Bayesian Games

How do we model uncertainty about the payoffs or (more generally) knowledge of the other players?

The traditional distinction (see Fudenberg and Tirole) is that uncertainty about payoffs is called **incomplete information**, and uncertainty about previous actions is called **imperfect information**.

Harsanyi invented the trick of converting games of incomplete information to games of imperfect information, by introducing **nature** as an additional player. For example, suppose we are uncertain about whether another player has low costs or high costs. This can be modelled as knowing that costs are low when nature chooses $\omega = L$, and that costs are high when nature chooses $\omega = H$.

Definition 25.1: A **Bayesian game** consists of

1. A finite set N (players),

2. A finite set Ω (states of nature),

and for each player $i \in N$,

3. A set A_i (actions),

4. A finite set T_i (**signals** that may be observed by player i) and a function $\tau_i : \Omega \to T_i$ (the **signal function** of player i),

5. A probability measure, p_i, on Ω (the **prior belief** of player i) for which $p_i(\tau_i^{-1}(t_i)) > 0$ holds for all $t_i \in T_i$.

6. A preference relation $\succsim_i$ on the set of probability measures over $A \times \Omega$, where $A = \times_{j \in N} A_j$.

Notes:

In part (4), it is assumed that signals "partition" the set of states, so there is no noise in the process generating signals from states. This is without loss of generality, because otherwise we can simply expand the set of states to include the signal.

In part (5), we assume that all signals occur with positive probability. This is without loss of generality, because otherwise we can exclude that signal from T_i. Also, priors can differ across players.

In part (6), preferences can depend on the state as well as the actions. Also, since the state is uncertain, preferences must be over lotteries of action-state combinations.

The signal, t_i, is usually called player i's **type**. The interpretation of Definition 25.1 is that a player can observe his/her type before choosing an action. Since other players might be unsure about player i's type, their optimal actions will depend on the entire mapping from t_i to a_i. Thus, when player i is considering his optimal action when he is type t_i, he must consider what he would have done had his type realization been different.

One can think of (i, t_i') and (i, t_i'') as representing two different players.

Some More Notation:

Denote the Bayesian game as G^*, and an action profile as a^*.

The set of action profiles is $\times_{j \in N}(\times_{t_j \in T_j} A_j)$.

The poterior beliefs of player i when his type is t_i, and an action profile a^*, generates a lottery over $A \times \Omega$, denoted by $L_i(a^*, t_i)$. Given (a^*, t_i), the probability of $\left[\left(a^*(j, \tau_j(\omega))_{j \in N}\right), \omega\right]$ is player i's posterior belief that the state is ω, conditional on the signal t_i. This probability is

$$\frac{p_i(\omega)}{p_i(\tau_i^{-1}(t_i))} \text{ if } \omega \in \tau_i^{-1}(t_i) \text{ and zero otherwise.}$$

Definition 26.1: A Nash equilibrium of a Bayesian game $G^* = \langle N, \Omega, (A_i), (T_i), (\tau_i), (p_i), (\succsim_i) \rangle$ is a Nash equilibrium of the strategic game G defined as follows.

1. The set of players is the set of all pairs, (i, t_i) for $i \in N$ and $t_i \in T_i$.

2. The set of actions for player (i, t_i) is A_i.

3. The preference ordering for player (i, t_i) is denoted by $\succsim^*_{(i,t_i)}$ and defined by

$$a^* \succsim^*_{(i,t_i)} b^* \text{ if and only if } L_i(a^*, t_i) \succsim_i L_i(b^*, t_i).$$

Example 27.1: Second-price auction with independent private values, drawn from a finite set, V, with probability distribution π.

$N = \{1, 2, ..., n\}$.

$\Omega = V^n$ (the set of valuation profiles).

$A_i = \mathbb{R}_+$ (any nonnegative bid is allowed)

$T_i = V$ (a player's signal is his/her valuation)

For any $\omega = (v_1, ..., v_n)$, the signal function is defined by $\tau_i(v_1, ..., v_n) = v_i$

For all i, priors are $p_i(v_1, ..., v_n) = \prod_{j=1}^{n} \pi(v_j)$

Given an action profile a^* and a signal $\overline{v}_i$, preferences for consumer i are represented by a utility function, $U_i(a^*, \overline{v}_i) =$

$$\sum_{\omega : v_i = \overline{v}_i} \left[\prod_{j \neq i} \pi(v_j) \right] \left(\overline{v}_i - \max_{j \in N \setminus i} a^*(j, \tau_j(\omega)) \right) \mathbf{1}_i(\omega)$$

where

$\mathbf{1}_i(\omega) = 1$ if player i is the lowest index for whom $a_i \geq a_j$ for all $j \in N$ (i wins)

$\mathbf{1}_i(\omega) = 0$ otherwise.

This game has a Nash equilibrium action profile given by $a^*(i, v_i) = v_i$. It is impossible to do better than bidding one's valuation, no matter what actions other players are choosing, so this clearly is a NE.

Discussion

If preferences can be represented by a von Neumann-Morgenstern utility function, then we can equivalently think of player i as choosing a plan of action (a mapping from $T_i \to A_i$) before learning his type. (Call the game G'.) Then the set of actions in G' is the set of action functions, based on G^*. Preferences in G' are given by ex ante expected utility, rather than conditional on the realized type. For any Nash equilibrium to G', $(a'_i(\cdot))_{i \in N}$, there is a Nash equilibrium of G^* such that $a'_i(t_i) = a^*(i, t_i)$ for all (i, t_i).

This interpretation of the game as G', with A'_i being the space of action functions, allows the easiest extension to model a continuum of types.

The model applies to games with uncertainty over other players' knowledge as well as uncertainty over other players' valuations. For example, we could have an auction where, in some state ω, everyone's signal allows them to know all of the valuations; however, players cannot infer the other players' signals, so no one knows whether other people know all of the valuations or just their own valuation.

Harsanyi advocates the assumption of common priors, $p_i = p$ for all $i \in N$. This is based on the idea that any apparent differences in prior beliefs should be due to unmodelled differences in information, so we should take the model back to the point in time when prior beliefs are the same for everyone.

There is a symmetry between people with zero information, from which we invoke the principle of insufficient reason.

1. This is a matter of faith, and cannot be proven.

2. Even if the Harsanyi doctrine is true at some level, the "correct" model is impossible to write down or analyze. (Do we have to consider a player's prior beliefs, before their DNA becomes part of them?) For the models we work with, we might want to allow differences in priors.

3. The principle of insufficient reason is controversial.