Bayesian Games

Moty Katzman

April 27, 2015

A Bayesian game is, roughly speaking, a game in which players do not have full information about the game they are playing.

Example. PC Bob faces an armed suspect called Alice and they have to decide simultaneously whether to shoot or not. The suspect is either a criminal (with probability p) or not (with probability $1 - p$). The sheriff would rather shoot if the suspect shoots, but not otherwise. A criminal would rather shoot even if the sheriff does not. An innocent suspect would rather not shoot even if the sheriff shoots.

The showdown– cont'd

Bob doesn't know which of the following two games is being played

Game against innocent person

	shoot	dont shoot
shoot	-3, -1	-1, -2
dont shoot	-2, -1	0, 0

Game against criminal

	shoot	dont shoot
shoot	0, 0	2, -2
dont shoot	-2, -1	-1, 1

he knows that he is playing the first with probability $1 - p$ and the second with probability p.

The showdown– cont'd

What should Bob do? He should put himself in Alice's shoes: there is no uncertainty for her. In the first game her *dont shoot* strategy dominates her *shoot* strategy and in the second game her *shoot* strategy dominates her *dont shoot* strategy. So Bob is really playing the following games

Game against innocent person

	shoot	dont shoot
dont shoot	-2, -1	0, 0

Game against criminal

	shoot	dont shoot
shoot	0, 0	2, -2

Now if Bob shoots, his expected payoff is $-(1-p)$ and $-2p$ if he doesnt. So he shoots when $-(1-p) \geq -2p$, i.e., when $p \geq 1/3$.

Definition. A (two-player) Bayesian game consists of:

(a) a set of actions available to each player (S and T),

(b) a finite set of types for each player (Θ_A and Θ_B),

(c) a set of possible states of the world Ω, and functions $\tau_A : \Omega \to \Theta_A$ and $\tau_B : \Omega \to \Theta_B$.

(d) a payoff function for each player whose domain is $S \times T \times \Omega$,

(e) all probabilities $\mathrm{Prob}(\omega|\theta)$ for all $\theta \in \Theta_A \cup \Theta_B$ and $\omega \in \Omega$.

PC Bob revisited:

(a) Both Alice and Bob have set of actions $S = T = \{\text{shoot}, \text{don't shoot}\}$.

(b) $\Theta_B = \{\text{PC}\}$ and $\Theta_A = \{\text{innocent}, \text{criminal}\}$.

(c) $\Omega = \{\omega_1, \omega_2\}$, $\tau_B(\omega_1) = \tau_B(\omega_2) = \text{PC}$, $\tau_A(\omega_1) = \text{innocent}$, $\tau_A(\omega_2) = \text{criminal}$.

(d) $u_A(\text{shoot}, \text{shoot}, \omega_1) = -3$,
$u_B(\text{shoot}, \text{shoot}, \omega_1) = -1$, etc.

(e) $\text{Prob}(\omega_1|\text{PC}) = 1 - p$, $\text{Prob}(\omega_2|\text{PC}) = p$,
$\text{Prob}(\omega_1|\text{innocent}) = 1$, $\text{Prob}(\omega_2|\text{innocent}) = 0$,
$\text{Prob}(\omega_1|\text{criminal}) = 0$, $\text{Prob}(\omega_2|\text{criminal}) = 1$.

Strategies in Bayesian Games

Definition. A *strategy* for a player in a Bayesian game is a function from the set of its types to the set of its actions. Using our notation for Alice and Bob, a strategy for Alice is an element of S^{Θ_A} and a strategy for Bob is an element in T^{Θ_B}. A strategy profile is a choice of a strategy for each player, i.e., an element in $S^{\Theta_A} \times T^{\Theta_B}$.

Given a strategy profile $(s(-), t(-)) \in S^{\Theta_A} \times T^{\Theta_B}$, Alice's expected payoff when she has type $\theta_A \in \Theta_A$ is

$$\sum_{\omega \in \Omega} \mathrm{Prob}(\omega | \theta_A) u_A(s(\theta_A), t(\tau_B(\omega)), \omega)$$

and Bob's expected payoff when he has type $\theta_B \in \Theta_B$ is

$$\sum_{\omega \in \Omega} \mathrm{Prob}(\omega | \theta_B) u_B(s(\tau_A(\omega)), t(\theta_B), \omega).$$

Bayes-Nash equilibrium

Definition. The strategy profile $(s(-), t(-)) \in S^{\Theta_A} \times T^{\Theta_B}$ is a *Bayes-Nash equilibrium* if $s = s(\theta_A)$ maximizes

$$\sum_{\omega \in \Omega} \text{Prob}(\omega|\theta_A) u_A(s, t(\tau_B(\omega)), \omega)$$

for all $\theta_A \in \Theta_A$ and $t = t(\theta_B)$ maximizes

$$\sum_{\omega \in \Omega} \text{Prob}(\omega|\theta_B) u_B(s(\tau_A(\omega)), t, \omega)$$

for all $\theta_B \in \Theta_B$.

This definition says that a strategy profile is a Bayes-Nash equilibrium if *each type separately* would not change its action as specified in the corresponding strategy.

So for practical uses we can look for Bayes-Nash equilibria by replacing a given Bayesian game with a new game whose set of players are all the types in the original game.

Example

Example. Alice has a dispute with Bob, who is either strong (S) or weak (W). Alice believes that Bob is strong with probability p. Each person can either fight or yield (henceforth abbreviated F and Y). The outcome of the confrontation is given as follows:

strong Bob

	F	Y
F	-1,1	1,0
Y	0,1	0,0

weak Bob

	F	Y
F	1,-1	1,0
Y	0,1	0,0

Here the sets of actions for both players is $\{F, Y\}$, $\Theta_A = \{A\}$ and $\Theta_B = \{S, W\}$ $\Omega = \{\omega_1, \omega_2\}$, $\tau_A(\omega_1) = \tau_A(\omega_2) = A$, $\tau_B(\omega_1) = S$, $\tau_B(\omega_2) = W$. Alice has two strategies: $(A \to F)$ and $(A \to Y)$. Bob has four strategies $(S \to F, W \to F)$, $(S \to F, W \to Y)$, $(S \to Y, W \to F)$, and $(S \to Y, W \to Y)$.We compute expected payoffs of all strategy profiles:

	$(S \to F, W \to F)$	$(S \to F, W \to Y)$
$(A \to F)$	1-2p, (1,-1)	1-2p,(1,0)
$(A \to Y)$	0, (1,1)	0, (1,0)

	$(S \to Y, W \to F)$	$(S \to Y, W \to Y)$
$(A \to F)$	1, (0,-1)	1, (0,0)
$(A \to Y)$	0, (0,1)	0, (0,0)

If $p < 1/2$, $1 - 2p > 0$ and the first row dominates the second, and we obtain a Bayes-Nash equilibrium at $((A \to F), (S \to F, W \to Y))$. If $p > 1/2$, $1 - 2p < 0$ and we obtain a Bayes-Nash equilibrium at $((A \to Y), (S \to F, W \to F))$.

Example: More information may hurt

Alice and Bob play one of the games below:

Game G_1

	L	M	R
U	1, 2x	1,0	1, 3x
D	2,2	0,0	0,3

Game G_2

	L	M	R
U	1, 2x	1,3x	1, 0
D	2,2	0,3	0,0

where $0 \leq x < 1/2$. Neither know which game is played, both assign probability $1/2$ to either game. Here each player has one type, $\Omega = \{G_1, G_2\}$.

Strategies can be identified with actions and the expected payoffs of this Bayesian game are given by

	L	M	R
U	1, 2x	1,3x/2	1, 3x/2
D	2,2	0,3/2	0,3/2

Column L dominates the others and we obtain a Bayes-Nash equilibrium (D,L) which results in a payoff of 2 for Bob.

Game G_1	L	M	R
U	1, 2x	1,0	1, 3x
D	2,2	0,0	0,3

Game G_2	L	M	R
U	1, 2x	1,3x	1, 0
D	2,2	0,3	0,0

Suppose now that Bob knows which game he is playing: now Bob has two types, say B_1 and B_2, $\tau_B(G_1) = B_1$ and $\tau_B(G_2) = B_2$. Also, now Bob has 9 strategies, $(L, L), (L, M), \ldots$.
For B_1, strategy R is a dominant strategy, and for B_2, M is a dominant strategy, so we can eliminate strategy L. Now Alice strategy U dominates, and we end up in a Bayes-Nash equilibrium $(U, (R, M))$ which give Bob a payoff of $3x < 2$!