CS 405/6001: Game Theory and Algorithmic Mechanism Design **Jul–Nov 2026**

Lecture 12: Bayesian Games and Bayesian Equilibrium

Lecturer: Swaprava Nath *Scribe(s): Claude Scribe*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor.*

12.1 A Different Kind of Uncertainty

In Lecture 11 the players knew the game. What they did not know was *where in that game* they were standing, and beliefs over information sets fixed that. Today the uncertainty moves one level up: a player may not know *which game she is playing at all*.

It is worth placing this in the map of everything we have seen.

- **Non-cooperative games.**
  - *Complete information* – players deterministically know which game they are playing. There may be chance moves, but their probabilities are known. Normal form games, appropriate for simultaneous-move single-stage interaction, with SDSE, WDSE, PSNE, MSNE and correlated equilibrium; and extensive form games, appropriate for multi-stage interaction, with SPNE for PIEFGs and behavioral strategies and PBE for IIEFGs.
  - *Incomplete information* – players do not deterministically know which game they are playing. They receive private signals, called **types**. This is today.
- **Cooperative games** – players form coalitions and utilities are defined over coalitions.
- **Other types** – repeated, stochastic, and so on.

We will study a special subclass: games with incomplete information and a **common prior**, due to Harsanyi (1967), also called **Bayesian games**.

12.2 An Example Before the Definition

Example 12.1 (Football game) *Two competing teams, ARG and FRA, are about to play each other. Each has privately settled on a gameplan: aim to win (Aggressive, A) or aim to draw (Passive, P). We call the gameplan the team's* ***type***. *It is a private signal, caused by external factors such as weather conditions, player injuries, or ground conditions. There are four type profiles, AA, AP, PA, PP.*

Each team then chooses an action: ATT (attack) or DEF (defend). The payoffs depend on the type profile, giving four different normal form games. Three of them are in Table 12.1; the PA matrix is the symmetric opposite of AP.

The three matrices are genuinely different games. In Γ_{AA}, ATT strictly dominates DEF for both teams ($1 > 0$ against ATT, $2 > 0$ against DEF), so the unique PSNE is (ATT, ATT) with payoff $(1, 1)$. In Γ_{PP},

		FRA ATT	FRA DEF
ARG	ATT	$1, 1$	$2, 0$
	DEF	$0, 2$	$0, 0$

		FRA ATT	FRA DEF
ARG	ATT	$2, 0$	$2, 1$
	DEF	$0, 1$	$0, 0$

		FRA ATT	FRA DEF
ARG	ATT	$0, 0$	$1, 0$
	DEF	$0, 1$	$-1, -1$

Table 12.1: The state games Γ_{AA} (left), Γ_{AP} (middle) and Γ_{PP} (right) of Example 12.1.

mutual ATT is the worst cell for both, $(-1, -1)$. So the right action depends on a type that nobody else observes.

This already tells us what a team is choosing. Not an action, since it does not know which matrix it is in. A *rule* that maps its own type to an action.

12.3 The Model

Definition 12.2 (Bayesian game) *A* ***Bayesian game*** *is a tuple* $\langle N, (\Theta_i)_{i\in N}, P, (\Gamma_\theta)_{\theta\in\times_{i\in N}\Theta_i}\rangle$ *where*

- N *is the set of players;*
- Θ_i *is the set of types of player* i*, and* $\Theta = \times_{i\in N}\Theta_i$*;*
- P *is a* ***common prior*** *distribution over* Θ *satisfying*

$$\sum_{\theta_{-i}\in\Theta_{-i}} P(\theta_i, \theta_{-i}) > 0, \qquad \forall \theta_i \in \Theta_i,\ \forall i \in N, \tag{12.1}$$

  i.e. the marginal of every type is positive (otherwise that type can be pruned from Θ_i*);*

- $\Gamma_\theta = \langle N, (A_i(\theta))_{i\in N}, (u_i(\theta))_{i\in N}\rangle$ *is the normal form* ***state game*** *for the type profile* θ*, with* $u_i : A \times \Theta \to \mathbb{R}$ *and* $A = \times_{i\in N} A_i$*. We assume* $A_i(\theta) = A_i$ *for all* θ*.*

The common prior is assumed to be common knowledge.

Stages of a Bayesian game.

1. $\theta = (\theta_i, \theta_{-i})$ is drawn according to P.
2. Each player observes her own type θ_i only.
3. Each player i picks an action $a_i \in A_i$.
4. Player i realizes payoff $u_i(a_i, a_{-i}; \theta_i, \theta_{-i})$.

Remark 12.3 (Why (12.1) is not a technicality) *Without it, the conditional* $P(\theta_{-i} \mid \theta_i)$ *we are about to write down is* $0/0$*, and a player of that type has no belief to best-respond to. It is the same* $0/0$ *we met in Lecture 11's doctor analogy: a room the booking system never actually sends anyone to, so there is no historical ratio to consult. The difference is that here we can simply legislate it away: a type with zero marginal never occurs, so deleting it changes nothing. We will see in Theorem 12.12 that the assumption is doing real work.*

12.4 Strategies and the Two Stages of Utility

Definition 12.4 (Strategy) *A **pure strategy** of player i is a map $s_i : \Theta_i \to A_i$. A **mixed strategy** is a map $\sigma_i : \Theta_i \to \Delta A_i$.*

Note the domain. A player with $|\Theta_i|$ types and $|A_i|$ actions has $|A_i|^{|\Theta_i|}$ pure strategies, not $|A_i|$: the type set is what the plan is a function of, not a label attached to it. With 3 types and 4 actions that is $4^3 = 64$ pure strategies, against 4 in any single state game.

Because the type is revealed part-way through, a player experiences utility at two different moments.

Definition 12.5 (Ex-ante utility) *The expected utility* before *observing one's own type:*

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta)\, u_i(\sigma(\theta); \theta) = \sum_{\theta \in \Theta} P(\theta) \sum_{(a_1,\ldots,a_n) \in A} \prod_{j \in N} \sigma_j(\theta_j)[a_j]\; u_i(a_1, \ldots, a_n; \theta_1, \ldots, \theta_n). \tag{12.2}$$

Once player i observes θ_i, her belief about the others' types is updated by Bayes' rule,

$$P(\theta_{-i} \mid \theta_i) = \frac{P(\theta_i, \theta_{-i})}{\sum_{\tilde{\theta}_{-i} \in \Theta_{-i}} P(\theta_i, \tilde{\theta}_{-i})}, \tag{12.3}$$

which is exactly where (12.1) is needed.

Definition 12.6 (Ex-interim utility) *The expected utility* after *observing one's own type:*

$$u_i(\sigma \mid \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i} \mid \theta_i)\, u_i(\sigma(\theta); \theta). \tag{12.4}$$

For independent types, observing θ_i carries no information about θ_{-i} and (12.3) reduces to $P(\theta_{-i})$. There is also an ex-post utility, which is the complete-information payoff once everything is revealed.

The two stages are linked by the identity

$$u_i(\sigma) = \sum_{\theta_i \in \Theta_i} P(\theta_i)\, u_i(\sigma \mid \theta_i), \tag{12.5}$$

that is, the ex-ante utility is the $P(\theta_i)$-weighted average of the ex-interim utilities. Everything in Section 12.6 rests on (12.5).

Example 12.7 (Numerical: forming the ex-interim belief) *Let $\Theta_1 = \Theta_2 = \{H, L\}$ with common prior $P(H,H) = 0.3$, $P(H,L) = 0.1$, $P(L,H) = 0.2$, $P(L,L) = 0.4$, the first coordinate being Player 1's type. If Player 1 observes $\theta_1 = H$, her marginal is $0.3 + 0.1 = 0.4$, so by (12.3)*

$$P(\theta_2 = H \mid \theta_1 = H) = \frac{0.3}{0.4} = \frac{3}{4}, \qquad P(\theta_2 = L \mid \theta_1 = H) = \frac{1}{4}.$$

Note that these are not the unconditional marginals of Player 2, which are 0.5 each: the types here are correlated, so learning one's own type is informative about the opponent.

12.5 Two Standing Examples

Example 12.8 (Two player bargaining game) *Player 1 is a seller and her type is the price at which she is willing to sell; Player 2 is a buyer and his type is the price at which he is willing to buy. Here* $\Theta_1 = \Theta_2 = \{1, 2, \ldots, 100\}$ *and* $A_1 = A_2 = \{1, 2, \ldots, 100\}$*. If the seller's bid is at most the buyer's bid, trade happens at the average of the two, otherwise no trade.*

Suppose the types are generated independently and uniformly, so $P(\theta_2 \mid \theta_1) = P(\theta_2) = 1/100$ *and* $P(\theta_1 \mid \theta_2) = P(\theta_1) = 1/100$ *for all* θ_1, θ_2*, and the common prior is* $P(\theta_1, \theta_2) = 1/10000$*. The payoffs are*

$$u_1(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \frac{a_1+a_2}{2} - \theta_1 & \text{if } a_2 \geqslant a_1, \\ 0 & \text{otherwise,} \end{cases} \qquad u_2(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \theta_2 - \frac{a_1+a_2}{2} & \text{if } a_2 \geqslant a_1, \\ 0 & \text{otherwise.} \end{cases}$$

Example 12.9 (Sealed bid auction) *Two players, both willing to buy an object. Their values* θ_i *and bids* b_i *lie in* $[0, 1]$*. The allocation is*

$$O_1(b_1, b_2) = \mathbf{1}\{b_1 \geqslant b_2\}, \qquad O_2(b_1, b_2) = \mathbf{1}\{b_2 > b_1\},$$

and the beliefs are uniform and independent, $f(\theta_2 \mid \theta_1) = f(\theta_1 \mid \theta_2) = f(\theta_1, \theta_2) = 1$*. The winner pays her own bid, so* $u_i(b_1, b_2; \theta_1, \theta_2) = O_i(b_1, b_2)(\theta_i - b_i)$*.*

12.6 Equilibrium: Ex-ante and Ex-interim

There are two natural places to demand optimality, and it is not obvious they agree.

Definition 12.10 (Nash equilibrium, ex-ante) (σ^*, P) *is a* ***Nash equilibrium*** *if*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geqslant u_i(\sigma_i', \sigma_{-i}^*), \qquad \forall \sigma_i',\ \forall i \in N. \tag{12.6}$$

Definition 12.11 (Bayesian equilibrium, ex-interim) (σ^*, P) *is a* ***Bayesian equilibrium*** *if*

$$u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* \mid \theta_i) \geqslant u_i(\sigma_i'(\theta_i), \sigma_{-i}^* \mid \theta_i), \qquad \forall \sigma_i',\ \forall \theta_i \in \Theta_i,\ \forall i \in N. \tag{12.7}$$

The Nash condition takes the expectation over $P(\theta)$; the Bayesian condition takes it over $P(\theta_{-i} \mid \theta_i)$, and it must hold *for every type separately.* In both definitions the deviation σ_i' on the right may be replaced by a pure action $a_i \in A_i$, for exactly the reason it could in the MSNE characterisation: a mixture cannot beat the best of the actions it mixes over.

Theorem 12.12 *In finite Bayesian games, a strategy profile is a Bayesian equilibrium if and only if it is a Nash equilibrium.*

Proof: $(\Rightarrow)$ Let (σ^*, P) be a Bayesian equilibrium and let σ_i' be any strategy. Using (12.5) and then (12.7) type by type,

$$u_i(\sigma_i', \sigma_{-i}^*) = \sum_{\theta_i \in \Theta_i} P(\theta_i)\, u_i(\sigma_i'(\theta_i), \sigma_{-i}^* \mid \theta_i) \ \leqslant \sum_{\theta_i \in \Theta_i} P(\theta_i)\, u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* \mid \theta_i) = u_i(\sigma_i^*, \sigma_{-i}^*).$$

A non-negatively weighted average of pointwise inequalities is an inequality; that is the whole of this direction.

($\Leftarrow$) By contraposition. Suppose (σ^*, P) is not a Bayesian equilibrium, so there exist $i \in N$, $\theta_i \in \Theta_i$ and $a_i \in A_i$ with

$$u_i(a_i, \sigma^*_{-i} \mid \theta_i) > u_i(\sigma^*_i(\theta_i), \sigma^*_{-i} \mid \theta_i). \tag{12.8}$$

Construct $\hat{\sigma}_i$ that agrees with σ^*_i at every other type and plays a_i for sure at θ_i:

$$\hat{\sigma}_i(\theta'_i) = \sigma^*_i(\theta'_i) \ \ \forall \theta'_i \in \Theta_i \setminus \{\theta_i\}, \qquad \hat{\sigma}_i(\theta_i)[a_i] = 1, \quad \hat{\sigma}_i(\theta_i)[b_i] = 0 \ \ \forall b_i \in A_i \setminus \{a_i\}.$$

Then, splitting the sum in (12.5) at θ_i,

$$\begin{aligned} u_i(\hat{\sigma}_i, \sigma^*_{-i}) &= \sum_{\tilde{\theta}_i \in \Theta_i} P(\tilde{\theta}_i)\, u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma^*_{-i} \mid \tilde{\theta}_i) \\ &= \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i)\, u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma^*_{-i} \mid \tilde{\theta}_i) \ + \ P(\theta_i)\, u_i(\hat{\sigma}_i(\theta_i), \sigma^*_{-i} \mid \theta_i) \\ &> \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i)\, u_i(\sigma^*_i(\tilde{\theta}_i), \sigma^*_{-i} \mid \tilde{\theta}_i) \ + \ P(\theta_i)\, u_i(\sigma^*_i(\theta_i), \sigma^*_{-i} \mid \theta_i) \ = \ u_i(\sigma^*_i, \sigma^*_{-i}). \end{aligned}$$

So σ^* is not a Nash equilibrium. ■

Remark 12.13 (Where the strict inequality comes from) *In the last display, every term of the first sum is unchanged, so the strictness has to come from the single term $P(\theta_i)\, u_i(\cdot \mid \theta_i)$. Combining (12.8) with $P(\theta_i) > 0$ is what turns $\geqslant$ into $>$. Drop the positive-marginal assumption (12.1) and the theorem is false: a type that occurs with probability zero could deviate profitably ex-interim without moving the ex-ante utility at all.*

Theorem 12.14 *Every finite Bayesian game has a Bayesian equilibrium. (Finite: the sets of players, actions and types are all finite.)*

Proof: The idea is to convert the Bayesian game into a complete-information normal form game, by treating *each type of every player as its own player*. We spell this out for two players $N = \{1, 2\}$, each with two types, $\Theta_1 = \{\theta^1_1, \theta^2_1\}$ and $\Theta_2 = \{\theta^1_2, \theta^2_2\}$; nothing in the argument below uses either the number of players or the number of types being two, so it extends verbatim to any finite N and any finite $(\Theta_i)_{i \in N}$.

Step 1: build the expanded (agent-form) game. Let the expanded player set be

$$\widetilde{N} \ = \ \bigcup_{i \in N} \Theta_i \ = \ \{\theta^1_1, \theta^2_1, \theta^1_2, \theta^2_2\}.$$

That is, every type of every player becomes its own player in a new, complete-information game. The action set available to the agent θ^j_i in the expanded game is exactly A_i, the action set that player i had under type θ^j_i in the original Bayesian game – this is well defined because we assumed $A_i(\theta) = A_i$ for all θ in Definition 12.2.

Step 2: define the expanded game's payoffs. We must specify a payoff function for each agent θ^j_i over action profiles of $\widetilde{N}$. For pure actions $a_{\theta^1_1} \in A_1$, $a_{\theta^2_1} \in A_1$, $a_{\theta^1_2} \in A_2$, $a_{\theta^2_2} \in A_2$ – one action chosen by each agent in $\widetilde{N}$ – define the payoff to agent θ^1_1 by

$$\tilde{u}_{\theta^1_1}\big(a_{\theta^1_1}, a_{\theta^2_1}, a_{\theta^1_2}, a_{\theta^2_2}\big) \ := \ P(\theta^1_2 \mid \theta^1_1)\, u_1\big(a_{\theta^1_1}, a_{\theta^1_2}; \theta^1_1, \theta^1_2\big) \ + \ P(\theta^2_2 \mid \theta^1_1)\, u_1\big(a_{\theta^1_1}, a_{\theta^2_2}; \theta^1_1, \theta^2_2\big), \tag{12.9}$$

where u_1 is player 1's payoff function from the original Bayesian game. In words: agent θ_1^1 only cares about her own action $a_{\theta_1^1}$ and about what the two possible agents θ_2^1, θ_2^2 standing in for the opponent's types do – weighted by how likely each of those types is, conditional on being type θ_1^1 herself. She does not care at all what $a_{\theta_1^2}$ is: that agent represents a version of player 1 with a different type, on a different branch of the tree of nature, and the two never interact. The payoffs of the other three agents $\theta_1^2, \theta_2^1, \theta_2^2$ are defined symmetrically, each using the appropriate conditional probabilities and the original u_1 or u_2.

Step 3: extend to mixed strategies, and identify the expanded game with the Bayesian game. Let $(\sigma_{\theta_1^1}, \sigma_{\theta_1^2}, \sigma_{\theta_2^1}, \sigma_{\theta_2^2})$ be a mixed strategy profile in the expanded game, where $\sigma_{\theta_i^j} \in \Delta A_i$. The expected payoff to agent θ_1^1 under this profile is, by definition of expected utility in a normal form game,

$$\begin{aligned}
&\tilde{u}_{\theta_1^1}\big(\sigma_{\theta_1^1}, \sigma_{\theta_1^2}, \sigma_{\theta_2^1}, \sigma_{\theta_2^2}\big) \\
&= \sum_{a_{\theta_2^2} \in A_2} \sum_{a_{\theta_1^2} \in A_1} \sum_{a_{\theta_2^1} \in A_2} \sum_{a_{\theta_1^1} \in A_1} \sigma_{\theta_1^1}(a_{\theta_1^1})\, \sigma_{\theta_2^1}(a_{\theta_2^1})\, \sigma_{\theta_1^2}(a_{\theta_1^2})\, \sigma_{\theta_2^2}(a_{\theta_2^2})\; \tilde{u}_{\theta_1^1}\big(a_{\theta_1^1}, a_{\theta_1^2}, a_{\theta_2^1}, a_{\theta_2^2}\big) \\
&= \sum_{a_{\theta_2^1} \in A_2} P(\theta_2^1 \mid \theta_1^1)\, \sigma_{\theta_2^1}(a_{\theta_2^1})\, u_1\big(a_{\theta_1^1}, a_{\theta_2^1}; \theta_1^1, \theta_2^1\big) \\
&\qquad + \sum_{a_{\theta_2^2} \in A_2} P(\theta_2^2 \mid \theta_1^1)\, \sigma_{\theta_2^2}(a_{\theta_2^2})\, u_1\big(a_{\theta_1^1}, a_{\theta_2^2}; \theta_1^1, \theta_2^2\big).
\end{aligned} \tag{12.10}$$

The first equality is just the definition of expected payoff over the joint mixed profile of the four agents in $\widetilde{N}$; substituting (12.9) for $\tilde{u}_{\theta_1^1}$ inside the sum and summing out the $\sigma_{\theta_1^1}(a_{\theta_1^1})$ and $\sigma_{\theta_1^2}(a_{\theta_1^2})$ factors that do not appear in either term of (12.9) – each such factor sums to 1 over its own action set, since it is a probability distribution – collapses the quadruple sum down to (12.10).

Now identify each mixed strategy $\sigma_{\theta_i^j}$ of agent θ_i^j in the expanded game with player i's choice of mixed action at type θ_i^j in the original Bayesian game, writing $\sigma_{\theta_i^j}(a_{\theta_i^j}) =: \sigma_i(\theta_i^j)[a_{\theta_i^j}]$. Under this identification, $a_{\theta_1^1} \in A_1$ in (12.10) is a free pure action that agent θ_1^1 plays, so (12.10) reads, for that fixed pure action $a_{\theta_1^1}$,

$$\sum_{\theta_2 \in \Theta_2} P(\theta_2 \mid \theta_1^1)\, u_1\big(a_{\theta_1^1}, \sigma_2(\theta_2); \theta_1^1, \theta_2\big) \;=\; u_1\big(a_{\theta_1^1}, \sigma_2 \mid \theta_1^1\big),$$

which is exactly the ex-interim utility (12.4) of player 1 at type θ_1^1, playing the pure action $a_{\theta_1^1}$ against player 2's strategy σ_2. Averaging over $a_{\theta_1^1}$ according to $\sigma_1(\theta_1^1)$ recovers

$$\tilde{u}_{\theta_1^1}\big(\sigma_{\theta_1^1}, \sigma_{\theta_1^2}, \sigma_{\theta_2^1}, \sigma_{\theta_2^2}\big) \;=\; u_1\big(\sigma_1(\theta_1^1), \sigma_2 \mid \theta_1^1\big) \;=\; u_1(\sigma_1, \sigma_2 \mid \theta_1^1), \tag{12.11}$$

where the last equality just records that $u_1(\sigma \mid \theta_1^1)$ depends on σ_1 only through $\sigma_1(\theta_1^1)$, by the definition of ex-interim utility. The same computation, mutatis mutandis, shows that every agent's expected payoff in the expanded game equals the corresponding ex-interim utility in the original game: $\tilde{u}_{\theta_i^j}(\cdot) = u_i(\sigma \mid \theta_i^j)$.

So a mixed strategy profile $(\sigma_{\theta_1^1}, \sigma_{\theta_1^2}, \sigma_{\theta_2^1}, \sigma_{\theta_2^2})$ of the expanded, complete-information game $\widetilde{N}$ is exactly the same object as a strategy profile $\sigma = (\sigma_1, \sigma_2)$ of the original Bayesian game, with $\sigma_i(\theta_i^j) := \sigma_{\theta_i^j}$, and under this identification the two games' payoffs agree, agent by agent.

Step 4: invoke Nash's theorem. The expanded game has a finite player set $\widetilde{N}$ (finite because N and every Θ_i are finite) and each agent θ_i^j has a finite action set A_i (finite because the original Bayesian game is finite). By Nash's theorem, every finite normal form game has a mixed strategy Nash equilibrium; let $(\sigma^*_{\theta_1^1}, \sigma^*_{\theta_1^2}, \sigma^*_{\theta_2^1}, \sigma^*_{\theta_2^2})$ be one. By definition of Nash equilibrium in the expanded game, no agent θ_i^j can improve $\tilde{u}_{\theta_i^j}$ by unilaterally changing her own mixed action, holding the other three agents' strategies fixed.

Step 5: translate back. Let σ^* be the strategy profile of the original Bayesian game corresponding to $(\sigma^*_{\theta_1^1}, \sigma^*_{\theta_1^2}, \sigma^*_{\theta_2^1}, \sigma^*_{\theta_2^2})$ under the identification of Step 3. Fix any player i, any type θ_i^j, and any alternative mixed action $\sigma_i'(\theta_i^j) \in \Delta A_i$ for that type. Deviating to $\sigma_i'(\theta_i^j)$ in the Bayesian game, holding every other type of every player fixed at σ^*, corresponds – again by the identification of Step 3, which showed the two games' payoffs agree termwise – exactly to agent θ_i^j unilaterally deviating to $\sigma_i'(\theta_i^j)$ in the expanded game, holding the other three agents fixed at σ^*. Since $(\sigma^*_{\theta_1^1}, \sigma^*_{\theta_1^2}, \sigma^*_{\theta_2^1}, \sigma^*_{\theta_2^2})$ is a Nash equilibrium of the expanded game, this deviation cannot increase $\tilde{u}_{\theta_i^j}$, i.e. by (12.11) it cannot increase $u_i(\sigma^* \mid \theta_i^j)$. Since i, θ_i^j, and the alternative $\sigma_i'(\theta_i^j)$ were arbitrary, this is exactly condition (12.7): σ^* is a Bayesian equilibrium of the original game. ■

Remark 12.15 (What the expanded game is really doing) *Two agents built from the same player, say θ_1^1 and θ_1^2, never move against each other and never appear in one another's payoff function – Eq. (12.9) for $\tilde{u}_{\theta_1^1}$ has no $a_{\theta_1^2}$ in it at all. They are the same physical player standing at two mutually exclusive branches of nature's move, so treating them as separate, non-interacting agents costs nothing: whichever branch is realized, only one of them ever actually acts. What each agent inherits from the original game is precisely her own type's conditional belief $P(\theta_{-i} \mid \theta_i)$ over the* other *players' types, which is exactly the ex-interim object Definition 12.6 was built around. This is the same device used in Lecture 11 to reduce equilibrium existence in general games with incomplete information to Nash's theorem: split one player into one player per information state, and a hard problem becomes one Nash's theorem already covers.*

12.7 Auctions as Bayesian Games

Return to Example 12.9. We look for a Bayesian equilibrium in linear strategies, $s_i(\theta_i) = \alpha_i \theta_i$ with $\alpha_i > 0$.

12.7.1 First price auction

Here $u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_1)\mathbf{1}\{b_1 \geqslant b_2\}$ and $u_2 = (\theta_2 - b_2)\mathbf{1}\{b_1 < b_2\}$. Fix Player 2 at $s_2(\theta_2) = \alpha_2\theta_2$. Player 1 with value θ_1 maximises her ex-interim utility:

$$\begin{aligned}\max_{b_1 \in [0,1]} \; u_1(b_1, \sigma_2^* \mid \theta_1) &= \max_{b_1 \in [0,1]} \int_0^1 f(\theta_2 \mid \theta_1)(\theta_1 - b_1)\mathbf{1}\{b_1 \geqslant \alpha_2\theta_2\}\, d\theta_2 \\ &= \max_{b_1 \in [0,1]} \; (\theta_1 - b_1)\frac{b_1}{\alpha_2} \qquad \Longrightarrow \qquad b_1 = \frac{\theta_1}{2}. \end{aligned} \tag{12.12}$$

The objective is a concave quadratic in b_1 with roots at 0 and θ_1, so the maximum sits at the midpoint. Notice that α_2 scales the objective but not the maximiser: Player 1's best reply does not depend on how aggressively Player 2 shades. By symmetry,

$$s_1^*(\theta_1) = \frac{\theta_1}{2}, \qquad s_2^*(\theta_2) = \frac{\theta_2}{2}$$

is a Bayesian equilibrium. In the Bayesian game induced by a uniform prior on the first price auction, bidding half of the true value is a Bayesian equilibrium. Shading by a factor of two is not timidity: it is the exact trade-off between winning less often and paying less when you do win.

12.7.2 Second price auction

Now the highest bidder wins but pays the second highest bid, so $u_1 = (\theta_1 - b_2)\mathbf{1}\{b_1 \geqslant b_2\}$ and $u_2 = (\theta_2 - b_1)\mathbf{1}\{b_1 < b_2\}$. Player 1 maximises

$$\int_0^1 f(\theta_2 \mid \theta_1)(\theta_1 - s_2(\theta_2))\mathbf{1}\{b_1 \geqslant s_2(\theta_2)\}\, d\theta_2 = \int_0^1 1 \cdot (\theta_1 - \alpha_2\theta_2)\mathbf{1}\Big\{\theta_2 \leqslant \frac{b_1}{\alpha_2}\Big\}\, d\theta_2 = \frac{1}{\alpha_2}\Big(b_1\theta_1 - \frac{b_1^2}{2}\Big),$$

which is increasing in b_1 up to $b_1 = \theta_1$ and gains nothing beyond it, so it is maximised at $b_1 = \theta_1$. Similarly $b_2 = \theta_2$.

The uniform prior was not essential. If θ_1 and θ_2 were drawn from an arbitrary but independent distribution with density f and cdf F, the objective would be

$$\int_0^{b_1/\alpha_2} f(\theta_2)(\theta_1 - \alpha_2\theta_2)\, d\theta_2 = \theta_1 F\left(\frac{b_1}{\alpha_2}\right) - \alpha_2 \int_0^{b_1/\alpha_2} \theta_2 f(\theta_2)\, d\theta_2.$$

Differentiating with respect to b_1,

$$\theta_1 \frac{1}{\alpha_2} f\left(\frac{b_1}{\alpha_2}\right) - \alpha_2 \cdot \frac{b_1}{\alpha_2} f\left(\frac{b_1}{\alpha_2}\right)\frac{1}{\alpha_2} = 0 \quad \Longrightarrow \quad f\left(\frac{b_1}{\alpha_2}\right)(b_1 - \theta_1) = 0, \tag{12.13}$$

so $b_1 = \theta_1$ whenever $f(b_1/\alpha_2) > 0$. The density factors out of the first-order condition entirely; all that survives is the requirement that it be positive. Similarly for Player 2. Hence for any independent positive prior, bidding one's true type is a Bayesian equilibrium of the induced Bayesian game in the second price auction.

12.8 Summary and What Comes Next

- A private signal is a **type**; a strategy is a map $\Theta_i \to A_i$; the common prior with positive marginals lets every type form a belief $P(\theta_{-i} \mid \theta_i)$.
- Utility is felt **ex-ante** (before the type) or **ex-interim** (after it), and the two are linked by (12.5).
- **Bayesian equilibrium** demands ex-interim optimality for every type. In finite Bayesian games it coincides with Nash equilibrium (Theorem 12.12), and it always exists (Theorem 12.14), by splitting each type into its own player in an expanded complete-information game and invoking Nash's theorem there.
- First price with uniform values: bid $\theta/2$. Second price: bid θ, for any independent positive prior.

That last line is a hint at the sequel. The second price rule made truth-telling optimal regardless of the prior, and the first price rule did not. From next week we stop analysing given games and start *designing* the rules so that a desired behaviour becomes equilibrium behaviour. That is mechanism design.

[MSZ] M. Maschler, E. Solan, S. Zamir, *Game Theory*, Cambridge University Press. Sections 9.4 and 10.5.

[Har67] J. C. Harsanyi, “Games with Incomplete Information Played by Bayesian Players”, *Management Science*, 1967–68.

[Nar] Y. Narahari, *Game Theory and Mechanism Design*, World Scientific and IISc Press.