

Bayesian Games and Equilibrium

P. Battigalli

Bocconi University

Game Theory: Analysis of Strategic Thinking

11 March , 2026

Abstract

In the previous lecture, we represented incomplete information using games with payoff uncertainty. This allows for a relatively straightforward extension of the rationalizability concept, which characterizes the information-dependent behavioral implications of rationality and common belief in rationality (also the pure SCE concept can be easily extended). *If rationalizability gives a unique prediction for each type of each player, the resulting profile of functions is an equilibrium in an obvious sense, extending the Nash equilibrium concept (NE). If the prediction is not unique, to obtain a meaningful extension of NE we have to enrich the mathematical structure with a description of players' information-dependent exogenous beliefs* (beliefs about parameters, or exogenous variables) and obtain a—so called—**Bayesian game**. Equilibria of Bayesian games, called **Bayesian equilibria**, are analyzed and illustrated. We also hint at rationalizability for Bayesian games. [These slides summarize and complement Section 8.4 and part of Section 8.5 of Ch. 8 of GT-AST.]

Games with Payoff Uncertainty and Rationalizability

- Recall: we can represent (simultaneous) strategic interaction under *incomplete information* with a **game with payoff uncertainty**

$$\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i : \Theta \times A \rightarrow \mathbb{R})_{i \in I} \rangle,$$

which is (informally) assumed to be common knowledge.

- Interpretation*: θ_0 affects the payoffs of somebody and nobody knows it, θ_i is i 's **private information** ($i \in I$), e.g., personal *traits* of i (tastes and other personal features like ability), but it could be also information privately acquired before the interaction; θ_i may be *payoff irrelevant* and yet relevant for strategic reasoning.
- We analyzed [directed] **rationalizability** for games with payoff uncertainty, $(\rho^m(\Theta \times A))_{m=1}^\infty [(\rho_\Delta^m(\Theta \times A))_{m=1}^\infty]$, to characterize the information-dependent *behavioral implications of Rationality and Common Belief in Rationality* [under transparent belief restrictions].

Equilibrium and Beliefs: the Simplest Case

- Rationalizability yields a set-valued prediction for each $i \in I$ and each information-type θ_i : $\theta_i \mapsto A_i^\infty(\theta_i) := \left(\text{proj}_{\Theta_i \times A_i} \rho^\infty(\Theta \times A) \right)_{\theta_i}$.
That is, the projection onto $\Theta_i \times A_i$ is a (typically *non-Cartesian*) subset of $\Theta_i \times A_i$, the section at θ_i of this subset is the **set** $A_i^\infty(\theta_i)$ **of rationalizable actions for type** θ_i .
- If rationalizability yields a *unique* action for every $\theta_i \in \Theta_i$ and $i \in I$,
 - we get as a *solution* a profile of *functions* $s = (s_i : \Theta_i \rightarrow A_i)_{i \in I}$ [for every $i \in I$ and $\theta_i \in \Theta_i$, $s_i(\theta_i)$ is the unique element of $A_i^\infty(\theta_i)$].
 - With this, it is easily verified that s is an **equilibrium** in an obvious sense: no player i , given any θ_i , has an incentive to deviate from the unique rationalizable action $s_i(\theta_i)$ given conjecture s_{-i} specifying how co-players are supposed to behave as a function of their private information. (Cf. examples in the previous lecture.)
- How can we define an equilibrium more generally, that is, when rationalizability does *not* yield a unique action for each θ_i of each i ?

The Need for Exogenous Beliefs: A numerical Example

Example

(Rowena *knows* the *true* payoff matrix, Colin does *not*):

$$\hat{G} :$$

θ'	ℓ	r
a	4,0	2,1
b	3,1	1,0

θ''	ℓ	r
a	1,1	0,0
b	0,1	2,0

► Is r a best response to $(s_1(\theta') = a, s_1(\theta'') = b)$? No, if $P^2(\theta') < \frac{1}{2}$, else Yes. ► In “equilibrium” $s_1(\theta') = a$ (dominance). Now suppose $P^2(\theta') < \frac{1}{2}$. - Then the b.r. of Colin (pl. 2) to any s_1 such that $s_1(\theta') = a$ is ℓ . - The b.r. of Rowena to ℓ given θ'' is a . ► “Equilibrium”: $s_1(\theta') = s_1(\theta'') = a, s_2 = \ell$.

Note: This makes sense if Rowena “knows” $P^2(\theta')$.

Scenario with Heterogeneous Populations

To give equilibrium conditions we can introduce beliefs $p^i \in \Delta(\Theta_{-i})$ for each i , and we have to say something about what each i believes about $(p^j)_{j \neq i}$, what each i believes about what each $j \neq i$ believes about $(p^k)_{k \neq j}$ etc.

Consider first the following scenario of *anonymous interaction with heterogeneous agents*:

- *distributed information* (neglect Θ_0);
- agent playing in role i drawn at random from large population i , **fraction/density** $q_i(\theta_i)$ of agents in role/population i have traits θ_i (e.g., ability, strength, or tastes), with Θ_i finite;
- statistical distributions $(q_i)_{i \in I} \in \times_{i \in I} \Delta(\Theta_i)$ are *commonly known*;
- For all $i \in I$ and $\theta_{-i} \in \Theta_{-i}$, the probability of facing co-players with private information θ_{-i} is

$$p^i(\theta_{-i}) := \prod_{j \neq i} q_j(\theta_j).$$

Simple Bayesian Games with Type-Independent Beliefs

- The interactive situation represented by $\Gamma = \langle I, (\Theta_i, A_i, u_i, p^i)_{i \in I} \rangle$ is assumed to be *transparent*. With this, we can meaningfully define as an **equilibrium** a profile of choice functions $(s_i)_{i \in I} \in \times_{i \in I} A_i^{\Theta_i}$ s.t.

$$\forall i \in I, \forall \theta_i \in \Theta_i, s_i(\theta_i) \in \arg \max_{a_i \in A_i} \mathbb{E}_{p^i, s_{-i}} (u_{i, \theta_i, a_i}). \quad (1)$$

[where $\mathbb{E}_{p^i, s_{-i}} (u_{i, \theta_i, a_i}) = \sum_{\theta_{-i} \in \Theta_{-i}} u_i(\theta_i, \theta_{-i}, a_i, s_{-i}(\theta_{-i})) p^i(\theta_{-i})$].

- **Note:** All “types” θ_i of i hold the *same belief* about types of others (we postpone the comment on “Bayesian”). When there are *private values*, this definition makes sense also without transparency of Γ :
 - Equilibria are stationary states of adaptive processes in situations of recurrent anonymous interaction where actions (or statistical distributions of actions) are observed ex post and all the individuals with the same private information take the same action.
 - When agents in role i with the same θ_i may take different actions, we obtain a notion of “randomized” equilibrium that extends *mixed* equilibrium to this environment.

A First-Price Auction with Independent Private Values

Example

A given object is offered for sale by means of a sealed-bid auction. The highest bidder gets the object and pays her bid (ties are broken at random) ▶ n competitors: $i = 1, \dots, n$ ▶ θ_i = value of the object for i in money units (private values), i is risk neutral ▶ $\Theta_i = [0, 1]$, uniform distribution \Rightarrow uniform product distribution of $\tilde{\theta}_{-i}$ on $[0, 1]^{n-1}$ (independent values) ▶ the payoff function is:

$$u_i(\theta, a) = \begin{cases} (\theta_i - a_i) \frac{1}{|\arg \max_j a_j|}, & \text{if } a_i = \max_j a_j \\ 0, & \text{if } a_i < \max_j a_j \end{cases}$$

Next we derive a symmetric, linear equilibrium s whereby $s_i(\theta_i) = \frac{n-1}{n}\theta_i$.

Derivation of the Symmetric Linear Equilibrium

A Conjecture-And-Verify Approach

- Assume i has a symmetric linear conjecture about competitors: for each $j \neq i$, $s_j(\theta_j) = k\theta_j$, where $k \in (0, 1)$,
- then $\mathbb{P}(\forall j \neq i, s_j(\theta_j) < a_i) = \mathbb{P}(\forall j \neq i, \theta_j < \frac{a_i}{k})$, and

$$\mathbb{P}\left(\forall j \neq i, \theta_j < \frac{a_i}{k}\right) = \begin{cases} \left(\frac{a_i}{k}\right)^{n-1}, & \text{if } \frac{a_i}{k} < 1 \\ 1, & \text{if } \frac{a_i}{k} \geq 1 \end{cases}$$

- Therefore i offers $a_i^*(\theta_i; k) := \min\left\{k, \arg \max_{a_i \geq 0} \left(\frac{a_i}{k}\right)^{n-1} (\theta_i - a_i)\right\}$.
- From FOC $\frac{\partial}{\partial a_i} \left(\frac{a_i}{k}\right)^{n-1} (\theta_i - a_i) = 0$ get $a_i^*(\theta_i; k) = \min\left\{\frac{n-1}{n}\theta_i, k\right\}$:
 - $\frac{\partial}{\partial a_i} \left(\frac{a_i}{k}\right)^{n-1} (\theta_i - a_i) = \frac{(n-1)}{k} \left(\frac{a_i}{k}\right)^{n-2} (\theta_i - a_i) - \left(\frac{a_i}{k}\right)^{n-1} = 0$,
 - $(n-1)(\theta_i - a_i) - a_i = 0$, $(n-1)\theta_i - na_i = 0$, $a_i = \frac{n-1}{n}\theta_i$.
- Symmetry implies $k = \frac{n-1}{n} \Rightarrow$ get $s_i(\theta_i) = a_i^*(\theta_i; \frac{n-1}{n}) = \frac{n-1}{n}\theta_i$.

Generalization: Heterogeneous Exogenous Beliefs

- A simple Bayesian game with type-independent beliefs specifies the same belief $p^i \in \Delta(\Theta_{-i})$ for all information-types $\theta_i \in \Theta_i$.
- But the information contained in θ_i may affect i 's beliefs about θ_{-i} , for example, because of *non-random, assortative matching*.
- Then specify, for each i , a **belief map** $\theta_i \mapsto p_{\theta_i}^i$, that is,

$$\left(p_{\theta_i}^i \right)_{\theta_i \in \Theta_i} \in \Delta(\Theta_{-i})^{\Theta_i}.$$

- *Equivalently*, specify, for each $i \in I$, a “**prior belief**” $P^i \in \Delta(\Theta)$ s.t. (in the finite case) $P^i(\{\theta_i\} \times \Theta_{-i}) > 0$ for each $\theta_i \in \Theta_i$. With this,

$$p_{\theta_i}^i(\theta_{-i}) = P^i(\theta_{-i}|\theta_i) = \frac{P^i(\theta_i, \theta_{-i})}{P^i(\{\theta_i\} \times \Theta_{-i})}.$$

- **Common Prior (on Θ) assumption:** There is some $P \in \Delta(\Theta)$ s.t. $p_{\theta_i}^i(\theta_{-i}) = P(\theta_{-i}|\theta_i)$ for all i, θ_i, θ_{-i} .
 - This *restrictive assumption* is reasonable in some applications: $\tilde{\theta}$ is drawn at random $\sim P$ (objective), θ_i is acquired information.

Simple Bayesian Games

- A (finite) **simple Bayesian game** is a (finite) mathematical structure

$$\Gamma = \langle I, \Theta_0, (\Theta_i, A_i, u_i, P^i)_{i \in I} \rangle$$

where $P^i \in \Delta(\Theta)$ is s.t. $\text{supp}(\text{marg}_{\Theta_i} P^i) = \Theta_i$, that is, (when Θ_i is finite) $P^i(\theta_i) := P^i(\Theta_0 \times \{\theta_i\} \times \Theta_{-i}) > 0$ for every $\theta_i \in \Theta_i$ (cf. Harsanyi 1967-68).

- **Note:** we put Θ_0 back into the picture. Maybe only θ_0 is directly payoff relevant, and θ_i is (thought to be) correlated with θ_0 .
- The map $\theta_i \mapsto P^i(\cdot | \theta_i)$ describes i 's *exogenous* beliefs as a function of i 's information-type. *Such maps are all that matters for strategic reasoning.* We informally assume that they are **transparent**: each player has a correct common belief of what such maps are.

Bayesian Equilibrium

Given $\Gamma = \langle I, \Theta_0, (\Theta_i, A_i, u_i, P^i)_{i \in I} \rangle$, we can give a meaningful definition of equilibrium:

Definition

A profile of functions $s = (s_i)_{i \in I} \in \times_{i \in I} A_i^{\Theta_i}$ is a **Bayesian equilibrium** if

$$\forall i \in I, \forall \theta_i \in \Theta_i, s_i(\theta_i) \in \arg \max_{a_i \in A_i} \mathbb{E}_{P^i(\cdot|\theta_i), s_{-i}} (u_{i, \theta_i, a_i}),$$

where

$$\mathbb{E}_{P^i(\cdot|\theta_i), s_{-i}} (u_{i, \theta_i, a_i}) = \sum_{\theta'_0, \theta'_{-i}} u_i(\theta'_0, \theta_i, \theta'_{-i}, a_i, s_{-i}(\theta'_{-i})) P^i(\theta'_0, \theta'_{-i} | \theta_i).$$

Why “Bayesian”?

- There was no reference to Bayes rule so far. Why “Bayesian” game and equilibrium? There are two reasons for this terminology:
 - 1 **Subjective probabilities:** In situations with *genuine incomplete information*, we interpret $P^i(\cdot|\theta_i)$ as a *subjective* probability measure. Bayes rule is the cornerstone of the subjectivist approach to probability and statistics. *Many theorists call “Bayesian” any model with subjective probabilities* (see Wikipedia entry).
 - 2 **Interesting special case:** Start with signal structure $\langle I, p, \Theta_0, (\pi_{i|0} : \Theta_0 \rightarrow \Delta(\Theta_i))_{i \in I} \rangle$, $p \in \Delta(\Theta_0)$ is a CP, $(\tilde{\theta}_i)_{i \in I}$ is a profile of signals, $\pi_{i|0}(\cdot|\theta_0) \in \Delta(\Theta_i)$ is the distribution of signal $\tilde{\theta}_i$ given θ_0 , signals of different players are *conditionally independent* given θ_0 : $P(\theta_0, \theta_i, \theta_{-i}) = \pi_{-i|0}(\theta_{-i}|\theta_0) \pi_{i|0}(\theta_i|\theta_0) p(\theta_0)$. Then, $P(\cdot|\theta_i)$ is determined by *Bayes rule* as follows:

$$P(\theta_0, \theta_{-i}|\theta_i) = \frac{\pi_{-i|0}(\theta_{-i}|\theta_0) \pi_{i|0}(\theta_i|\theta_0) p(\theta_0)}{\sum_{\theta'_0 \in \Theta_0} \pi_{i|0}(\theta_i|\theta'_0) p(\theta'_0)}.$$

Recall: It is possible that u_i depend *only* on θ_0 .

Ex Ante Strategic (or Normal) Form

- Consider, maybe just as a *metaphor*, the following **ex ante interpretation of Γ** : θ_i is acquired information, a signal; P^i is i 's *ex ante* belief (prior); decision rule $s_i : \Theta_i \rightarrow A_i$ is i 's “strategy” specifying *ex ante* what to do as a function of signal θ_i . To ease notation, let $S_i := A_i^{\Theta_i}$ denote the set of such “strategies”.
- For each profile $s = (s_i)_{i \in I} \in \times_{i \in I} S_i$, we can compute the **ex ante expected payoff** of each player i : let $s(\theta) = (s_i(\theta_i))_{i \in I}$, then

$$U_i(s) = \mathbb{E}_{P^i, s}(u_i) = \sum_{(\theta_0, \theta) \in \Theta_0 \times \Theta} u_i(\theta_0, \theta, s(\theta)) P^i(\theta_0, \theta).$$

Definition

The **ex ante strategic form** of Bayesian game Γ is the simultaneous-move game $\mathcal{AS}(\Gamma) = \langle I, (S_i, U_i)_{i \in I} \rangle$.

- **Remark** For any profile $s^* = (s_i^*)_{i \in I}$ in Γ , s^* is a Bayesian equilibrium of Γ IFF s^* is a Nash equilibrium of $\mathcal{AS}(\Gamma)$.

Example: Ex Ante Strategic Form

- Γ is \hat{G} of the previous example with the addition of $P^2(\theta') = p$.

$\theta' (p)$	ℓ	r
a	4,0	2,1
b	3,1	1,0

$\theta'' (1-p)$	ℓ	r
a	1,1	0,0
b	0,1	2,0





- The ex ante strategic form $\mathcal{AS}(\Gamma)$ is

$s_1 \setminus s_2$	ℓ	r
▶ a.a	$3p + 1, 1 - p$	$2p, p$
▶ a.b	$4p, 1 - p$	$2, p$
b.a	$2p + 1, 1$	$p, 0$
b.b	$3p, 1$	$2 - p, 0$

- where $x.y = x$ if θ' , y if θ'' .
- a dom b given $\theta' \Rightarrow$ a.a dom. b.a and a.b dom. b.b (delete);
- if $p < \frac{1}{2}$, ℓ dom $r, \Rightarrow (a.a, \ell)$;
- if $p > \frac{1}{2}$, r dom $\ell, \Rightarrow (a.b, r)$.

Computation. Rationalizability. Mixed actions.

- Finding the NEs of the ex ante strategic form $\mathcal{AS}(\Gamma)$ is a valid *method to compute* the *Bayesian equilibria* of Γ .
- Applying, instead, rationalizability to $\mathcal{AS}(\Gamma)$ (see previous example) may leave out some of the rationalizable profiles of Γ (=consistent with RCBR under transparency of $(P^i)_{i \in I}$). Apply directed rationalizability with $\Delta_{\theta_i} = \{P^i(\cdot|\theta_i)\}$ for all $i \in I$ and $\theta_i \in \Theta_i$ to get $\rho_{i,\Delta}(C_{0,-i}) =$
$$= \left\{ (\theta_i, a_i) : \exists \mu^i \in \Delta(C_{0,-i}), \text{marg}_{\Theta_{0,-i}} \mu^i = P^i(\cdot|\theta_i), a_i \in r_i(\mu^i, \theta_i) \right\}$$
and $\rho_{\Delta}(C) = \Theta_0 \times (\times_{i \in I} \rho_{i,\Delta}(C_{0,-i}))$; iterate ρ_{Δ} from $\Theta \times A$.
- We can also define and compute Bayesian equilibria in *mixed actions* $(\sigma_i : \Theta_i \rightarrow \Delta(A_i))_{i \in I}$, e.g., when different agents in population i with type θ_i take different actions, among which they are indifferent.
- The Bayesian equilibria in mixed actions σ of Γ can be recovered from the mixed equilibria α of $\mathcal{AS}(\Gamma)$: for every $i \in I$, $\theta_i \in \Theta_i$, and $a_i \in A_i$, $\sigma_i(\theta_i)(a_i) = \alpha_i \left(\left\{ s_i \in A_i^{\Theta_i} : s_i(\theta_i) = a_i \right\} \right)$.

-  BATTIGALLI, P., E. CATONINI, AND N. DE VITO (2025): *Game Theory: Analysis of Strategic Thinking*. Typescript, Bocconi University.
-  BATTIGALLI, P. (2025): *Mathematical Language and Game Theory*. Typescript, Bocconi University.
-  HARSANYI, J. (1967-68): "Games of Incomplete Information Played by Bayesian Players. Parts I, II, III," *Management Science*, 14, 159-182, 320-334, 486-502.
-  WIKIPEDIA, *Bayesian Probability*.
https://en.wikipedia.org/wiki/Bayesian_probability

FIRST-PRICE IPV AUCTION

(2-bidders case: i and $-i$)

θ_i - best reply to linear conjecture $\sigma_{-i}(\theta_{-i}) = k\theta_{-i}$

