Multistage Games with Payoff Uncertainty

Pierpaolo Battigalli
Bocconi University
Game Theory: Analysis of Strategic Thinking

22 April, 2026

Abstract

This lecture introduces the theory of multistage games with observed actions and *either incomplete information*, *or asymmetric information about an initial chance move*. It is focused on the description of these games and on the role of Bayes rule to model updated beliefs about co-players' types. The analytical representation blends the parameterized payoff functions (with differential information about parameters) used to study static games with incomplete information in Chapter 8 with the sequential game-tree representation of multistage games of Chapter 9. Graphically, the game is represented as a *forest* of trees, one for each possible profile of payoff functions, connected by information sets showing that (active) players cannot distinguish the different types of co-players. Conditional probability systems about the types and strategies of co-players yield "personal assessments" comprising conditional probabilities of actions and conditional probabilities of nodes in the information sets. Personal assessments must satisfy Bayes rule.

[These slides summarize and, in part, complement Sections 1, 2 and part of section 3 of Chapter 15 of GT-AST.]

Introduction

- We want to study multistage games with observed actions and incomplete information.
  - The same analytical tools can be applied to study games with complete, but imperfect and asymmetric information, e.g., Poker.
- To represent **incompleteness of information** in *static games*:
  - we used parameterized payoff functions $(u_i : \Theta \times A \to \mathbb{R})_{i \in I}$,
  - $\theta = (\theta_i)_{i \in I} \in \Theta$, where each $i \in I$ knows (only) θ_i (for simplicity, *we neglect residual uncertainty*).
- To represent the rules of interaction of a **multistage game with observed actions**:
  - we used a structure $\langle I, (A_i, \mathcal{A}_i(\cdot))_{i \in I} \rangle$ with A_i=i's potentially feasible actions, $\mathcal{A}_i(\cdot)$=i's constraint correspondence;
  - we derived a **tree** of histories $(\bar{H}, \preceq)$, with $\preceq$="prefix of", and $\bar{H} = H \cup Z$ (Z: terminal histories).
- To represent *incomplete information* (payoff uncertainty), posit $(u_i : \Theta \times Z \to \mathbb{R})_{i \in I}$.

Graphical representation: a simple example

- We get a *forest of game trees* indexed by $\theta \in \Theta$, each one with payoff functions $\left(u_{i,\theta} : Z \to \mathbb{R}\right)_{i \in I}$.
- Asymmetric information is graphically represented by **information sets**: for every $i \in I$, $\bar{\theta}_i \in \Theta_i$, and $h \in H$, the *indistinguishable* "nodes" $\left(\left(\bar{\theta}_i, \theta_{-i}\right), h\right)$ $\left(\theta_{-i} \in \Theta_{-i}\right)$ are connected by a dashed line. In Figure A, $i = 2$, $\Theta_2 = \left\{\bar{\theta}_2\right\}$, $\Theta_1 \cong \Theta = \left\{\theta', \theta''\right\}$, $h = (r)$.

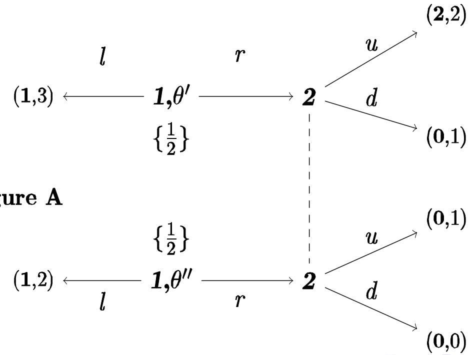

Figure A

Intuitive analysis of the example

- In Figure A we also represent the initial, exogenous belief of pl. 2 about θ. His (posterior) belief upon observing r may well be different: r is dominated for θ'', but it is justifiable for θ'; if 2 strongly believes in 1's rationality, then the posterior belief is $\mu_2\left(\theta''|r\right) = 1 - \mu_2\left(\theta'|r\right) = 0$.

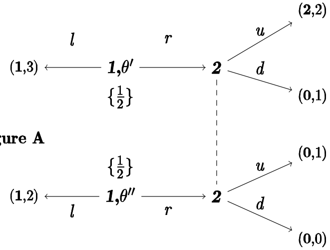

Figure A

- Also, d is conditionally dominated by u. A rational pl. 2 would choose u after r. Anticipating this, a rational player 1 of type θ' chooses r.
- Intuitive solution: $(r'.\ell'', u)$ [where $r'.\ell''$=(r if θ', ℓ if θ'')]

Multistage games with payoff uncertainty

A **multistage game with payoff uncertainty** (and observed actions) is a structure

$$\hat{\Gamma} = \left\langle I, (\Theta_i, A_i, \mathcal{A}_i(\cdot), u_i)_{i \in I} \right\rangle$$

where

- $\left\langle I, (A_i, \mathcal{A}_i(\cdot))_{i \in I} \right\rangle$ is a **multistage game tree;**
- $\theta_i \in \Theta_i$ are information **types** (traits, tastes, private information);
- $\Theta := \times_{i \in I} \Theta_i$ (more generally, $\Theta \subseteq \times_{i \in I} \Theta_i$);
- $u_i : \Theta \times Z \to \mathbb{R}$ is i's parameterized **payoff function;**
  - u_i is derivable from $g : \Theta \times Z \to Y$ and $v_i : \Theta_i \times Y \to \mathbb{R}$;
  - generalization: $\mathcal{A}_i(\cdot, \cdot) : \Theta_i \times H \rightrightarrows A_i$, $a_i \in \mathcal{A}_i(\theta_i, h)$ feasible for θ_i given h;
  - $\Theta = \{\bar{\theta}\}$ singleton (or θ-indep. payoff) $\Rightarrow$ *complete information.*
- To ease notation, we neglect residual uncertainty (Θ_0 is a singleton, not shown).

An example

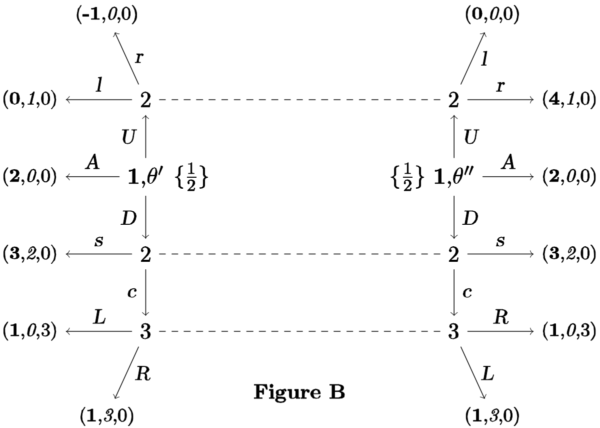

Figure B

- $\Theta = \{\theta', \theta''\} \cong \Theta_1$, only pl. 1 is informed.
  - Dashed lines join "nodes" (θ, h) in the same information set of the active player (pl. 2 and 3 have the same information).
  - (Numbers in $\{\cdot\}$: exogenous common prior on Θ, not part of $\hat{\Gamma}$.)

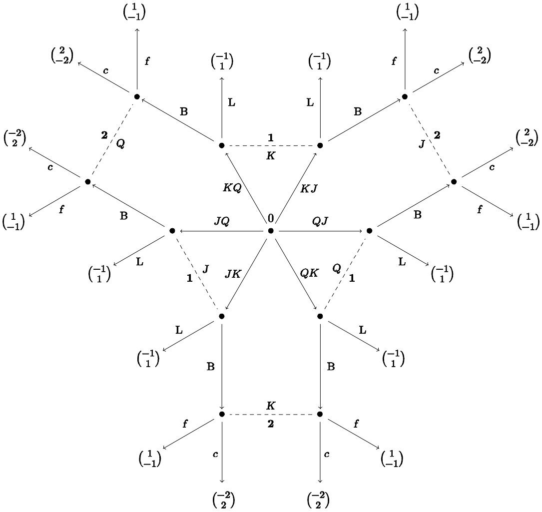

- **Mini-Poker:** θ_i is acquired information (not personal trait).
  - $\Theta_i = \{J, Q, K\}$, $\Theta = \{JK, JQ, KQ, KJ, QJ, QK\} \subset \Theta_1 \times \Theta_2$.
  - Initial node of pl. 0=chance move (random order of cards).

Personal Assessments

- For simplicity, focus on two-person games.
  - With complete information, we could represent a conjecture of i as $\beta^i = \left(\beta^i(\cdot|h)\right)_{h\in H} = \beta_{-i} \in \times_{h\in H}\Delta\left(\mathcal{A}_{-i}(h)\right)$, which is like a behavior strategy of $-i$, and we obtained a subjective decision tree for i.
  - With incomplete info, we need a **personal assessment** $\left(\beta^i, \mu_i\right)$**:** **conjecture** β^i must have the form
$$\beta^i = \left(\beta^i(\cdot|\theta_{-i}, h)\right)_{\theta_{-i}\in\Theta_{-i}, h\in H} \in \left(\times_{h\in H}\Delta\left(\mathcal{A}_{-i}(h)\right)\right)^{\Theta_{-i}},$$
because pl. i of type θ_i expects that $-i$'s behavior depends of his type θ_{-i}; also, we have to specify updated beliefs about θ_{-i} given each $h \in H$: a **system of beliefs**
$$\mu_i = \left(\mu_i(\cdot|h)\right)_{h\in H} \in \left(\Delta\left(\Theta_{-i}\right)\right)^H$$
(we also write $\mu_i(\cdot|\theta_i, h)$ to emphasize that θ_i is known/given).
  - **Note:** μ_i should be derived from β^i (if possible) using *Bayes rule*.
- From $\left(\beta^i, \mu_i\right)$ we will derive conditional values for type θ_i.

Conditional Probability Systems (CPSs)

- In the (rationalizability) analysis of *static* games with *incomplete* information, we considered conjectures $\mu^i \in \Delta\left(\Theta_{-i} \times A_{-i}\right)$.
- In the (rationalizability) analysis of *multistage* games with *complete* information, we considered CPSs $\mu^i \in \Delta^H\left(S_{-i}\right)$.
- In the (rationalizability) analysis of *multistage* games with *incomplete* information, we can use CPSs $\bar{\mu}^i \in \Delta^H\left(\Theta_{-i} \times S_{-i}\right)$, where—as before—$S_{-i} = \times_{h \in H} \mathcal{A}_{-i}(h)$ are the co-player's pure strategies (we write $\bar{\mu}^i$ to distinguish from systems of beliefs $\mu_i \in \left(\Delta\left(\Theta_{-i}\right)\right)^H$).
- We can *derive a personal assessment* $\left(\beta^i, \mu_i\right)$ from a CPS $\bar{\mu}^i$: $\forall \left(\theta_{-i}, h\right) \in \Theta_{-i} \times H$, $\forall a_{-i} \in \mathcal{A}_{-i}(h)$, $\mu_i\left(\theta_{-i} | h\right) = \bar{\mu}^i\left(\left\{\theta_{-i}\right\} \times S_{-i}(h) | h\right)$, if $\bar{\mu}^i\left(\left\{\theta_{-i}\right\} \times S_{-i}(h) | h\right) > 0$ then

$$\beta^i\left(a_{-i} | \theta_{-i}, h\right) = \frac{\bar{\mu}^i\left(\left\{\theta_{-i}\right\} \times S_{-i}\left(h, a_{-i}\right) | h\right)}{\bar{\mu}^i\left(\left\{\theta_{-i}\right\} \times S_{-i}(h) | h\right)}.$$

Example of Personal Assessment: Signaling

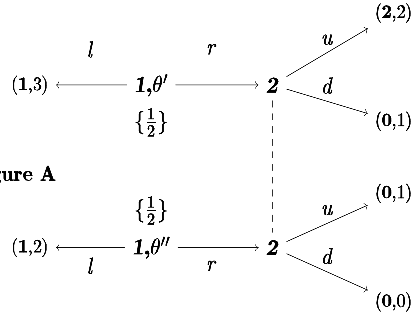

Figure A

- $\mu_2\left(\theta'|\varnothing\right) = p\left(\theta'\right) = 1/2,$
- $\beta^2\left(r|\theta'\right) > 0,\ \beta^2\left(r|\theta''\right) = 0,$
- $\mu_2\left(\theta'|r\right) = \frac{\beta^2(r|\theta')p(\theta')}{\beta^2(r|\theta')p(\theta')+\beta^2(r|\theta'')p(\theta'')} = 1.$

Example of Personal Assessment: Poker

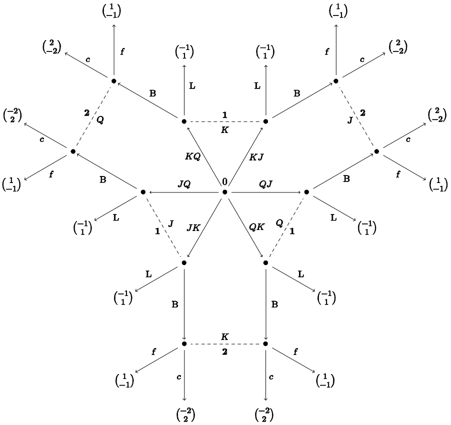

- $p(\theta) = 1/6$ for each $\theta \in \{JK, JQ, KQ, KJ, QJ, QK\}$.
- $\beta^2(\mathrm{B}|K_1) = \beta^2(\mathrm{B}|Q_1) = 1$, $\beta^2(\mathrm{B}|J_1) = q \in (0,1)$.
- $\mu_2(J_1|Q_2, \mathrm{B}) = \frac{1}{6}q / \left(\frac{1}{6}q + \frac{1}{6}\right) = q/(q+1)$.

Rehearsing Bayes Rule: Preliminaries

- For the first time, we are *truly* using Bayes rule!
- Abstract *finite* setting: **uncertainty space**

$$\Omega = \Theta \times X,$$

e.g., $\theta \in \Theta$ is the composition of a k-color n-balls urn, $x \in X$ is a sequence of colors of balls drawn from the urn.

- The **likelihood function** $\theta \mapsto P(\cdot|\theta) \in \Delta(X)$ gives the probability law for $x \in X$ for each parameter value $\theta \in \Theta$.
- As in *Bayesian statistics*, we posit $P(\cdot) \in \Delta(\Theta \times X)$ and interpret $P(x|\theta) = P(\{(\theta, x)\} \mid \{\theta\} \times X)$, where

$$P(\{(\theta, x)\} \mid \{\theta\} \times X) = \frac{P(\{(\theta, x)\})}{P(\{\theta\} \times X)} =: \frac{P(\theta, x)}{P(\theta)}$$

(recall: $P(E|F) = P(E \cap F|F) = P(E \cap F)/P(F)$ if $P(F) > 0$).

Rehearsing Bayes Theorem (or Rule)

- The joint probability of θ and x can be expressed as

$$P(x|\theta)P(\theta) = P(\theta, x) = P(\theta|x)P(x).$$

- The marginal, or **predictive** probability of x is

$$P(x) := P(\Theta \times \{x\}) = \sum_{\theta' \in \Theta} P(\theta', x) = \sum_{\theta' \in \Theta} P(x|\theta')P(\theta').$$

- Therefore (**Bayes theorem**) the "posterior" probability of any $\theta \in \Theta$ given any evidence $x \in X$ is

$$P(\theta|x) = \frac{P(\theta, x)}{P(x)} = \frac{P(x|\theta)P(\theta)}{\sum_{\theta' \in \Theta} P(x|\theta')P(\theta')}.$$

  - (1) Bayes rule only uses the likelihood function $\theta \mapsto P(\cdot|\theta)$ and the "prior" $\mathrm{marg}_\Theta P \in \Delta(\Theta)$.
  - (2) In statistics, we may assume $P(x) > 0$ for all x, *not in GT* where x is endogenous (actions).

Bayes Rule: Connection to Games

- Take the perspective of an *external observer* of a multistage game with incomplete information.
- The "probability law" is $\beta = (\beta(\cdot|\theta, h))_{\theta\in\Theta, h\in H} \in (\times_{h\in H}\Delta(\mathcal{A}(h)))^{\Theta}$ (obtained from a profile of type-dependent behavior strategies).
- There is an exogenous prior $p \in \Delta(\Theta)$; write $\beta(a|\theta, \varnothing) = \beta(a|\theta)$; for each $a \in \mathcal{A}(\varnothing)$ such that $\sum_{\theta'\in\Theta} \beta(a|\theta') p(\theta') > 0$ and each $\theta \in \Theta$,

$$\mu(\theta|a) = \frac{\beta(a|\theta) p(\theta)}{\sum_{\theta'\in\Theta} \beta(a|\theta') p(\theta')}.$$

- Let $\mu(\cdot|h) \in \Delta(\Theta)$ be given; for each $a \in \mathcal{A}(h)$ such that $\sum_{\theta'\in\Theta} \beta(a|\theta', h) \mu(\theta'|h) > 0$ and each $\theta \in \Theta$,

$$\mu(\theta|h, a) = \frac{\beta(a|\theta, h) \mu(\theta|h)}{\sum_{\theta'\in\Theta} \beta(a|\theta', h) \mu(\theta'|h)}.$$

References

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.