Game Theory 2025

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Plan for Today

So far, our players didn't know the strategies of the others, but they did know the rules of the game (i.e., how actions determine outcomes) and everyone's incentives (i.e., their utility functions).

Today we are going to change this and introduce *uncertainty:*

- Idea: *epistemic types*
- Model: *Bayesian games*
- Solution concept: *Bayes-Nash equilibrium*

This (and more) is also covered in Chapter 7 of the *Essentials*.

K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multi-disciplinary Introduction*. Morgan & Claypool Publishers, 2008. Chapter 7.

Modelling Uncertainty

We are only going to model uncertainty about utility functions.

Is this not too restrictive? No! Example:

Suppose Rowena is uncertain whether Colin has action M available. She can simply assume he does, but entertain the possibility that he assigns very low utility to any outcome involving M:

	L	R
T	1, 2	5, 8
B	11, 4	7, 6

$\sim$

	L	M	R
T	1, 2	3, -100	5, 8
B	11, 4	9, -100	7, 6

Epistemic Types

The main new concept for today is that of a player's (epistemic) *type*. This encodes all the private information for that player.

- When the game is played, you know your own type with certainty but only have probabilistic knowledge about the types of the others.
- Your own utility depends on your own type. Still might reason about a game *before* you observe your type (example: make conditional plan ahead of collecting information).
- Your own utility also depends on the types of others (example: utility of winning an auction depends on knowledgeability of rival bidders).

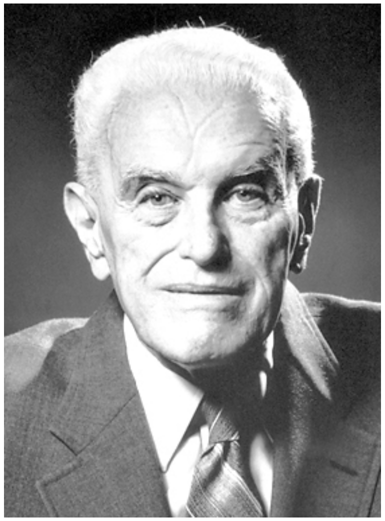

John C. Harsanyi (1920–2000)

J.C. Harsanyi. Games with Incomplete Information Played by "Bayesian" Players. Part I: The Basic Model. *Management Science*, 14(3):159–182, 1967.

Bayesian Games

A *Bayesian game* is a tuple $\langle N, \boldsymbol{A}, \boldsymbol{\Theta}, p, \boldsymbol{u} \rangle$, where

- $N = \{1, \ldots, n\}$ is a finite set of *players*;
- $\boldsymbol{A} = A_1 \times \cdots \times A_n$, with A_i the set of *actions* of player i;
- $\boldsymbol{\Theta} = \Theta_1 \times \cdots \times \Theta_n$, with Θ_i the set of possible *types* of player i;
- $p : \boldsymbol{\Theta} \to [0, 1]$ is a *common prior* (probability distribution) over $\boldsymbol{\Theta}$;
- $\boldsymbol{u} = (u_1, \ldots, u_n)$ is a profile of *utility functions* $u_i : \boldsymbol{A} \times \boldsymbol{\Theta} \to \mathbb{R}$.

We assume that also $\boldsymbol{A}$ and $\boldsymbol{\Theta}$ are *finite* (generalisations are possible).

Player i *knows* $\boldsymbol{\Theta}$ and p, and *observes* her own type $\theta_i \in \Theta_i$, but not $\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}$. She *chooses* an action a_i, giving rise to the profile $\boldsymbol{a} \in \boldsymbol{A}$. Actions are played simultaneously. Player i receives payoff $u_i(\boldsymbol{a}, \boldsymbol{\theta})$.

Remark: If $|\Theta_i| = 1$ for all $i \in N$ (if everyone's type is unambiguous), this reduces to the familiar definition of a normal-form game.

Knowledge of the State of the World

The common prior p is only defined on $\boldsymbol{\Theta}$. But we can also infer:

- Probability $p(\theta_i)$ of player i having type θ_i:

$$p(\theta_i) \quad = \sum_{\boldsymbol{\theta'} \in \boldsymbol{\Theta} \text{ s.t. } \theta'_i = \theta_i} p(\boldsymbol{\theta'})$$

- Conditional probability $p(\boldsymbol{\theta}_{-i} \mid \theta_i)$ of the other players having the types as indicated by $\boldsymbol{\theta}_{-i}$, given that player i has type θ_i:

$$p(\boldsymbol{\theta}_{-i} \mid \theta_i) \quad = \quad \frac{p(\boldsymbol{\theta})}{p(\theta_i)}$$

  This is all player i knows upon observing her own type.

Strategies

A *pure strategy* for player i now is a function $\alpha_i : \Theta_i \to A_i$ for picking the action she will play once she observes her own type.

A *mixed strategy* for i is a probability distribution $s_i \in S_i = \Pi(A_i^{\Theta_i})$ over the space of her pure strategies. Three ways to think about this:

- Mapping pure strategies to probabilities:
  $s_i : (\Theta_i \to A_i) \to [0, 1]$
- Mapping types to probability distributions over actions:
  $s_i : \Theta_i \to (A_i \to [0, 1])$
- Mapping pairs of types and actions to probabilities:
  $s_i : \Theta_i \times A_i \to [0, 1]$

Write $s_i(a_i \mid \theta_i) = \frac{s_i(\theta_i, a_i)}{p(\theta_i)}$ for the probability of player i playing the action a_i in case she has type θ_i and uses strategy s_i.

Three Notions of Expected Utility

Player i's *ex-post expected utility* is her expected utility given everyone's strategies $\boldsymbol{s}$ and types $\boldsymbol{\theta}$:

$$u_i(\boldsymbol{s}, \boldsymbol{\theta}) \quad = \quad \sum_{\boldsymbol{a} \in \boldsymbol{A}} \left[u_i(\boldsymbol{a}, \boldsymbol{\theta}) \cdot \prod_{j \in N} s_j(a_j \mid \theta_j) \right]$$

Player i's *ex-interim expected utility* is her expected utility given everyone's strategies $\boldsymbol{s}$ and her own type θ_i:

$$u_i(\boldsymbol{s}, \theta_i) \quad = \sum_{\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}} u_i(\boldsymbol{s}, (\theta_i, \boldsymbol{\theta}_{-i})) \cdot p(\boldsymbol{\theta}_{-i} \mid \theta_i)$$

Player i's *ex-ante expected utility* is her expected utility given everyone's strategies $\boldsymbol{s}$, before observing her own type:

$$u_i(\boldsymbol{s}) \quad = \quad \sum_{\theta_i \in \Theta_i} u_i(\boldsymbol{s}, \theta_i) \cdot p(\theta_i) \quad = \quad \sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} u_i(\boldsymbol{s}, \boldsymbol{\theta}) \cdot p(\boldsymbol{\theta})$$

Remark: We again use u_i both for plain utility and for expected utility.

Exercise

Verify that our two alternative definitions of ex-ante expected utility indeed coincide. In other words, prove the following for all $\boldsymbol{s}$*:*

$$\sum_{\theta_i \in \Theta_i} u_i(\boldsymbol{s}, \theta_i) \cdot p(\theta_i) \quad = \quad \sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} u_i(\boldsymbol{s}, \boldsymbol{\theta}) \cdot p(\boldsymbol{\theta})$$

Bayes-Nash Equilibria

Consider a Bayesian game $\langle N, \boldsymbol{A}, \boldsymbol{\Theta}, p, \boldsymbol{u}\rangle$ with strategies $s_i \in S_i$.

We say that strategy $s_i^\star \in S_i$ is a *best response* for player i to the (partial) strategy profile $\boldsymbol{s}_{-i}$ if $u_i(s_i^\star, \boldsymbol{s}_{-i}) \geqslant u_i(s_i', \boldsymbol{s}_{-i})$ for all $s_i' \in S_i$.

We say that profile $\boldsymbol{s} = (s_1, \ldots, s_n)$ is a *Bayes-Nash equilibrium*, if s_i is a best response to $\boldsymbol{s}_{-i}$ for every player $i \in N$.

Remark: The definitions on this slide are essentially copies of the definitions we had used to introduce mixed Nash equilibria. Only the type of game and the notion of expected utility have changed.

Note: Best responses are defined via *ex-ante* expected utility ($\hookrightarrow$).

Discussion

You need to think about what strategy to use *after* you observe your own type. So why define best responses via *ex-ante* expected utility?

Answer: Keep in mind that strategies s_i are 'conditional'. They fix a plan for how to play for any type θ_i you might end up observing.

So when you optimise to find your best response to $\boldsymbol{s}_{-i}$, you in fact are solving an *independent optimisation problem* for every possible type:

$$s_i^\star \in \operatorname*{argmax}_{s_i \in S_i} u_i(s_i, \boldsymbol{s}_{-i}) = \operatorname*{argmax}_{s_i \in S_i} \sum_{\theta_i \in \Theta_i} \underbrace{u_i((s_i, \boldsymbol{s}_{-i}), \theta_i)}_{\substack{\text{does not depend} \\ \text{on } s_i(\_, \theta_i') \\ \text{for } \theta_i' \neq \theta_i}} \cdot p(\theta_i)$$

Thus, in case we have $p(\theta_i) > 0$ for all $i \in N$ and all $\theta_i \in \Theta_i$, we can equivalently define BNE via *ex-interim* expected utility:

$$\boldsymbol{s} \text{ is a BNE } \underline{\text{iff}} \;\; s_i \in \operatorname*{argmax}_{s_i' \in S_i} u_i((s_i', \boldsymbol{s}_{-i}), \theta_i) \text{ for all } i \in N$$

Example: Fight!

You (Player 1) are considering to have a fight with Player 2, who could be of the *weak* or the *strong* type. (Your own type is clear to everyone.)

	F	$\overline{\text{F}}$
F	1, -2	2, -1
$\overline{\text{F}}$	-1, 2	0, 0

$\theta_2 = weak$

	F	$\overline{\text{F}}$
F	-2, 1	2, -1
$\overline{\text{F}}$	-1, 2	0, 0

$\theta_2 = strong$

Let p be the probability (common prior) that Player 2 is $weak$.

<u>Exercise:</u> *Analyse the game for the special cases of $p = 1$ and $p = 0$!*

Exercise: Compute the Bayes-Nash Equilibria

So we have $A_1 = A_2 = \{\mathsf{F}, \overline{\mathsf{F}}\}$, $\Theta_1 = \{\bot\}$, and $\Theta_2 = \{weak, strong\}$. Pure strategies are of the form $\alpha_i : \Theta_i \to A_i$. Here they are:

- Player 1: *do-fight*, *don't-fight*
- Player 2: *always-fight*, *fight-if-strong*, *fight-if-weak*, *never-fight*

So there might be up to $2 \times 4 = 8$ pure Bayes-Nash equilibria ...

Let $p = p(\bot, weak)$ be the probability of Player 2 being weak.

Writing θ_2 for $\boldsymbol{\theta} = (\bot, \theta_2)$, the *ex-ante expected utility* of player i for strategy profile $\boldsymbol{s}$ is $u_i(\boldsymbol{s}) = p \cdot u_i(\boldsymbol{s}, weak) + (1-p) \cdot u_i(\boldsymbol{s}, strong)$.

Recall that $\boldsymbol{s}$ is a *Bayes-Nash equilibrium* iff $s_i \in \operatorname{argmax}_{s_i' \in S_i} u_i(s_i', \boldsymbol{s}_{-i})$.

For $p = 0$ and $p = 1$ we've analysed the game already (previous slide).

Now: *For any given $p \in (0,1)$, compute all pure Bayes-Nash equilibria!*

Solution: Pure Bayes-Nash Equilibria

Two pure strategies of Player 2 are *strictly dominated* for any $p \in (0, 1)$:

- Player 1: *do-fight*, *don't-fight*
- Player 2: *always-fight*, *fight-if-strong*, ~~*fight-if-weak*~~, ~~*never-fight*~~

Intuitively, pure BNE will depend on p:

- $(\textit{do-fight}, \textit{fight-if-strong})$ for *high* p (when Player 2 is likely weak)
- $(\textit{don't-fight}, \textit{always-fight})$ for *low* p (when Player 2 is likely strong)

In fact, *fight-if-strong* is always the best response to *do-fight* and *always-fight* is always the best response to *don't-fight*.

So need to determine for which values of p the opposite holds as well:

$$\begin{aligned} u_1(\textit{do-fight}, \textit{fight-if-strong}) &\geqslant u_1(\textit{don't-fight}, \textit{fight-if-strong}) \\ 2p - 2(1-p) &\geqslant 0p - 1(1-p) \quad \Rightarrow \text{BNE for } p \geqslant \tfrac{1}{3} \end{aligned}$$

Same approach: other BNE for $p \leqslant \frac{1}{3}$ $[-1p - 1(1-p) \geqslant 1p - 2(1-p)]$

Thus: two BNE for $p = \frac{1}{3}$, and otherwise exactly one BNE.

Translation

Bayesian games are convenient for reasoning about strategic behaviour in the presence of uncertainty, but—in principle—the same reasoning could be carried out using simple normal-form games.

We can translate $\langle N, \boldsymbol{A}, \boldsymbol{\Theta}, p, \boldsymbol{u} \rangle$ to $\langle N^\star, A^\star, \boldsymbol{u}^\star \rangle$ as follows:

- $N^\star := N$ — same set of players
- $A_i^\star := \{a_i^\star \mid a_i^\star : \Theta_i \to A_i\}$ — actions are pure Bayesian strategies
- $u_i^\star : \boldsymbol{a}^\star \mapsto u_i(\boldsymbol{a}^\star)$ — utility is *ex-ante* expected utility

Remark: The Bayes-Nash equilibria of the original Bayesian game now correspond to the Nash equilibria of its translation.

Exercise: *Express the fighting game for $p = \frac{1}{2}$ as a normal-form game!*

Solution: Translation to Normal Form

Recall that actions now are of the form $a_i^\star : \Theta_i \to A_i$, with $A_i = \{\mathsf{F}, \overline{\mathsf{F}}\}$. As $|\Theta_1| = 1$ and $|\Theta_2| = 2$, we get a $2^1 \times 2^2$ normal-form game.

For $p = \frac{1}{2}$, we can compute expected utilities for all combinations:

	AlwF	FiS	FiW	NevF
F	$-\frac{1}{2}$, $-\frac{1}{2}$	0, 0	$\frac{3}{2}$, $-\frac{3}{2}$	2, -1
$\overline{\mathsf{F}}$	-1, 2	$-\frac{1}{2}$, 1	$-\frac{1}{2}$, 1	0, 0

Note that (as seen before) *fight-if-weak* and *never-fight* are dominated.

So the only pure NE is (*do-fight*, *fight-if-strong*), corresponding to the pure BNE we computed before for any $p \geqslant \frac{1}{3}$.

Existence of Bayes-Nash Equilibria

Recall that here we are only dealing with *finite* games. Thus:

Corollary 1 *Every Bayesian game has a Bayes-Nash equilibrium.*

Proof: Follows immediately from (i) our discussion of how to translate Bayesian games into normal-form games and (ii) Nash's Theorem on the existence of Nash equilibria. ✓

Summary

This has been an introduction to games of *incomplete information*, modelled in the form of *Bayesian games*. We have seen:

- definition of the model, based on *epistemic types* θ_i, with utilities based on $(\theta_1, \ldots, \theta_n)$ and a common prior on the full type space
- three notions of *expected utility: ex-post*, *ex-interim*, *ex-ante*
- *Bayes-Nash equilibrium:* a solution concept defined in terms of best responses relative to *ex-ante* expected utility
- *translation* to *complete-information* normal-form games is possible

What next? Modelling sequential actions via games in extensive form.