

General Bayesian Games and Equilibrium

P. Battigalli

Bocconi University

Game Theory: Analysis of Strategic Thinking

23 March, 2026

Abstract

“*Incomplete information*” means lack of common knowledge of the rules of the game (e.g., of the outcome function), or of players’ preferences over (lotteries of) outcomes. The most basic representation of strategic interaction with incomplete information, a structure called *game with payoff uncertainty*, allows meaningful strategic analysis (e.g., a generalization of rationalizability and self-confirming equilibrium), but does not allow traditional equilibrium analysis, according to which each player is somehow able to “divinate” the decision functions of co-players. Such more standard analysis can be performed by adding to the basic structure a *belief structure*, which is an implicit representation of players’ exogenous interactive beliefs, thus obtaining a so called “*Bayesian*” game. Equilibria of Bayesian games can be computed as Nash equilibria of their strategic forms.

[These slides summarize and complement Sections 8.5-8.6 of Ch. 8 of GT-AST.]

Hierarchies of Beliefs

- Strategic reasoning should presumably consider not only (exogenous) beliefs about θ , but also beliefs about the (exogenous) beliefs of other players. Consider first a *finite, two-person game with distributed knowledge* of θ (for notational simplicity): $j \neq i$, in general, j is not only uncertain about θ_j but *also about* $p^i \in \Delta(\Theta_j)$.
- $p^i \in \Delta(\Theta_j)$ = **first-order** beliefs of i about θ_j , these are beliefs about the “primitive uncertainty.”
 - j is uncertain about $(\theta_i, p^i) \in \Theta_i \times \Delta(\Theta_j)$
 - $\Rightarrow j$ holds beliefs $q^j \in \Delta(\Theta_i \times \Delta(\Theta_j))$ = **second-order** beliefs, or beliefs about (primitive uncertainty and) the first-order beliefs of others.
 - By coherence: $\text{marg}_{\Theta_i} q^j = p^j$ [if q^j has finite supp., $p^j(\theta_i) = \sum_{\tilde{p}^i} q^j(\theta_i, \tilde{p}^i)$].
- No reason to stop at the second order \Rightarrow *beliefs hierarchies!*

Bayesian Games: The “Ex Ante” Metaphor

- To avoid hierarchies of beliefs, which are complex mathematical objects, Harsanyi suggested to use an *implicit representation* of beliefs about beliefs by means of a mathematical structure very similar to the one we used to define correlated equilibrium (6.2 of GT-AST, Lecture 9), but with a different interpretation.
- *Metaphor*: a **state** of the world $\omega \in \Omega$ is selected “at random,” each $i \in I$ initially holds “prior” $p_i \in \Delta(\Omega)$ then gets “signal” $t_i = \tau_i(\omega) \in T_i$ called the **type** of i , which also includes i 's private information about θ : $\theta_i = \vartheta_i(t_i)$.
- Why is it only a metaphor? Because *we do not really assume that players start all symmetrically ignorant and then “learn” their types*; we only claim that, *for equilibrium analysis*, it is “as if” this were the case.

General Bayesian Games: Definition

- To ease notation, assume *distributed knowledge* of θ (neglect Θ_0).

Definition

A **Bayesian Game** is a structure [assume *finite sets* for simplicity]

$$BG = \langle I, \Omega, (\Theta_i, T_i, A_i, \vartheta_i, \tau_i, p_i, u_i)_{i \in I} \rangle$$

where

- ▶ $\forall i \in I, \vartheta_i : T_i \rightarrow \Theta_i, \tau_i : \Omega \rightarrow T_i$ (onto), $\forall t_i \in T_i, p_i(\tau_i^{-1}(t_i)) > 0$
(no player ex ante rules out any type of hers),
- ▶ $\forall i \in I, u_i : \Theta \times A \rightarrow \mathbb{R}$.

- $(\vartheta_i \circ \tau_i)(\omega) = \vartheta_i(\tau_i(\omega))$ = **private information** of i about θ at **state** ω .
- The sub-structure $\langle I, \Omega, (\Theta_i, T_i, \tau_i, \vartheta_i, p_i)_{i \in I} \rangle$ is the **belief structure** (also called “type space”) of the Bayesian game.

Types As “Tastes & Thoughts”

- **Types** *à la* Harsanyi: $t_i \in T_i$ is called the **type** of player i , it determines i 's “tastes and thoughts.”
- Beside i 's private information ($\theta_i = \vartheta_i(t_i)$), t_i determines i 's *beliefs about every exogenous unknown*, such as θ_j ($j \neq i$) or opponents' beliefs about exogenous unknowns.
- Exogenous unknowns depend on ω . Beliefs of player i at state ω are given by $p_i(\cdot | \tau_i(\omega))$: for all t_i and $E \subseteq \Omega$

$$p_i(E | t_i) = \sum_{\omega' \in E \cap \tau_i^{-1}(t_i)} \frac{p_i(\omega')}{p_i(\tau_i^{-1}(t_i))}$$

(recall, we assumed $p_i(\tau_i^{-1}(t_i)) > 0$ for each t_i).

- We also write $p_i(t_i)[E] = p_i(E | t_i)$ to emphasize that i 's (subjective) probability of E is a function of t_i .

Bayesian Games: Comments

- **Comment 1:** We assume informally that the situation represented by BG is *commonly known*, or at least transparent to the players.
[Transparent=True and Commonly Believed]
- **Comment 2:** *Interpretation:*
 - (II) “genuine” Incomplete Information vs
 - (AIC) Asymmetric Information about the realization of an initial move of Chance.
 - There are important and relevant *differences* in interpretation, but the *Bayesian-game* mathematical structure *may represent both!*
- **Comment 3:** A Bayesian game is “**simple**” if, for each i , Harsanyi-types t_i coincide with information types θ_i ($\forall i, T_i = \Theta_i$ and ϑ_i is the identity), which means that exogenous beliefs are pinned down by private information. *In economic models it is often unclear whether types just represent private information, or they are also a parameterization of subjective beliefs going beyond private information.*

States and Beliefs About Types (Optional)

- **Preliminary:** Fix $f : X \rightarrow Y$ (*finite* sets for simplicity) and $\mu \in \Delta(X)$. The induced probability measure on Y is given by the **pushforward map** $\mu \mapsto \mu \circ f^{-1}$:

$$\forall E_Y \subseteq Y, (\mu \circ f^{-1})(E_Y) := \mu(f^{-1}(E_Y)) = \sum_{x:f(x) \in E_Y} \mu(x)$$

- Each map $\omega \mapsto p_i(\tau_i(\omega))[\cdot]$ is “transparent”. This will allow to unravel a *hierarchy of beliefs* from each type t_i of each player i .
- (Two players, $i \neq j$) First, derive from $p_i(t_i) \in \Delta(\Omega)$ the **beliefs** $\hat{p}_i(t_i) := p_i(t_i) \circ \tau_j^{-1} \in \Delta(T_j)$ **of a type t_i about the co-player’s type**:

$$\forall t_j \in T_j, \hat{p}_i(t_i)[t_j] = p_i(t_i)[\tau_j^{-1}(t_j)].$$

First-Order Belief Maps (Optional)

- Next derive the **first-order beliefs** $p_i^1(t_i) := \hat{p}_i(t_i) \circ \vartheta_j^{-1} \in \Delta(\Theta_j)$ of each t_i about θ_j :

$$\begin{aligned}\forall \bar{\theta}_j \in \Theta_j, p_i^1(t_i)[\bar{\theta}_j] &= \sum_{t_j: \vartheta_j(t_j) = \bar{\theta}_j} \hat{p}_i(t_i)[t_j] \\ &= \hat{p}_i(t_i)[\vartheta_j^{-1}(\bar{\theta}_j)].\end{aligned}$$

- We derive the **first-order belief map**

$$p_i^1 : T_i \rightarrow \Delta(\Theta_j).$$

Second-Order Belief Maps (Optional)

- The functions $(\vartheta_j, p_j^1) : T_j \rightarrow \Theta_j \times \Delta(\Theta_i)$ ($j = 1, 2$) are transparent.
- Therefore, the **second-order beliefs** of each t_i are given by $p_i^2(t_i) = \hat{p}_i(t_i) \circ (\vartheta_j, p_j^1)^{-1} \in \Delta(\Theta_j \times \Delta(\Theta_i))$:

$$\begin{aligned} \forall (\bar{\theta}_j, \bar{p}_j^1) \in \Theta_j \times \Delta(\Theta_i), p_i^2(t_i)[\bar{\theta}_j, \bar{p}_j^1] &= \sum_{t_j: (\vartheta_j(t_j), p_j^1(t_j)) = (\bar{\theta}_j, \bar{p}_j^1)} \hat{p}_i(t_i)[t_j] \\ &= \hat{p}_i(t_i)[(\vartheta_j, p_j^1)^{-1}(\bar{\theta}_j, \bar{p}_j^1)]. \end{aligned}$$

- It can be verified that $p_i^1(t_i)[\cdot] = \text{marg}_{\Theta_j} p_i^2(t_i)[\cdot]$.
- We derive the **second-order belief map**

$$p_i^2 : T_i \rightarrow \Delta(\Theta_j \times \Delta(\Theta_i))$$

Hierarchies of Exogenous Beliefs (Optional)

- The full recursive construction is as follows. Suppose that, for *each* $t_j \in T_j$, $p_j^1(t_j)[\cdot]$, \dots , $p_j^k(t_j)[\cdot]$ (beliefs of t_j up to order k) have been determined.

- Then we have a map

$$(\vartheta_j, p_j^1, p_j^2, \dots, p_j^k) : T_j \rightarrow \Theta_j \times \Delta(\Theta_i) \times \Delta(\Theta_i \times \Delta(\Theta_j)) \times \dots$$

- and the $(k+1)$ -**order beliefs** for of *each* type t_i of i are

$$p_i^{k+1}(t_i) = \hat{p}_i(t_i) \circ (\vartheta_j, p_j^1, p_j^2, \dots, p_j^k)^{-1}.$$

- It can be verified that $p_i^k(t_i)[\cdot]$ can be obtained from $p_i^{k+1}(t_i)[\cdot]$ via marginalization, as it should be.
- We derive the $(k+1)^{th}$ -order belief map $p_i^{k+1} : T_i \rightarrow \Delta(\Theta_j \times \Delta(\Theta_i) \times \Delta(\Theta_i \times \Delta(\Theta_j))) \times \dots$

- In general, players' choices depend not only on their basic private information, but more generally on their *types*.
- Types may be interpreted as true information (possibly correlated with the information of others) or simply as a “parameterization” of beliefs about θ and beliefs about co-players' beliefs.
- We now revert to notation $p_i(\cdot|t_i)$, more suggestive of the “information interpretation” of types and a bit easier to parse.

Definition of Bayesian Equilibrium

- Fix a profile of decision functions $\sigma_{-i} = (\sigma_j : T_j \rightarrow A_j)_{j \neq i}$.
- The **expected payoff** for **type** t_i **of choosing** a_i given σ_{-i} is

$$\begin{aligned} & \mathbb{E}_{\sigma_{-i}} (u_{i,a_i} | t_i) \\ &= \sum_{\omega \in \Omega} p_i(\omega | t_i) u_i(\vartheta_i(t_i), \vartheta_{-i}(\tau_{-i}(\omega)), a_i, \sigma_{-i}(\tau_{-i}(\omega))) \\ &= \sum_{t_{-i} \in T_{-i}} \hat{p}_i(t_{-i} | t_i) u_i(\vartheta_i(t_i), \vartheta_{-i}(t_{-i}), a_i, \sigma_{-i}(t_{-i})) \end{aligned}$$

where $\hat{p}_i(t_{-i} | t_i) = p_i(\tau_{-i}^{-1}(t_{-i}) | t_i)$ (cf. pages above).

Definition

A **Bayesian Equilibrium** of *BG* is a profile of decision functions $(\sigma_i : T_i \rightarrow A_i)_{i \in I}$ [often called “strategies”] such that

$$\forall i \in I, \forall t_i \in T_i, \sigma_i(t_i) \in \arg \max_{a_i \in A_i} \mathbb{E}_{\sigma_{-i}} (u_{i,a_i} | t_i).$$

Strategic-Form Payoffs

- Bayesian equilibrium of BG can be equivalently restated as a Nash equilibrium of an associated auxiliary game with *complete* information: the *ex ante* strategic form (there is also an *interim* strategic form, we do not consider it in these slides). Therefore we often say “**Bayes-Nash**” equilibrium.
- **Ex ante strategic form.** It refers to the metaphor that was previously introduced to explain the elements of the Bayesian game: $\sigma_i : T_i \rightarrow A_i$ is a **strategy** (contingent plan of action) formulated by i *ex ante*. The expected payoff induced by (σ_i, σ_{-i}) is

$$U_i(\sigma_i, \sigma_{-i}) =$$

$$\sum_{\omega \in \Omega} p_i(\omega) u_i(\vartheta_i(\tau_i(\omega)), \vartheta_{-i}(\tau_{-i}(\omega)), \sigma_i(\tau_i(\omega)), \sigma_{-i}(\tau_{-i}(\omega))) \quad (1)$$

Strategic Form and Equilibrium

Definition

(Strategic form) The **ex ante strategic form of BG** is the static game $\langle I, (\Sigma_i, U_i)_{i \in I} \rangle$, where, for each $i \in I$, U_i is defined by eq. (1) and $\Sigma_i := (A_i)^{T_i}$.

Theorem

(Bayesian and Nash equilibrium) A profile $(\sigma_i)_{i \in I}$ is a Bayesian equilibrium of BG if and only if it is a Nash equilibrium of the ex ante strategic form of BG (game $\langle I, (\Sigma_i, U_i)_{i \in I} \rangle$).

- Prove the theorem as an exercise: recall, for all $t_i \in T_i$, $\hat{p}_i(t_i) := p_i(\tau_i^{-1}(t_i)) > 0$; with this, use “book-keeping” and expected-utility tricks.
- **Hint:** Obtain $U_i(\sigma_i, \sigma_{-i}) = \sum_{t_i \in T_i} \hat{p}_i(t_i) \sum_{t_{-i} \in T_{-i}} \hat{p}_i(t_{-i} | t_i) u_i(\vartheta_i(t_i), \vartheta_{-i}(t_{-i}), \sigma_i(t_i), \sigma_{-i}(t_{-i}))$, maximize each t_i -term separately w.r.t. a_i . ♡

Example

- Consider the following game with payoff uncertainty where pl. 1 (row) is informed and pl. 2 (col.) is not (**key** payoffs of **pl. 2** in **bold**):

$$\hat{G}$$

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

- Let $\Theta_1 \cong \Theta$, $\Omega = \{\omega', \omega''\}$, $\vartheta_1(\omega') = \theta'$, $\vartheta_1(\omega'') = \theta''$, $\tau_1(\omega') = t'_1$, $\tau_1(\omega'') = t''_1$, $\tau_2(\omega') = \tau_2(\omega'') = \bar{t}_2$, $p_1(\theta') = p_2(\theta') = p$. Then the ex ante strategic form is

$\sigma_1 \backslash \sigma_2$	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	$3p, 1$	$2 - p, 0$

[Unique equilibrium (a.a, ℓ) obtained by iterated dominance IFF $p < 1/2$]

Bayesian and correlated equilibrium

- **Special case:** *complete information*, Θ singleton, or (pseudo more generally) for all $i \in I$, $a \in A$, $\theta', \theta'' \in \Theta$,

$$u_i(\theta', a) = u_i(\theta'', a).$$

- Even in this case, we may have a Bayesian game BG (called “**Bayesian elaboration**” of the complete information game G) with $|\Omega| > 1$ and $|T_i| > 1$ for some i .
- If BG has *common prior* (CP: $\forall i, p_i = p \in \Delta(\Omega)$) a Bayesian equilibrium of such BG is a **correlated equilibrium** of G ; *without CP*, an equilibrium of such BG is called **subjective correlated equilibrium** of G .

Equivalence of Bayesian Equilibrium and Rationalizability

Theorem

Fix a (finite) game with payoff uncertainty \hat{G} .

(1) For every Bayesian game BG based on \hat{G} , every Bayesian equilibrium σ of BG and every state $\omega \in \Omega$, the corresponding profile of information types and actions is rationalizable in \hat{G} :

$$(\vartheta_i(\tau_i(\omega)), \sigma_i(\tau_i(\omega)))_{i \in I} \in \rho^\infty(\Theta \times A).$$

(2) Conversely, there is a BG based on \hat{G} [possibly with heterogeneous priors] and a Bayesian equilibrium σ of BG such that




$$(\vartheta_i \circ \tau_i, \sigma_i \circ \tau_i)_{i \in I}(\Omega) = \rho^\infty(\Theta \times A),$$

where $\omega \mapsto (\vartheta_i \circ \tau_i, \sigma_i \circ \tau_i)_{i \in I}(\omega) = (\vartheta_i(\tau_i(\omega)), \sigma_i(\tau_i(\omega)))_{i \in I}$.

[A similar result holds for Directed Rationalizability.]

Sketch of Proof

- Use results on rationalizability and sets with the best reply property (BRP). In GT-AST you can find the proof for the **special case** when there is *complete information* and Bayesian equilibrium=subjective correlated equilibrium.
- Recall: $\hat{p}_i(\cdot|t_i) \in \Delta(T_{-i})$ beliefs of type t_i about co-players' types (under complete information, types are payoff-irrelevant).
- 1 Fix subjective CE σ . Let $C_i = \sigma_i(T_i) \subseteq A_i$ for each i . For each t_i , action $\sigma_i(t_i) \in C_i$ is justified by the following conjecture:
 $\hat{p}_i(\cdot|t_i) \circ (\sigma_{-i})^{-1} \in \Delta(\sigma_{-i}(T_{-i})) = \Delta(C_{-i})$. Thus, $C \subseteq \rho(C)$ (BRP) (with $C = \times_{i \in I} C_i$), which implies $\sigma(T) = C \subseteq \rho^\infty(A)$.
- 2 Find BG s.t. ... Let $C = \times_{i \in I} C_i = \rho^\infty(A)$. Then $C = \rho(C)$. Let $\Omega = \times_{i \in I} T_i$, let each T_i be a "copy" of C_i via a bijection σ_i :
 $T_i \xleftrightarrow{\sigma_i} C_i$. Each $a_i \in C_i$ is BR to some conjecture $\beta_i(a_i) \in \Delta(C_{-i})$.
Let $\hat{p}_i(t_{-i}|t_i) = \beta_i(\sigma_i(t_i))(\sigma_{-i}(t_{-i}))$ for all t_i and t_{-i} . Prior: e.g.,
 $p_i(\cdot) = \frac{1}{|T_i|} \sum_{t_i \in T_i} \hat{p}_i(\cdot|t_i)$. ♡

-  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.