

©Martin Osborne and Ariel Rubinstein. These **lecture notes** are distributed for the exclusive use of students in Econ 501, Princeton University Fall 2001. They are basically inserts from our book *A Course of Game Theory*, MIT Press 1994 combined with some additional comments.

Lecture 13: Bayesian Games
Readings: Osborne and Rubinstein Ch2.5

Bayesian Games: Strategic Games with Imperfect Information

A Bayesian game is a model which is designed to analyze situations in which there is asymmetry in the information held by the players regarding elements which are relevant to the play of the game. Those elements may not be “payoff relevant” necessarily. They could reflect asymmetry in knowledge or be relevant in equilibrium due to the fact that because of unspecified reasons they became relevant to the way that players play the game.

We have in mind situations which contain more information than the strategic game. The model of *Bayesian Game* contains the fundamentals from which the strategic considerations involved in such a situation are derived:

- a finite set N (the set of players)
- a finite set Ω (the set of states)
and $\forall i \in N$
- a set A_i (the set of actions available to player i)
- a finite set T_i (the set of signals that may be observed by player i) and a function $\tau_i: \Omega \rightarrow T_i$ (the signal function of player i) with a full range (by that we exclude the disturbing case that some signals never realize).
- a probability measure p_i on Ω (the prior belief of player i) for which $p_i(\tau_i^{-1}(t_i)) > 0$ for all $t_i \in T_i$
- a preference relation \succsim_i on the set of probability measures over $A \times \Omega$ (i 's preferences).

The names of the signals will not play any role in the forthcoming analysis and we can identify each $t_i \in T_i$ with the cell of the partition of Ω where the signal t_i is received. If $\tau_i(\omega) = \{\omega\}$ for all $\omega \in \Omega$ then player i has full information about the state of nature.

Often a Bayesian game is described not in terms of an underlying state space Ω , but as a “reduced form” in which the basic primitive that relates to the players’ information is the profile of the sets of possible types. Then $\Omega = \times_{i \in N} T_i$ and $\tau_i(t) = t_i$ for $\forall i$.

The interpretation we follow is that once player i receives the signal $t_i \in T_i$ then he deduces that the state is in the set $\tau_i^{-1}(t_i) = \{\omega \in \Omega | \tau_i(\omega) = t_i\}$. His *posterior belief* about the state that has been realized, is reduced from the Bayesian formula, that is, he assigns to each state $\omega \in \tau_i^{-1}(t_i)$ the probability of ω conditional on $\tau_i^{-1}(t_i)$ that is $cond_i(\omega|t_i) = p_i(\omega)/p_i(\tau_i^{-1}(t_i))$.

Frequently the model is used in situations in which a state of nature is a profile of parameters of the players’ preferences (for example, profiles of their valuations of an object). However, the model is much more general and we will see it might capture situations in which each player is uncertain about what the others *know*.

The idea that a situation in which the players are unsure about each other's characteristics can be modeled as a Bayesian game, in which the players' uncertainty is captured by a probability measure over some set of "states", is due to Harsanyi (1967/68). Harsanyi assumes that the prior belief of every player is the same, arguing that all differences in the players' knowledge should be derived from an objective mechanism that assigns information to each player, not from differences in the players' initial beliefs.

The assumption of a common prior belief has strong implications for the relationship between the players' posterior beliefs. (For example (see exercise below) after a pair of players receive their signals it cannot be "common knowledge" between them that player 1 believes the probability that the state of nature is in some given set to be α and that player 2 believes this probability to be $\beta \neq \alpha$, though it is possible that player 1 believes the probability to be α , player 2 believes it to be β , and one of them is unsure about the other's belief.)

We now turn to a definition of equilibrium for a Bayesian game. In any given play of a game each player knows his type and does not need to plan what to do in the hypothetical event that he is of some other type. Consequently, one might think that an equilibrium should be defined for each state of nature in isolation. However, in any given state a player who wishes to determine his best action may need to hold a belief about what the other players would do in other states, since he may be imperfectly informed about the state. Further, the formation of such a belief may depend on the action that the player himself would choose in other states, since the other players may also be imperfectly informed.

This leads us to the following definition:

Definition: A Nash equilibrium of a Bayesian game $\langle N, \Omega, (A_i), (T_i), (\tau_i), (p_i), (\succsim_i) \rangle$ is a Nash equilibrium of the strategic game G_1 defined as follows.

- The set of players T is the set of all pairs (i, t_i) for $i \in N$ and $t_i \in T_i$.
- $A_{(i, t_i)} = A_i$.
- The preference ordering $\succsim_{(i, t_i)}$ of each player (i, t_i) is defined by $a^* \succ_{(i, t_i)} b^*$ iff $L_i(a^*, t_i) \succ_i L_i(b^*, t_i)$, where $L_i(a^*, t_i)$ is the lottery over $A \times \Omega$ that assigns probability $cond_i(\omega|t_i)$ to $((a^*(j, \tau_j(\omega)))_{j \in N}, \omega)$ if $\omega \in \tau_i^{-1}(t_i)$, zero otherwise.

Note that to determine whether an action profile is a Nash equilibrium of a Bayesian game we need to know only how each player in the Bayesian game compares lotteries over $A \times \Omega$ in which the distribution over Ω is the same: a player never needs to compare lotteries in which this distribution is different. Thus from the point of view of Nash equilibrium the specification of the players' preferences in a Bayesian game contains more information than is necessary. (This redundancy has an analog in a strategic game: to define a Nash equilibrium of a strategic game we need to know only how any player i compares any outcome (a_{-i}, a_i) with any other outcome (a_{-i}, b_i) .)

An alternative approach is equivalent for the case we employ Bayesian updating and expected utility is one where we analyze the Bayesian game as the following strategic game G_2 :

- The set of players is once again N .
- The set of action for player i is A_i^T that is a plan what action to take contingent of the signal.

- The preference ordering \succeq_i over the space of all profile of strategies $s = (s_i)_{i \in N}$ is induced from the utility function $v_i(s) = \text{Ex}_\omega u_i(\omega, (s_i(t_i(\omega)))_{i \in N})$.

Note that under this definition a strategy is a contingent plan what action to take as a function of the signal received. It would fit a scenario that each player has to choose his plan ahead of him receiving the information about his type. Otherwise, this notion of strategy conflicts with the interpretation of a strategy as a course of action. A player will never have to choose a plan contingent on the signal as always he will confront a certain action, and thus the strategy has to be thought as the beliefs of the players about what the player would do in all contingencies (though in any instance he is aware of one signal).

Proposition: Let $\langle N, \Omega, (A_i), (T_i), (\tau_i), (p_i), (\succeq_i) \rangle$ be a Bayesian game. The profile of strategies $(s_i)_{i \in N}$ is a Nash equilibrium of G_2 iff the profile of strategies $(s_{(i,t_i)} = s_i(t_i))$ is a Nash equilibrium for G_1 .

In the following, we will see three examples of Bayesian games. The games will differ in the type of asymmetric information.

Example : A Dispute

The scenario: Players 1 and 2 are involved in a dispute. The outcome of the dispute if reaches confrontation depends on the relative power of the two. Player 1 does not know the relative power and Player 2 knows who is the stronger. Note that in any particular realization of the world, player 2 is in a certain state of knowledge, but the analysis of player 1's behavior must rely on player 1's considerations about player 2's behavior in the two states.

The scenario is modeled by the following Bayesian game:

- $N = \{1, 2\}$
- $\Omega = \{S, W\}$ (a state is interpreted as the relative strength of player 2)
- $A_i = \{Fight, Yield\}$
- $\tau_1(\omega) = \Omega$ and $\tau_2(\omega) = \omega$
- common beliefs: $p(S) = \alpha$ and $p(W) = 1 - \alpha$.
- The payoffs are given by the following tables

$u_i(a, S)$	F	Y
F	-1, 1	1, 0
Y	0, 1	0, 0

$u_i(a, W)$	F	Y
F	1, -1	1, 0
Y	0, 1	0, 0

The corresponding G_1 game is : $T = \{1, (2, S), (2, W)\}$, $A_{(i,t_i)} = \{F, Y\}$ for all players and $v_1(a) = \alpha u_1(a_1, a_S) + (1 - \alpha) u_1(a_1, a_W)$ and $v_{(2,X)}(a) = u_2(a_1, a_X, X)$ for $X = W, S$.

(the payoff of one type of player 2 does not depend on actions of the other type).

It is easy to verify that for $\alpha < 1/2$ the only Nash equilibrium is (F, F, Y) and for $\alpha > 1/2$ the only NE is (Y, F, F) .

The alternative game G_2 is:

	FF	FY	YF	YY
F	$1 - 2\alpha, 2\alpha - 1$	$1 - 2\alpha, \alpha$	$2\alpha - 1, \alpha - 1$	$1, 0$
Y	$0, 1$	$0, \alpha$	$0, 1 - \alpha$	$0, 0$

Where XY is “do X if S and do Y if W ”

Example 2: Three experts

The proper treatment of a patient is 0 or 1 with equal probabilities. The patient is tested by three experts who examine him independently. The majority of recommendations will determine the treatment. Each expert comes to a conclusion, 0 or 1. The conclusion is stochastic: it is the “correct” one with probability $q > 1/2$. All experts are “nice people” and are only interested to increase the probability that the patient will get the correct treatment (there is an underlying symmetry in the loss function).

We can model this situation as the Bayesian game in which

- $N = \{1, 2, 3\}$
 - $\Omega = \{0, 1\} \times \{-, +\}^N$ (the state $(0, -, +, +)$ means that the agent’s right treatments is 0 and expert 1 got it wrong)
 - $A_i = \{0, 1\}$
 - $\tau_i(x, t_1, t_2, t_3) = \begin{cases} x & \text{if } t_i = + \\ 1 - x & \text{if } t_i = - \end{cases}$
 - common beliefs: $p(x, t_1, t_2, t_3) = q^{\#}(1 - q)^{\#}/2$.
 - Payoffs: all players share the same payoffs.
- $$u_i(x, a_1, a_2, a_3) = \begin{cases} 1 & \text{if } \#\{i|a_i = x\} > 1 \\ 0 & \text{if otherwise} \end{cases}$$

The BNE are the NE in the G_2 game where each of the players has to choose one of four contingent strategies:

- “ T ”= always report the signal,
- “ NT ”= always mis-report the signal,
- “ C_n ”= always report n

Claim 1 (C_0, C_0, C_0) is a BNE (no player can influence the consequence once the two other declare 0).

Claim 2 (T, T, T) is a BNE

Claim 3 (C_1, C_0, T) is a BNE Player 3 plays a best response: using T the probability of right treatment is q . If he deviates to NT the probability is reduced to $1 - q$ and if he uses A_n it is reduced to $1/2$.

To see that player 1’s strategy is best response, note that if he deviates to C_0 or NT he reduces the probability of a correct treatment to $1/2$. If he uses T , he affects the outcome only in the following two events:

- $(1, -, ., +)$ for bad
- $(0, +, ., -)$ for good.

The two events are equally likely the deviation to T is not strictly beneficial.

Claim 4 (exercise) Show that the game does not have any equilibrium with NT. Is NT a dominated strategy (namely, is there a strategy such that whatever the other players do, NT is inferior to that strategy)?

Example 3: Electronic game (A variant of Rubinstein (AER 1989))

A Bayesian game can be used to model not only situations in which each player is uncertain about the other players' payoffs, but also situations in which each player is uncertain about the other players' *knowledge*.

Two players can choose one of two actions *attack* or *not*. An attack will be successful only if both attack and the state of nature will be right (happens with probability 1/2). A failed attack cause a loss of $L > M$. The following is the payoff of a player choosing the row conditional that the other chooses the column:

u_i at G	A	N
A	M	$-L$
N	0	0

u_i at B	A	N
A	$-L$	$-L$
N	0	0

Only player 1 knows whether the conditions to attack are good or not. To provoke coordination a protocol of communication is set: if the machine of player 1 gets the signal that the state of the world is G (and only then !) the machine sends *automatically* a message (blip) to the machine of player 2. However, each message may be interrupted with (small) probability $\mu > 0$. Therefore a protocol has been devised so that machine 2 sends a confirmation to machine 1's message. But, in order for player 2 to be sure that player 1 knows that the message has arrived machine 1 sends a confirmation to the confirmation and so on.... To summarize, the protocol of communication is such that each machine automatically confirms any message it gets until a message does not go through. The machines work fast and when they stop, (an event which happens with probability 1) each player observes the number of messages that his machine sent and only takes an action (attack or not). Here is the Bayesian game which seems to model such a scenario:

- $N = \{1, 2\}$
- $A_i = \{attack, not\}$
- $\Omega = \{(t_1, t_2) | t_1 = t_2 \text{ or } t_1 = t_2 + 1\}$ (t_i is the number of messages sent by the machine of player i).
- $\tau_i(t_1, t_2) = t_i$ (When 1 notes the number 3 on his machine he does not know whether player 2 got 2 messages and his last message did not go through or player 2 got 3 messages and his last confirmation was lost). In this game an infinite number of types $\{1(q), 2(q) | q = 0, 1, 2, \dots\}$.
- $p(0, 0) = 1/2, p(t_1, t_2) = \mu(1 - \mu)^{t_1+t_2-1}/2$.

Note, here the only differences between types is in the knowledge!

Proposition: The only BNE of this Bayesian game is the one where all types do not attack.

Thus, even if μ is very small, the above protocol of communication does not enable coordination!.)

Proof: Consider any BNE of the game. It is dominating for type 1(0) to play n . As to

type 2(0), in equilibrium, if he attacks, with probability $\mu/2/(\mu/2 + 1/2) = \mu/(1 + \mu)$ he get M and with probability $1/(1 + \mu)$ he gets at most $-L$. Thus, it is optimal for him not to attack. Now, order the types, 1(0),2(0),1(1),2(1).... assume that all types up to $i(t)$ play n . Consider the next type $j(t)$ or $j(t + 1)$. If he does not attack he gets 0. If he attacks he gets $-L$ with probability $\mu/((1 - \mu)\mu + \mu) = 1/(2 - \mu)$ and with the smaller probability $(1 - \mu)/(2 - \mu)$ he gets at most M . Thus it must be that he does not attack as well.

Problem set 13

1. **(Exercise)** Consider the case that two players have common prior on the space of states of nature, where each state represents the objective value of a commodity.

Show that it impossible that player 1 thinks after any signal he gets that the expected value is strictly above p^* whereas player 2 thinks after any signal he gets that the expected value is strictly below p^* .

Explain why the common prior assumption is crucial for this result.

2. **(Exercise)** Construct and prove the one player analogous result of the proposition about the equivalence between the equilibria of the games G_1 and G_2 described above.

3. **(Exercise)** Show that more information may hurt a player by constructing a two-player Bayesian game with the following features. Player 1 is fully informed while player 2 is not; the game has a unique Nash equilibrium, in which player 2's payoff is higher than his payoff in the unique equilibrium of any of the related games in which he knows player 1's type.

4. **(Exercise)** Each of two players receives a ticket on which there is a number in some finite subset S of the interval $[0, 1]$. The number on a player's ticket is the size of a prize that he may receive. The two prizes are identically and independently distributed, with distribution function F . Each player is asked independently and simultaneously whether he wants to exchange his prize for the other player's prize. If both players agree then the prizes are exchanged; otherwise each player receives his own prize. Each player's objective is to maximize his expected payoff. Model this situation as a Bayesian game and show that in any Nash equilibrium the highest prize that either player is willing to exchange is the smallest possible prize.

(See <http://www.u.arizona.edu/~chalmers/papers/envelope.html> for some philosophical discussion of the related two-envelope paradox).

5. **Exercise (Second-price auction)** Consider a variant of the second-price sealed-bid auction described in the previous lecture in which each player i knows his own valuation v_i but is uncertain of the other players' valuations. Suppose that the set of possible valuations is the finite set V and each player believes that every other player's valuation is drawn independently from the same distribution over V . Players maximize their expected "surplus".

- ▲ Model this situation as the Bayesian game .
- ▲ Show that the game has a BNE where each player bids his valuation. Is it the unique BNE?
- ▲ Show that bidding is a weakly dominant action for each type of each player to bid his valuation.