical qualification if $Y(\omega)$ is a continuum.)

The consistency (or the existence of common prior), is a quite strong assumption. It assumes that differences in beliefs (i.e. in probability assessments) are due *only* to differences in information; players having precisely *the same information* will have precisely *the same beliefs*. It is no surprise that this assumption has strong consequences the most known of which is due to Aumann [1]: Players with consistent beliefs *cannot agree to disagree* that is: If, at some state of the world, it is commonly known that one player assigns probability q_1 to an event E and another player assigns probability q_2 to the same event, then it must be the case that $q_1 = q_2$. Variants of this result appear under the title of “*No trade theorems*” (see e.g. Milgrom and Stokey [14]): Rational players with consistent beliefs cannot both believe that they can gain from a trade or a bet between them.

The plausibility and the justification of the common prior assumption was extensively discussed in the literature (see e.g. Harsanyi [10], Gul [10], Aumann [4]). It is sometimes referred to in the literature as the *Harsanyi doctrine*. Here we only make the observation that within the set of BL -subspaces in Ω , the set of consistent BL -subspaces is a set of measure zero. To see the idea of the proof, consider the Following example:

Example 8. Generalization of Examples 5 and 6.

Consider a BL -subspace as in Examples 5 and 6 but with type sets:

$$T_I = \{I_1, I_2\} = \{[\alpha_1\omega_{11}, (1 - \alpha_1)\omega_{12}], [\alpha_2\omega_{21}, (1 - \alpha_2)\omega_{22}]\}$$

$$T_{II} = \{II_1, II_2\} = \{[\beta_1\omega_{11}, (1 - \beta_1)\omega_{21}], [\beta_2\omega_{12}, (1 - \beta_2)\omega_{22}]\}.$$

belief space $\Omega(S, I)$. Now we add the the action sets and the payoff functions. These are actually part of the description of the state of nature: The mapping $s : \Omega \rightarrow S$ assigns to each state of the world ω the ‘game-form’ $s(\omega)$ played at this state. To emphasize this interpretation of $s(\omega)$ as a game-form, we denote it also as Γ_ω :

$$\Gamma_\omega = (I, A_i(t_i(\omega))_{i \in I}, (u_i(\omega))_{i \in I})$$

Where $A_i(t_i(\omega))$ is the actions set (pure strategies) of player i at ω and $u_i(\omega) : A(\omega) \rightarrow \mathbb{R}$ is his payoff function, where $A(\omega) = \times_{i \in I} A_i(t_i(\omega))$, is the set of action profiles at state ω . Note that while the actions of a player depends only on his type, his payoff depends on the actions and types of all players. For a vector of actions $a \in A(\omega)$, we write $u_i(\omega; a)$ for $u_i(\omega)(a)$. Given a BL -subspace $Y \subseteq \Omega(S, I)$ we define *the Bayesian game on Y* as follows:

Definition 14. The Bayesian game on Y is a vector payoff game in which:

- $I = \{1, \dots, n\}$ – the players set.
- Σ_i – the strategy set of player i , is the set of mappings

$$\sigma_i : Y \longrightarrow A_i \text{ which are } T_i\text{-measurable.}$$

In particular:

$$t_i(\omega_1) = t_i(\omega_2) \implies \sigma_i(\omega_1) = \sigma_i(\omega_2).$$

Let $\Sigma = \times_{i \in I} \Sigma_i$.

- The payoff function u_i for player i is a *vector valued function*, $u_i = (u_{t_i})_{t_i \in T_i}$ where u_{t_i} (the payoff function of player i of type t_i) is a mapping

$$u_{t_i} : \Sigma \longrightarrow \mathbb{R}$$

defined by:

$$u_{t_i}(\sigma) = \int_Y u_i(\omega; \sigma(\omega)) dt_i(\omega) \tag{11}$$

Note that u_{t_i} is T_i -measurable as it should be. When Y is a finite BL -subspace, the above defined Bayesian game is an n -person ‘game’ in which the payoff for player i is a vector with a payoff for each one of his types (therefore, a vector of dimension $|T_i|$). It becomes a ‘regular game-form’ for a given state of the world ω since then the payoff to player i is $u_{t_i(\omega)}$. However, these game-forms are not regular games since they are ‘interconnected’; the players do not know which of these ‘games’ they are playing (since they do not know the state of the world ω). Thus, just like a Harsanyi game, a Bayesian game on a BL -subspace Y consists of a family of ‘connected’ game-forms, one for each $\omega \in Y$. However, unlike a Harsanyi game, a Bayesian game has no chance move that chooses the state of the world (or the vector of types). A way to transform a Bayesian game into a regular game was suggested by R. Selten and was named by Harsanyi as the “Selten game G^{**} ” (see Harsanyi [11] p. 496). This is a game with $|T_1| \cdot |T_2| \cdot \dots \cdot |T_n|$ players (one for each type) in which each player $t_i \in T_i$ chooses a strategy and then *selects* his $(n - 1)$ partners, one from each T_j ; $j \neq i$, according to his beliefs t_i .

9.1 Bayesian equilibrium.

Although a Bayesian game is not a regular game, the Nash equilibrium concept based on the notion of *best reply* can be adapted yielding the solution concept of *Bayesian Equilibrium* (also called *Nash-Bayes Equilibrium*).

Definition 15. A vector of strategies $\sigma = (\sigma_1, \dots, \sigma_n)$, in a Bayesian game, is called a *Bayesian equilibrium* if for all i in I and for all t_i in T_i ,

$$u_{t_i}(\sigma) \geq u_{t_i}(\sigma_{-i}; \tilde{\sigma}_i), \quad \forall \tilde{\sigma}_i \in \Sigma_i, \quad (12)$$

where, as usual, $\sigma_{-i} = (\sigma_j)_{j \neq i}$ denotes the vector of strategies of the players other than i .

Thus, a Bayesian equilibrium specifies a behavior for each player which is a *best reply* to what he believes is the behavior of the other players that is, a best reply the strategies of the other players given his type. In a game with complete information, which corresponds to a *BL*-subspace with one state of the world ($Y = \{\omega\}$), as there is only one type of each player, and the beliefs are all probability one on a singleton, the Bayesian equilibrium is just the well known Nash equilibrium.

Remark 16. It is readily seen that when Y is finite, any Bayesian equilibrium is a Nash equilibrium of the “Selten game G^{**} ” in which each type is a player who selects the types of his partners according to his beliefs. Similarly, we can transform the Bayesian game into an ordinary game in strategic form by defining the payoff function to player i to be $\tilde{u}_i = \sum_{t_i \in T_i} \gamma_{t_i} u_{t_i}$ where γ_{t_i} are strictly positive. Again, independently of the values of the constants γ_{t_i} , any Bayesian equilibrium is a Nash equilibrium of this game and vice versa. In particular if we choose the constants so that $\sum_{t_i \in T_i} \gamma_{t_i} = 1$, we obtain the game suggested by Aumann and Maschler in 1967 (see Aumann and Maschler [8], p. 95) and again, the set of Nash equilibria of this game is precisely the set of Bayesian equilibria.

The Harsanyi game revisited

As we observed in Example 5. on page 14, the belief structure of a consistent *BL*-subspace is the same as in a Harsanyi game *after the chance move choosing the types*. That is, the embedding of the Harsanyi game as a *BL*-subspace in the universal belief space is only at *the interim stage*, *after* the moment that each player gets to know his type. The Harsanyi game on the other hand is at *the ex ante stage*, *before* a player knows his type. Then, what is the relation between the Nash equilibrium in the Harsanyi game, at the ex ante stage, and the equilibrium at the interim stage, namely the Bayesian equilibrium of the corresponding *BL*-subspace? This is an important question to the embedding of the Harsanyi game in the UBS since, as we said before, the chance move choosing the types *does not appear explicitly in the UBS*. The answer to this question was given by Harsanyi (1967-8) (assuming that each type t_i has a positive probability):

Theorem 17. (*Harsanyi*) *The set of Nash equilibria of a Harsanyi game is identical to the set of Bayesian equilibria of the equivalent BL-subspace in the UBS.*

In other words, this theorem states that any equilibrium in the ex ante stage is also an equilibrium at the interim stage and vice versa.

In modeling situations of incomplete information, the interim stage is the natural one; A player *knows his beliefs (type)*, then why should he analyze the situation, as Harsanyi suggests, from the ex ante point as if his type was not known to him and he could be as well of other type? Theorem 17 provides a technical answer to this question: The equilibria are *the same* in both games and the equilibrium strategy of the ex ante game specifies for each type precisely his equilibrium strategy at the interim stage. In that respect, for a player who knows his type, the Harsanyi model is just *an auxiliary game* to compute his equilibrium behavior. Of course the deeper answer to the question above comes from the interactive nature of the situation: Even though player i knows he is of type t_i , he knows that his partners do not know that, and they may consider the possibility that he is of type \tilde{t}_i , and since this affects their behavior, the behavior of type \tilde{t}_i is also relevant to player i who knows he is of type t_i . Finally, Theorem 17 makes the Bayesian equilibrium the natural extension of the Nash equilibrium concept to games with incomplete information for consistent or inconsistent beliefs, when the Harsanyi ordinary game model is not available.

9.1.1 Examples of Bayesian equilibria.

In Example 6 on page 15, there two players of two types each, and with *inconsistent* mutual beliefs given by

	Π_1	Π_2
I_1	3/7	4/7
I_2	2/3	1/3

Beliefs of player I

	Π_1	Π_2
I_1	1/2	4/5
I_2	1/2	1/5

Beliefs of player II

Assume that payoff matrices for the four type profiles are:

		II	
		L	R
I	T	2, 0	0, 1
	B	0, 0	1, 0

G_{11} : Payoffs when $t = (I_1, II_1)$

		II	
		L	R
I	T	0, 0	0, 0
	B	1, 1	1, 0

G_{12} : Payoffs when $t = (I_1, II_2)$

		II	
		L	R
I	T	0, 0	0, 0
	B	1, 1	0, 0

G_{21} : Payoffs when $t = (I_2, II_1)$

		II	
		L	R
I	T	0, 0	2, 1
	B	0, 0	0, 2

G_{22} : Payoffs when $t = (I_2, II_2)$

As the beliefs are inconsistent they cannot be presented by a Harsanyi game yet, we can compute the Bayesian equilibrium of this Bayesian game. Let (x, y) the strategy of player I which is:

- Play the mixed strategy $[x(T), (1 - x)(B)]$ when you are of type I_1 .
- Play the mixed strategy $[y(T), (1 - y)(B)]$ when you are of type I_2 .

and let (z, t) the strategy of player II which is:

- Play the mixed strategy $[z(L), (1 - z)(R)]$ when you are of type II_1 .
- Play the mixed strategy $[t(L), (1 - t)(R)]$ when you are of type II_2 .

For $0 < x, y, z, t < 1$, each player of each type must be indifferent between his two pure actions which yields the values in equilibrium:

$$x = \frac{3}{5}, \quad y = \frac{2}{5}, \quad z = \frac{7}{9}, \quad t = \frac{2}{9}$$

There is no ‘expected payoff’ since this is a Bayesian game and not a game; the expected payoffs depend on the *actual* state of the world i.e., the actual types of the players and the actual payoff matrix. For example, the state of the world is $\omega_{11} = (G_{11}; I_1, II_1)$, the expected payoffs are:

$$\pi(\omega_{11}) = \left(\frac{3}{5}, \frac{2}{5} \right) G_{11} \begin{pmatrix} 7/9 \\ 2/9 \end{pmatrix} = \left(\frac{46}{45}, \frac{6}{45} \right)$$

Similarly:

$$\pi(\omega_{12}) = \left(\frac{3}{5}, \frac{2}{5} \right) G_{12} \begin{pmatrix} 2/9 \\ 7/9 \end{pmatrix} = \left(\frac{18}{45}, \frac{4}{45} \right)$$

$$\pi(\omega_{21}) = \left(\frac{2}{5}, \frac{3}{5} \right) G_{21} \begin{pmatrix} 7/9 \\ 2/9 \end{pmatrix} = \left(\frac{21}{45}, \frac{21}{45} \right)$$

$$\pi(\omega_{22}) = \left(\frac{2}{5}, \frac{3}{5} \right) G_{22} \begin{pmatrix} 2/9 \\ 7/9 \end{pmatrix} = \left(\frac{28}{45}, \frac{70}{45} \right)$$

However, these are the *objective* payoffs as viewed by the analyst; they are viewed differently by the players. For player i of type t_i the relevant payoff is his subjective payoff $u_{t_i}(\sigma)$ defined in equation (11) (on page 19). For example, at state ω_{11} (or ω_{12}) player I believes that with probability $3/7$ the state is ω_{11} in which case his payoff is $46/45$ and with probability $4/7$ the state is ω_{12} in which case his payoff is $18/45$. Therefore his *subjective* expected payoff at state ω_{11} is $3/7 \times 46/45 + 4/7 \times 18/45 = 2/3$. Similar computations show that in states ω_{21} or ω_{22} player I ‘expects’ a payoff of $7/15$ while player II ‘expects’ $3/10$ at states ω_{11} or ω_{21} and $86/225$ in states ω_{12} or ω_{22} .

Bayesian equilibrium is widely used in *Auction Theory* which constitute of an important and successful applications of the theory of games with incomplete information. The simplest example is that of two buyers bidding in a first price auction for an indivisible object. If each buyer i has a private value v_i for the object (which is independent of the private value v_j of the other buyer), and if he further believes that v_j is random with uniform probability distribution on $[0, 1]$, then this is a Bayesian game in which the type of a player is his private valuation that is, the type sets are $T_1 = T_2 = [0, 1]$ which is a continuum. This is a consistent Bayesian game (that is, a Harsanyi game) since the beliefs are derived from the uniform probability distribution on $T_1 \times T_2 = [0, 1]^2$. A Bayesian equilibrium of this game is that in which each player bids half of his private value: $b_i(v_i) = v_i/2$ (see e.g. Wolfstetter [25], chapter III). Although auction theory was developed far beyond this simple example, almost all models studied so far, are Bayesian games with consistent beliefs, actually Harsanyi games. The main reason of course is that consistent Bayesian games are more manageable as they can be described as an equivalent ordinary game in strategic form. However, inconsistent beliefs are rather plausible and existing in the market place in general and even more so in auction situations. An example of that is in case of collusion

of bidders: When a bidding ring is formed, it may well be the case that some of the bidders outside the ring are not aware of its existence and behave under the belief that all bidders are competitive. The members of the ring may or may not know whether or not the other bidders know about the ring, or they may be uncertain about it. This rather plausible mutual belief situation is typically inconsistent and has to be treated as an inconsistent Bayesian game for which a Bayesian equilibrium is to be found.

9.2 Bayesian equilibrium and correlated equilibrium

Correlated equilibrium was introduced in Aumann (1974) as the Nash equilibrium of a game extended by adding to it random events about which the players have partial information. Basically, starting from an ordinary game, Aumann added a probability space and information structure and obtained a game with incomplete information the equilibrium of which he called a correlated equilibrium of the original game. Since the Nash equilibrium of a game with incomplete information is the Bayesian equilibrium, this suggests that concept of correlated equilibrium is closely related to that of Bayesian equilibrium. In fact Aumann noticed that and discussed it in a second paper entitled “Correlated Equilibrium as an Expression of Bayesian Rationality” (Aumann [3]). In this section we remind briefly, through an example the concept of Correlated equilibrium and state formally its relation to the concept of Bayesian equilibrium.

Example 18. Consider a two person game with actions $\{T, B\}$ for player 1 and $\{L, R\}$ for player 2 with corresponding payoffs given in the following matrix:

		2	
		L	R
1	T	6, 6	2, 7
	B	7, 2	0, 0

G : Payoffs of the basic game.

This game has three Nash equilibria: (T, R) with payoff $(2, 7)$, (B, L) with payoff $(7, 2)$ and the mixed equilibrium $([\frac{2}{3}(T), \frac{1}{3}(B)], [\frac{2}{3}(L), \frac{1}{3}(R)])$ with payoff $(4\frac{2}{3}, 4\frac{2}{3})$. Suppose that we add to the game a chance move that chooses an element in $\{T, B\} \times \{L, R\}$ according to the following probability distribution μ :

		L	R
		1/3	1/3
B	1/3	0, 0	

μ : probability distribution on $\{T, B\} \times \{L, R\}$.

Let us now extend the game G to a game with incomplete information G^* in which a chance move chooses an element in $\{T, B\} \times \{L, R\}$ according to the probability distribution above. Then player 1 is informed of the first (left) component of the chosen element and player 2 is informed of the second (right) component. Then each player chooses an action in G and the payoff is made. If we interpret the partial information as a suggestion which action to choose, then it is readily verified that following the suggestion is a Nash equilibrium of the

extended game yielding a payoff (5, 5). This was called by Aumann a correlated equilibrium of the original game G . In our terminology, the extended game G^* is a Bayesian game and its Nash equilibrium is its Bayesian equilibrium. Thus what we have here is that a correlated equilibrium of a game is just the Bayesian equilibrium of its extension to a game with incomplete information. We now make this a general formal statement. For simplicity, we use the Aumann model of game with incomplete information.

Let $G = (I, (A_i)_{i \in I}, (u_i)_{i \in I})$ be a game in strategic form where I is the set of players, A_i is the set of actions (pure strategies) of player i and u_i is his payoff function.

Definition 19. Given a game in strategic form G , an incomplete information extension (the I -extension) of the game G is the game G^* given by:

$$G^* = (I, (A_i)_{i \in I}, (u_i)_{i \in I}, (Y, p)), (\pi_i)_{i \in I},$$

where (Y, p) is a finite probability space and π_i is a partition of Y (the information partition of player i).

This is an Aumann model of incomplete information and, as we noted before, it is also a Harsanyi type based model in which the type of player i at state $\omega \in Y$ is $t_i(\omega) = \pi_i(\omega)$, and a strategy of player i is a mapping from his type set to his mixed actions: $\sigma_i : T_i \rightarrow \Delta(A_i)$.

We identify a correlated equilibrium in the game G by the probability distribution μ on the vectors of actions $A = A_1 \times \dots \times A_n$. Thus $\mu \in \Delta(A)$ is a correlated equilibrium of the game G if when $a \in A$ is chosen according to μ and each player i is suggested to play a_i , his best reply is in fact to play the action a_i .

Given a game with incomplete information G^* as in definition 19, any vector of strategies of the players $\sigma = (\sigma_1, \dots, \sigma_n)$ induces a probability distribution on the vectors of actions $a \in A$. we denote this as $\mu_\sigma \in \Delta(A)$.

We can now state the relation between Correlated and Bayesian equilibria:

Theorem 20. *Let σ be a Bayesian equilibrium in the game of incomplete information $G^* = (I, (A_i)_{i \in I}, (u_i)_{i \in I}, (Y, p)), (\pi_i)_{i \in I}$, then the induced probability distribution μ_σ is a correlated equilibrium of the basic game $G = (I, (A_i)_{i \in I}, (u_i)_{i \in I})$.*

The other direction is:

Theorem 21. *Let μ be a correlated equilibrium of the game $G = (I, (A_i)_{i \in I}, (u_i)_{i \in I})$, then G has an extension to a game with incomplete information $G^* = (I, (A_i)_{i \in I}, (u_i)_{i \in I}, (Y, p)), (\pi_i)_{i \in I}$ with a Bayesian equilibrium σ for which $\mu_\sigma = \mu$.*

10 Concluding remarks and future directions

10.1 The consistency assumption

To the heated discussion of the merits and justification of the consistency assumption in economic and game theoretical models, we would like to add couple of remarks: In our opinion, the appropriate way of modeling incomplete information situation is at the *interim stage* that is, when a player *knows his own beliefs (type)*. The Harsanyi, *ex ante* model is just an auxiliary construction for the analysis. Actually this was also the view of Harsanyi

who justified his model by proving that it provides the same equilibria as the interim stage situation it generates (Theorem 17). The “Harsanyi doctrine” says roughly that our models “*should be consistent*” and if we get an inconsistent model it must be the case that it not a ‘correct’ model of the situation at hand. This becomes less convincing if we agree that the interim stage is what we are interested in: Not only that “most” mutual beliefs are inconsistent as we saw in section 8, but it is hard to argue convincingly that the model in Example 5 on page 14 describes an ‘adequate’ mutual belief situation while the model in Example 6 on page 15 is not; the only difference between the two being that in one model, a certain type beliefs are $[\frac{3}{5}\omega_{11}, \frac{2}{5}\omega_{21}]$ while in the other model his beliefs are $[\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}]$.

Another related point is the fact that if players beliefs are the data of the situation (in the interim stage) then these are typically imprecise and rather vague to measure. Therefore any meaningful result of our analysis should be robust to small changes in the beliefs. This cannot be achieved within the consistent belief systems which are a thin set of measure zero in the Universal Belief Space.

10.2 Knowledge and Beliefs

Our interest in this article was mostly in the notion *beliefs* of players and less in the notion of *knowledge*. These are two related but different notions. Knowledge is defined through a knowledge operator satisfying some axioms (see article 22 in this collection). Beliefs are defined by means of probability distributions. Aumann’s model, discussed in section 4, has both elements: The knowledge was generated by the partitions of the players while the beliefs were generated by the probability P on the space Y (and the partitions). Being interested in the subjective beliefs of the player we could use “at state of the world $\omega \in \Omega$ player i knows the event $E \subseteq \Omega$ ” to mean “at state of the world $\omega \in \Omega$ player i assigns to the event $E \subseteq \Omega$ probability 1”. However, in the universal belief space, “belief with probability 1” does not satisfy a central axiom of the knowledge operator namely: If $\omega \in \Omega$ player i knows the event $E \subseteq \Omega$ then $\omega \in E$. That is, if a player ‘knows’ an event then this event in fact happened. In the universal belief space where all coherent beliefs are possible, in a state $\omega \in \Omega$ a player may assign probability 1 to the event $\{\omega'\}$ where $\omega' \neq \omega$. In fact, if in a BL -subspace Y the condition $\omega \in P_i(\omega)$ is satisfied for all i and all $\omega \in Y$, then belief with probability 1 is a knowledge operator on Y . This in fact was the case in Aumann’s and in Harsanyi’s models where by construction, the support of the beliefs of a player in the state ω , always included ω . For a detailed discussion of the relationship between knowledge and beliefs in the universal belief space see Vassilakis and Zamir [24].

10.3 Future directions

We have not said much about the existence of Bayesian equilibrium, mainly because it was not studied enough and there are no general results, especially in the non consistent case. We can readily see that a Bayesian game on a finite BL -subspace in which each state of nature $s(\omega)$ is a finite game-form, has a Bayesian equilibrium in mixed strategies. This can be proved for example by transforming the Bayesian game into an ordinary finite game (see Remark 16) and apply Nash Theorem for finite games. For games with incomplete information with continuum of strategies and payoff functions not necessarily continuous, there are no general existence results. Even in consistent auction models, existence was

proved for specific models separately (see Milgrom and Weber [20], Maskin and Riley [15], Reny and Zamir [22]). Establishing general existence results for large families of Bayesian games is clearly an important future direction of research. Since, as we argued before, most games are Bayesian games, the existence of Bayesian equilibrium should, and could, reach at least the level of generality available for existence of Nash equilibrium.

References

- [1] Aumann, R. (1976), “Agreeing to disagree”, *Annals of Statistics* 4:1236-1239.
- [2] Aumann, R. (1974), “Subjectivity and Correlation in Randomized Strategies”, *Journal of Mathematical Economics* 1:67–96.
- [3] Aumann, R. (1987), “Correlated Equilibrium as an Expression of Bayesian Rationality” *Econometrica* 55:1–18.
- [4] Aumann, R. (1998), “Common priors: A reply to Gul” *Econometrica* 66:929–938.
- [5] Aumann, R. (1999a), “Interactive epistemology I: Knowledge”, *International Journal of Game Theory* 28:263–300.
- [6] Aumann, R. (1999b), “Interactive epistemology I: Probability”, *International Journal of Game Theory* 28:301–314.
- [7] Aumann, R. and A. Heifetz (2002), “Incomplete Information”, in: R. Aumann and S. Hart eds., *Handbook of Game Theory* (Elsevier) Vol. 3, Chapter 43, 1666–1686.
- [8] Aumann, R. and M. Maschler (1995), “Repeated Games with Incomplete Information”, MIT Press, Cambridge, Massachusetts.
- [9] Brandenburger, A. and E. Dekel (1993), “Hierarchies of beliefs and common knowledge”, *Journal of Economic Theory* 59:189–198.
- [10] Gul, F. (1998) “A comment on Aumann’s Bayesian view”, *Econometrica* 66:923–927.
- [11] Harsanyi, J.(1967-8), “Games with incomplete information played by ‘Bayesian’ players”, parts I-III, *Management Science* 8:159-182, 320-334, 486–502.
- [12] Heifetz, A. (1993), “The Bayesian formulation of incomplete information, the non-compact case”, *International Journal of Game Theory* 21:329–338.
- [13] Heifetz, A. and P. Mongin (2001), “Probability Logic for Type Spaces”, *Games and Economic Behavior* 35:31–53.
- [14] Heifetz, A. and D. Samet (1998), “Topology-free topology of beliefs”, *Journal of Economic Theory* 82:324–341.
- [15] Maskin, E. and J. Riley (2000), “Asymmetric Auctions”, *Review of Economic Studies* 67:413–438.
- [16] Meier, M. (2001), “An infinitary probability logic for type spaces”, CORE Discussion Paper 2001/61.
- [17] Mertens, J-F., S. Sorin and S. Zamir (1994), “Repeated Games, Part A: Background Material”, CORE Discussion Paper No. 9420.
- [18] Mertens, J-F. and S. Zamir (1985), “Foundation of Bayesian analysis for games with incomplete information”, *International Journal of Game Theory* 14:1–29.
- [19] Milgrom, P. R. and N. Stokey (1982), “Information, Trade and Common Knowledge”, *Journal of Economic Theory* 26:17-27.

- [20] Milgrom, P. R. and R. J. Weber (1982), “A Theory of Auctions and Competitive Bidding”, *Econometrica* 50:1089–1122.
- [21] Nyarko, Y. (1991), “Most games violate the Harsanyi Doctrine”, C.V. Starr working paper #91-39, NYU.
- [22] Reny, P. and S. Zamir (2004), “On the existence of pure strategy monotone equilibria in asymmetric first price auctions”, *Econometrica* 72:1105–1125.
- [23] Sorin, S. and S. Zamir (1985), “A 2-person game with lack of information on $1\frac{1}{2}$ sides”, *Mathematics of Operations Research*, 10:17–23.
- [24] Vassilakis, S. and S. Zamir (1993), “Common beliefs and common knowledge”, *Journal of Mathematical Economics* 22:495–505.
- [25] Wolfstetter, E. (1999), “Topics in Microeconomics”, Cambridge University Press, Cambridge, UK.