Introduction to Game Theory: Incomplete Information
Lecture 11, *Experimental Econ. & Psychology*

Pierpaolo Battigalli
Bocconi University

14 October 2020

Abstract
In games with **complete information**, all aspects of the interactive situation ("game"), including players' personal features such as their preferences and abilities, are *common knowledge* (everybody knows them, everybody knows that everybody knows, etc.). **Incomplete information** means that some aspects of the interactive situation are not commonly known. In this case, we want to describe what is commonly known, and to distinguish it from what is privately known by each player. We also want to describe players' beliefs about the exogenous and unknown aspects of the game, besides their beliefs about behavior, as well as beliefs about beliefs. Incomplete information is the norm, both in real life and in experiments. Incompleteness of information affects the way we analyze games in important ways.

Introduction

- Whether a solution or equilibrium concept is consistent with incomplete information is a matter of *interpretation*, we must look at the conceptual *motivations*:
- The iterated deletion of never-best replies analyzed in Lectures 9-10 represents the behavioral implications of rationality and common belief in rationality *under complete information* (common knowledge of the outcome function and utility functions).
- *Nash equilibrium* can be motivated as an "*obvious way to play* the game": e.g., what comes out of strategic reasoning (see above), or a "self-enforcing agreement". Also this makes sense *under the complete information assumption*.
- *Deductive interpretation of NE*: Unique outcome of the iterated deletion of never-best replies, see above.

- *Self-enforcing agreement interpretation of NE*: Again, we need complete information (or maybe something "close" to it) in order to make sense of this interpretation.
- Consider the following game with the assumption that *utility coincides with own monetary payoff*. Then (t, ℓ) is a Pareto Nash equilibrium. Is it self-enforcing? The *agreement* (t, ℓ) *is self-enforcing if there is common belief that there is no incentive to deviate from* (t, ℓ).

	ℓ	r
t	100,100	0,99
b	99,0	99,99

- Would Rowena (row player) play t if she is not sure of the payoff function of Colin (column player)? What if she is not sure that Colin is sure of her payoff function? What if...?

Static environments with Incomplete Information

- Rules of the game $\Rightarrow$ **outcome function** $\pi : A \to Y$.
- Each player $i \in I$ ranks (lotteries over) consequences according to (the expectation of) a vNM utility function $v_i : Y \to \mathbb{R}$.
- In environments with **incomplete information** there is *lack of common knowledge of* π and/or $(v_i)_{i \in I}$.
- Such situation can be described with *parametrized utility (of actions) functions*

$$u_i : \Theta \times A \to \mathbb{R},$$

with

  - $\theta \in \Theta$ parameter affecting payoffs and utilities,

$$\theta = (\theta_0, (\theta_i)_{i \in I}) \in \Theta = \Theta_0 \times (\times_{i \in I} \Theta_i)$$

  - $i \in I$ *knows* only θ_i =private info. of i about $(u_j)_{j \in I}$, the **type** of i.

- *Intuition*: it is common knowledge that $\theta \in \Theta$, Θ_i represents what is commonly deemed possible about i's attributes, the "larger" Θ_i the more uncertain are the other players about i's type.
- If Θ_i is a singleton ($i \in I$), that is, $\Theta_i = \left\{\bar{\theta}_i\right\}$, it means that what i knows is *common knowledge* (it is common knowledge that $\theta_i = \bar{\theta}_i$) and Θ_i can be neglected: indeed, $\Theta_0 \times \left(\times_{j\in I\setminus\{i\}}\Theta_j\right)$ and Θ have the same cardinality; hence, they are (intuitively) isomorphic.
- Θ_0 represents the *residual uncertainty* that would remain if the players could pool their private information.
- We often *focus* on the case where Θ_0 is a singleton: there is *no residual uncertainty* after pooling private information (in this case it is said that there is "**distributed knowledge**" of θ). Thus, we will often neglect Θ_0.

Private and interdependent values

We distinguish between the case of **private values**, where u_i depends only on θ_i, and **interdependent values**, where u_i may depend on the whole θ.

- **Private values:** *Common knowledge of* π but lack of common knowledge of $(v_i)_{i \in I}$
  - (common knowledge that) each i knows his vNM utility function v_i $\Rightarrow$ parametrized representation $v_i : \Theta_i \times Y \to \mathbb{R}$.
  - Note: $\{w_i \in \mathbb{R}^Y : \exists \theta_i \in \Theta_i, w_i = v_i(\theta_i, \cdot)\}$ is the set of utility functions that each $j \neq i$ thinks i might have $\Rightarrow$ get

$$u_i(\theta_i, a) = v_i(\theta_i, \pi(a))$$

  - Note: *under private values* we may assume w.l.o.g. that there is *distributed knowledge of* θ (Θ_0 singleton).
- Typically *in experiments* outcome (monetary payoffs) function $\pi : A \to Y$ (with $Y \subseteq \mathbb{R}^I$) is made common knowledge, but the preferences of others (v_{-i}) are unknown $\Rightarrow$ *private values*.

- **Interdependent values:** *lack of common knowledge* of π (maybe π depends on personal features such as "ability").
    - if *common knowledge of* $(v_i)_{i\in I}$ (just for simplicity) $\Rightarrow$ parametrized representation $\pi : \Theta \times A \to Y$ ($\{f \in Y^A : \exists\theta \in \Theta, f = \pi(\theta, \cdot)\}$ is the set of possible outcome functions) $\Rightarrow$ get

$$u_i(\theta, a) = v_i(\pi(\theta, a)).$$

    - More generally, if *neither* π *nor* $(v_i)_{i\in I}$ *is common knowledge*, each v_i is parametrized by θ_i and

$$u_i(\theta, a) = v_i(\theta_i, \pi(\theta, a)).$$

- *Experiments* sometimes create situations with an unknown outcome function, e.g., to study behavior in **common-value auctions**: the monetary value of the object on sale is the same for all subjects and it is unknown to them, subjects obtain private information correlated with such value.

Example

Cournot oligopoly *model (quantity setting): firm* $i = 1, ..., n$ *produces* $q_i \geq 0$ *units of homogeneous good*

- Inverse demand $P(Q) = [\bar{p} + \theta_0 - Q]^+$ (with $[x]^+ := \max\{0, x\}$, $Q = \sum_{i=1}^n q_i$*)*
- Cost function *of firm i:* $C_i(q_i, \theta_i) = \theta_i q_i$, $0 \leq q_i \leq \bar{q}$ ($\bar{q}$=common capacity),
- Common knowledge of selfish risk neutrality *and of sets* $\Theta_0, \Theta_1, ..., \Theta_n$
- Utility/payoff *of i:*

$$u_i(\theta_0, \theta_i, q_1, ..., q_n) = \left(\left[\bar{p} + \theta_0 - \sum_{j=1}^n q_j \right]^+ - \theta_i \right) q_i,$$

- *There are private values and distributed knowledge of* θ **if** *there is common knowledge of market demand (*Θ_0 *singleton)*

Example

Team production: Team agents $i = 1, ..., n$, i exerts effort $e_i \geq 0$

- Cost of effort *(in units of output)* $C_i(e_i, k_i) = k_i e_i^2$, $k_i \in K_i \subseteq \mathbb{R}_+$
- Production function*:* $y = \prod_{i=1}^{n} e_i^{p_i}$, $p_i \in P_i \subseteq \mathbb{R}_+$
- $\theta_i = (k_i, p_i) \in K_i \times P_i = \Theta_i$
- Common knowledge of (output-)risk neutrality *and of sets* $\Theta_i = K_i \times P_i$
- Utility *of i:* $u_i(k_1, p_1, ..., k_n, p_n, e_1, ..., e_n) = \frac{1}{n} \prod_{j=1}^{n} e_j^{p_j} - k_i e_i^2$
- Private values iff *sets $P_1, ..., P_n$ are singletons (productivities are common knowledge), otherwise* interdependent *values*

Static games with uncertainty

- We can represent (simultaneous) strategic interaction under *incomplete information* with the mathematical structure

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

  which is (informally) assumed to be common knowledge. This is called **game with uncertainty**.

- *Interpretation*: θ_0 affects the payoff/utility of somebody (if $\theta_0' \neq \theta_0''$, then $\exists i \in I, u_i(\theta_0', \cdot) \neq u_i(\theta_0'', \cdot)$). But part, or all, of i's private information θ_i *may be payoff irrelevant*. Yet even payoff-irrelevant information may be strategically relevant (e.g., θ_i may be the report to i by an art expert about the autenticity of a painting on auction).

- Games with uncertainty are sufficient to describe certain aspects of strategic thinking, specifically, *rationality and common belief in rationality*, by an extension of the algorithm of iteratively eliminating never-best replies.
- Write $\mathrm{B}_i(E)$ for "*i believes E*" (with prob. 1), and $\mathrm{B}(E) = \bigcap_{i \in I} \mathrm{B}_i(E)$ for "*everybody believes E*," R_i for "*i is rational*," $R = \bigcap R_i$ for "*everybody is rational*."
- What actions of i are consistent with R (rationality), $\mathrm{B}(R)$ (mutual belief in rationality), $\mathrm{B}(\mathrm{B}(R))$, $\mathrm{B}(\mathrm{B}(\mathrm{B}(R)))$ ... $R \cap \mathrm{CB}(R)$? [$\mathrm{CB}\,(E)$=common belief of E.]

Example

Assume: CK that utility=payoff. Possible payoff functions given by the following tables. Player 1 (Rowena) knows θ *while player 2* (Colin) does not *($\Theta \approx \Theta_1$)*

$\hat{G}^1$:

θ'	ℓ	r
t	4,0	2,1
b	3,1	1,0

θ''	ℓ	r
t	2,0	0,1
b	0,1	1,2

▶ $R_1 \Rightarrow [t \text{ if } \theta']$, *because* t dominates b given $\theta = \theta'$ *(recall, Row. knows θ)* $\Rightarrow (\theta', b)$ *is inconsistent with rationality (delete).*▶ $R_2 \cap B_2(R_1) \Rightarrow r$, *because* $u_2(\theta, x, \ell) < u_2(\theta, x, r)$ *for all* $(\theta, x) \neq (\theta', b)$ *(those consistent with R_1).*▶ $R_1 \cap B_1(R_2) \cap B_1(B_2(R_1)) \Rightarrow$ *Row. picks best reply to r given θ* $\Rightarrow [b \text{ if } \theta = \theta'']$.

Example

Assume: CK that utility=payoff. Players 1 and 2 receive an envelope. Envelope of i contains θ_i Euros, with $\theta_i = 1, ..., K$. Each player can offer to exchange (OE) by paying transaction cost $\varepsilon > 0$ (small). Exchange executed IFF both offer:

$\hat{G}^2$:

$a_i \backslash a_j$	OE	N
OE	$\theta_j - \varepsilon$	$\theta_i - \varepsilon$
N	θ_i	θ_i

Note: *A rational player i* offers to exchange only *if she assigns positive probability to event* $[\theta_j > \theta_i] \cap [a_j = OE]$.▶ $R_i \Rightarrow [a_i = N \text{ if } \theta_i = K]$ *because OE is dominated in this case.*▶ $R_i \cap \mathrm{B}_i(R_j) \Rightarrow [a_i = N \text{ if } \theta_i = K - 1]$ *because ...*▶ $R_i \cap \mathrm{B}_i(R_j) \cap \mathrm{B}_i(\mathrm{B}_j(R_i)) \Rightarrow [a_i = N \text{ if } \theta_i = K - 2]$ *because ...*▶ *It can be shown that:* $R \cap \mathrm{CB}(R) \Rightarrow (\forall \theta_i, a_i = N \text{ given } \theta_i)$ (*no-trade!*).

First-order beliefs (static games)

- To ease notation, disregard residual uncertainty Θ_0.
- The primitive uncertainty space of pl. i (person playing in role i) includes the unknown parameters (about which others have private information) besides the actions of others: $\Theta_{-i} \times A_{-i}$.
- Therefore, 1^{st}-**order beliefs about others** are probabilistic beliefs $\alpha_{i,-i} \in \Delta(\Theta_{-i} \times A_{-i})$.
- The **plan** of i yields her (typically certain) prediction about her own behavior: $\alpha_{i,i} \in \Delta(A_i)$.
- 1^{st}-**order beliefs** are subjective probability measures $\alpha_i = \alpha_{i,i} \times \alpha_{i,-i}$ (self *vs* others *independence*), space $\Delta_i^1 \subseteq \Delta(\Theta_{-i} \times A)$.
- Suppose there are *private values* (u_i depends only on θ_i), *why should i care about* θ_{-i}? For *strategic reasoning*: e.g.,
  - if my oligopolistic competitor has low marginal cost, her output is more likely to be high;
  - if my competitor in an auction values the object on sale a lot, she is likely to bid high.

Second-order beliefs (static games)

- If the beliefs of others matter for psychological reasons, $u_i : \Theta \times A \times \Delta^1_{-i} \to \mathbb{R}$, 2^{nd}-order beliefs are necessary to compute expected utility.
- Even if preferences over outcomes are belief-independent, 2^{nd}-order beliefs can be used, for example, to express the belief in the rationality of others (only triples $(\theta_j, a_j, \alpha_j)$ s.t. a_j is a best reply to $\alpha_{j,-j}$ for θ_j are possible).
- **2^{nd}-order beliefs about others** are subjective prob. measures $\beta_{i,-i} \in \Delta\left(\Theta_{-i} \times A_{-i} \times \Delta^1_{-i}\right)$ from which we can derive 1^{st}-order belief $\alpha_{i,-i}$ by marginalization as $\alpha_{i,-i}(\theta_{-i}, a_{-i}) = \beta_{i,-i}\left(\{\theta_{-i}\} \times \{a_{-i}\} \times \Delta^1_{-i}\right)$.
- Together with i's own plan we get the overall **2^{nd}-order belief** $\beta_i = \alpha_{i,i} \times \beta_{i,-i} \in \Delta\left(A_i \times \Theta_{-i} \times A_{-i} \times \Delta^1_{-i}\right)$ (self *vs* others independence). The space of 2^{nd}-order beliefs of i is Δ^2_i.

Best replies (static games)

- In games with uncertainty and asymmetric information i's best replies of i to her subjective beliefs depend on type θ_i. In traditional GT:

$$\bar{u}_i\left(\theta_i, a_i, \alpha_{i,-i}\right) = \sum_{\theta_{-i}, a_{-i}} u_i\left(\theta_i, \theta_{-i}, a_i, a_{-i}\right) \alpha_{i,-i}\left(\theta_{-i}, a_{-i}\right),$$

$$\begin{array}{rccl} BR_i: & \Theta_i \times \Delta_i^1 & \rightrightarrows & A_i, \\ & \left(\theta_i, \alpha_{i,-i}\right) & \mapsto & \arg\max_{a_i \in A_i} \bar{u}_i\left(\theta_i, a_i, \alpha_{i,-i}\right). \end{array}$$

- **Exercise:** Find the best-reply correspondences of the row and column players in game $\hat{G}^1$.
- In PGT (under *own-plan independence* of psychological utility):

$$\bar{u}_i\left(\theta_i, a_i, \beta_{i,-i}\right) = \mathbb{E}_{\beta_{i,-i}}\left(u_i\left(\theta_i, \cdot, a_i, \cdot, \cdot\right)\right),$$

$$\begin{array}{rccl} BR_i: & \Theta_i \times \Delta_{i,-i}^2 & \rightrightarrows & A_i, \\ & \left(\theta_i, \beta_{i,-i}\right) & \mapsto & \arg\max_{a_i \in A_i} \bar{u}_i\left(\theta_i, a_i, \beta_{i,-i}\right). \end{array}$$

Predictions, solutions

- By eliciting beliefs, BR correspondences may be used to derive some predictions. Yet, the standard approach of GT is to use solution concepts, and specifically some notion of equilibrium.
- Game theorists devised clever ways to extend the traditional equilibrium analysis to games with incomplete information. This requires a specification of the possible exogenous beliefs of players. This is an important part of game theory studied by several Nobel prize winners (the latest are Milgrom and Wilson). But it is not crucial to derive predictions in the experiments we are interested in.
- We can extend the algorithm of iterated deletion of never-best replies to allow for incomplete information:
  - 1. Eliminate all type-action pairs (θ_i, a_i) s.t. a_i is never a best reply for θ_i, that is, $a_i \notin BR_i(\theta_i, \alpha_{i,-i})$ for all $\alpha_{i,-i}$ (in PGT, $a_i \notin BR_i\left(\theta_i, \beta_{i,-i}\right)$ for all $\beta_{i,-i}$).
  - $n > 1$. Eliminate all the (remaining) pairs (θ_i, a_i) s.t. a_i is not a best reply for θ_i to beliefs consistent with steps $1, ..., n-1$.

Exercises

- **Exercise:** For game $\hat{G}^1$, perform the iterated deletion of pairs (θ_1, a_1) for the pl. 1 (row) and actions a_2 for pl. 2 (col.), *carefully explaining the steps*.
- **Exercise:** For game $\hat{G}^2$, try prove that the iterated deletion of pairs $(\theta_i, a_i) \in \{1, ...K\} \times \{OE, N\}$ eventually yields N for every i and θ_i. (If you know what is a proof by mathematical induction, you can try that; otherwise, try at least to provide an intuitive argument.)

Dynamic games with uncertainty

- In the lecture on dynamic games we learned that we can represent a dynamic game form with player set I (plus, possibly the chance player 0) specifying the tree $\bar{H}$ of histories (sequences of action profiles) allowed by the rules of the game, where $\bar{H}$ is partitioned into non-terminal histories (H) and terminal histories (Z).
- What was explained above for static games can be extended to dynamic games, letting the outcome depend on the terminal history: $\pi : Z \to Y$.
- If the outcome function π is not commonly known, let it depend on parameter vector θ, about which players may have differential knowledge: $\pi : \Theta \times Z \to Y$.
- If the probabilities of chance moves are not commonly known, let p_0 depend on θ as well: $p_0 = \left(p_0\left(\cdot|h, \theta\right)\right)_{h \in H_0, \theta \in \Theta}$.
- Extend the notions of best reply and rational planning to such incomplete-information environment.

References

- Battigalli, P. (2020): *Game Theory: Analysis of Strategic Thinking*. Typescript, Bocconi University. [Downloadable from webpage, optional.]
- Battigalli, P. (2020): *Mathematical Language and Game Theory*. Typescript, Bocconi University. [Downloadable from webpage, optional.]
- Battigalli, P., C. Corrao, and M. M. Dufwenberg (2019): "Incorporating Belief-Dependent Motivation in Games," *Journal of Economic Behavior & Organization*, **167**, 185-218. [Downloadable from webpage, optional.]
- Battigalli, P., and M. Dufwenberg (2020): "Belief-Dependent Motivations and Psychological Game Theory," *Journal of Economic Literature*, forthcoming.