Introduction to Game Theory

Lecture 2

Lorenzo Rocco

Galilean School - Università di Padova

March 2017

Games in Extensive Form

The most accurate description of a game is its extensive form.
In the extensive form we explicit:

- who the players are (including possibly a factious player - Nature)
- when each player plays
- what each player knows when he happens to play
- what actions are available to each player when he moves
- what are the payoffs for each possible outcome of the game.

Decision trees

- root : initial node that states who is the first mover
- nodes : points at which players take decisions
- branches : actions that can be taken at each node
- terminal nodes : outcomes of the game
- leafs : payoffs
- information sets: sets of nodes indistinguishable to a player

Examples of games in extensive form

- Extensive form the entry game
- Extensive form of the last hand of "briscola"
- Extensive form of the Prisoner's Dilemma

Information sets

- Information sets are collections of nodes where a given player has the right to take an action at a given moment of the game.
- Information sets represent what information a player has regarding the past history of the game when he is to move.
- All nodes belonging to an information set are indistinguishable from player's viewpoint: the player cannot tell, basing on his or her own information, what particular node he or she is at.
- The set of actions available is the same at each node of a given information (otherwise the player could infer which node he/she is playing at)

Strategies

A strategy s to player i is a function that associates an action (among those available) to each possible information set of a player

$$s_i : H_i \longrightarrow A_i$$

- A strategy is an action plan that tells the player what to do in any possible circumstance
- A strategy is an instruction manual. The set of strategies is a library.

Example: define the set of strategies of the sequential Prisoner's Dilemma
Example: define the set of strategies of the simultaneous Prisoner's Dilemma

Strategies, outcomes and payoffs

- Each strategy profile (a collection of the strategies taken by each player) determines the path followed along the game tree and the outcome of the game.
- Given that a payoff profile is associated to each outcome, a payoff profile is associated to each strategy profile

Note: the domain of the payoff function is the set of all possible strategy profiles

Extensive and normal form

Every game in extensive form could be transformed in normal form:

- a matrix with players strategies
- and the payoffs corresponding to each strategy profile

The normal form is a simplification. From the normal form, without further information, it is not possible to derive the original extensive form (to every normal form correspond multiple possible extensive forms).
Note: in the normal form it is as if players simultaneously take a strategy, an instruction manual, and ask a manager to actually play the game

The first concept of solution: equilibrium in dominant strategies.

The first and weakest solution concept is that of solution in dominant strategies

Dominant strategies

Definition

A strategy $s_i \in S_i$ is a strictly dominant strategy for player i if for all $s_i' \neq s_i$

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \text{ for all } s_{-i} \in S_{-i}$$

If a strategy is strictly better than all others, regardless of what other players do, then it must be the choice of a rational player.

Definition

A game has a unique solution in dominant strategy if there exists a strategy profile $s = (s_1, ..., s_n)$ such that $s_i \in S_i$ is a dominant strategy for all $i \in N$.

Unfortunately, strictly dominant strategies rarely exist.

Prisoner's dilemma

	Don't confess	*Confess*
Don't confess	$-2, -2$	$-10, -1$
Confess	$-1, -10$	$-5, -5$

The second concept of solution: iterated elimination of dominated strategies

The second concept of solution, which has a wider use, is that of iterated elimination of (strictly) dominated strategies

Dominated strategies

Much more common than strictly dominant strategies are (strictly) dominated strategies:

Definition

A strategy s_i is strictly dominated for player i if there exists a strategy s_i' such that

$$u_i(s_i, s_{-i}) < u_i(s_i', s_{-i}) \text{ for all } s_{-i}$$

Definition

A strategy s_i is weakly dominated for player i in game Γ_N if there exists a strategy s_i' such that

$$u_i(s_i, s_{-i}) \leq u_i(s_i', s_{-i}) \text{ for all } s_{-i}$$

Iterated Elimination of Strictly Dominated Strategies

- Strictly dominated strategies can be ruled out, basing on the principle of rationality
- After the elimination of some strictly dominated strategies, other strictly dominated strategies might emerge and can be further deleted

	L	C	R
U	$4,3$	$3,5$	$2,4$
M	$9,4$	$2,5$	$3,4$
D	$5,3$	$0,2$	$2,3$

- Iterated elimination grounds on the common knowledge of rationality: each iteration requires that CK of rationality goes one level deeper. If player i's drops his strictly dominated strategies, and knows that all other players are rational and know that he is rational, then player i might forecast that his opponents will drop their own dominated strategies which emerge after his elimination.

- Game: Each student chooses a natural between 1 and 100. The winner is the one whose bid is closest to $\frac{2}{3}$ of the average of all bids in the class.

How much do players believe in other players' rationality?

	A	B
a	$1000, 0$	$1, 1$
b	$-1000, 0$	$2, 1$

Iterated Elimination of Weakly Dominated Strategies

- Elimination of weakly dominated strategies cannot ground on rationality alone: there is at least one s_{-i} for which a weakly dominated strategy is equivalent to another undominated strategy. However weakly dominated strategies can be ruled out if players make mistakes...
- The iterated elimination of weakly dominated strategy is a concept even harder to justify: if players make mistakes it cannot be grounded of rationality; moreover the induced solution may depend on the order of elimination

	L	R
U	$5, 1$	$4, 0$
M	$6, 0$	$3, 1$
D	$6, 4$	$4, 4$

- $M \rightarrow R \rightarrow (DL)$
- $U \rightarrow L \rightarrow (DR)$

The third concept of solution: equilibrium in rationalizable strategies

Less stringent than dominates strategies are rationalizable strategies $->$ more widely usable concept

Rationalizable Strategies

- Iterated elimination of dominated strategy uses players' rationality and common knowledge of other players rationality
- The same idea can be pushed further and look at strategies that are never best response

Definition

A strategy σ_i is a *best response* for player i to the other players' strategies σ_{-i} if

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}) \text{ for all } \sigma'_i$$

Definition

A strategy σ_i is *never a best response* if there is no σ_{-i} for which σ_i is a best response.

Whatever the conjecture of player i about his rivals' play σ_{-i}, a never best response is always worse than another. Equivalently a never best response cannot be justified, whatever player's i conjectures about σ_{-i}.

Definition

The strategies that survive the iterated elimination of strategies that are never a best response are defined as player $i's$ rationalizable strategies.

The concept of rationalizable strategies reduces the set of "reasonable" alternatives that rational players might adopt: this fact reduces the set of possible outcomes in a game played by rational players. However many outcomes can be based on rationalizable strategies.

- Note: if a strategy is strictly dominated it is never a best response (but the reverse is not true in general).
- Note: therefore the set of rationalizable strategies must be smaller than the set of undominated strategies
- Note: the order of removal of never best response strategies does not affect the set of possible outcomes

Example 1

	b_1	b_2	b_3	b_4
a_1	$0, 7$	$2, 5$	$7, 0$	$0, 1$
a_2	$5, 2$	$3, 3$	$5, 2$	$0, 1$
a_3	$7, 0$	$2, 5$	$0, 7$	$0, 1$
a_4	$0, 0$	$0, -2$	$0, 0$	$10, -1$

What is the set of rationalizable strategies?
Note: there are no strictly dominated strategies

Example 2 - joint venture

Two firms agree on jointly perform a project. Each puts a certain level of effort s_i. The revenues from the joint venture are $4[s_1 + s_2 + bs_1 s_2]$ with $0 \leqslant b \leqslant \frac{1}{4}$. Revenues are equally split between the two partners so that each firm's payoff is $\pi_i = 2[s_1 + s_2 + bs_1 s_2] - s_i^2$.
What are firms' best responses?
What are firms' never best resposes?
What are firms' sets of rationalizable strategies that survive iterated deletion of never best responses?

The fourth solution concept: The Nash Equilibrium

The Nash Equilibrium has been defined by John Nash.
It is not a rule to find a solution (as it is the case for iterated elimination), but a definition which allows to check whether an outcome is an equilibrium.

Nash Equilibrium

The concept of Nash Equilibrium is the most widely used solution concept

Definition

A strategy profile s is a Nash equilibrium if for every i

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i'$$

- Each player adopts his best response to the strategy *actually* played by his opponents.
- While rationalization is based on rationality alone, the concept of NE requires that players *correctly forecast* what strategies will actually be played by their rivals.
- In other words NE requires not only that a strategy is optimal for at least some rational conjectures, but also that the conjectures are correct!!!

Nash Equilibrium - Examples

	S	C	D
a	$4, 4$	$5, 3$	$1, 2$
b	$3, 2$	$4, 5$	$3, 6$

	L	R
U	$3, 4$	$4, 6$
D	$2, 6$	$5, 4$

Nash Equilibrium

- Note: by definition, each strategy part of a Nash equilibrium is rationalizable. Therefore, outcomes in rationalizable strategies exists "at least as often" as Nash equilibria.
- Note: there might exist many Nash equilibria. In this case the assumption of corrected conjectures is very strong. Players are assumed to correctly forecast which equilibrium will be played, while GT is unable to predict which NE will actually be played.

	T	W
T	$2, 1$	$0, 0$
W	$0, 0$	$1, 2$

(Battle of the sexes)

Alternative definition of Nash Equilibrium

Definition

player i's best response set $b_i(s_{-i})$ is the set

$$b_i(s_{-i}) = \{s_i \in S_i : u(s_i, s_{-i}) \geqslant u(s_i', s_{-i}) \text{ for all } s_i'\}$$

defined for any given s_{-i}

Definition

The strategy profile $(s_1, ...s_I)$ is a Nash Equilibrium if and only if

$$s_i \in b_i(s_{-i}) \text{ for } i = 1, ..., I$$

Remark

- According to the previous definition, Nash equilibria are intersection points of (all) best response functions/correspondences.
- With two players, it might be convenient to determine players' best responses. The full set of Nash equilibria can be found by intersecting best responses.
- It can be done with both descrete and continous action/strategy sets
- Example: the intersection of best response strategies in the prisoners' dilemma.

Interpretation of Nash equilibrium concept

Is the concept of Nash Equilibrium "reasonable"?

1. NE as a consequence of rational inference (but rationality only defines rationalizable strategies!)
2. NE as a necessary condition when it is a unique obvious outcome of the game (rational players must understand it and expect that other players understand it as well)
3. NE as a focal point: certain outcomes are focal for cultural reasons, or have same natural appeal (if so players must expect that their rival adopt the corresponding strategies)
4. NE as a self-enforcing agreements (players engage in pre-play non-binding communication, agree on an outcome and expect that rivals subsequently play it - possible only if it is in their self-interest)
5. NE as a steady state of some dynamic adjustment process, where players follow some simple rules of thumb to forecast opponent's play (for instance Cournot tatonnement - players expect that rivals' adopt today the same action adopted yesterday).

Examples of games with complete information

- War of attrition
  - 2 players decide simultanously how much time to wait for a prize t_1, t_2
  - the first player to quit, lose and the other wins the prize
  - prize valuation are θ_1 and θ_2
  - player i payoff is thus

$$u_i = \begin{cases} \theta_i - t_{-i} & \text{if} \quad t_i > t_{-i} \\ -t_i & \text{if} \quad t_i < t_{-i} \\ \frac{1}{2}\theta_i - t_{-i} & \text{if} \quad t_i = t_{-i} \end{cases}$$

- Second-price sealed-bid auction
  - n players simoultanously put bids in sealed envelopes. The highest bidder gets the auctioned object and pays the second-highest bid
  - players valuations of the object are common knowledge and are $\theta_1 > \theta_2 > ... > \theta_n > 0$
  - in case of ties, the object is assigned to the bidder with lowest rank

Fact

This game has many Nash equilibria, including equilibria where the object is won by the n-th bidder. Most of these equilibria are however implausible as bidding own valuation is a weakly dominant strategy.

Exercise

Find all the Nash equilibria of a second-price sealed-bid auction with two bidders.
(Hint: Construct the players' best response functions)

- first-price sealed-bid action
  - n players simoultanously put bids in sealed envelopes. The highest bidder gets the auctioned object and pays his own bid
  - players valuations of the object are common knowledge and are $\theta_1 > \theta_2 > ... > \theta_n > 0$
  - in case of ties, the object is assigned to the bidder with lowest rank

Fact

In first-price sealed-bid actions, bidding own valuation is not a Nash equiblrium. However in all equilibria player 1 wins (consider any profile where someone else wins and gets a nonnegative payoff: player 1 may always deviate by rising his bid)

Remark on auctions

- "distinguished" Nash equilirium in second price auctions: $(b_1, .., b_n) = (v_1, v_2, .., v_n)$, player 1 wins and pay v_2
- "distinguished" Nash equilirium in second price auctions: $(b_1, .., b_n) = (v_2, v_2, b_3.., b_n)$, player 1 wins and pay v_2
- This is an example of the revenue equivalence theorem: all standard auctions produce the same outcome and the same revenues to the seller.

Games with incomplete information - The Png Settlement game

Players: Plaintiff and Defendant
Plaintiff claims that the defendant was negligent in providing safety equipment at a chemical plant, a charge which is true with probability $q = 0.13$ while the defendant is blameless with probability $1 - q$

Order of play
1) the plaintiff decides to Sue or to "Grumble"
2) the defendant Offers a settlement of $S = 0.15$ or Resist and goes to trial $S = 0$
3.1) if the defendant Resists, the plaintiff can choose either of Drop the case, with no legal costs for both players $P = 0$ and $D = 0$, or Try it, with legal costs of $P = 0.1$ and $D = 0.2$
3.2) if the defendant Offers a settlement, the plaintiff either agrees to Settle or Refuses and goes to trial
4) if the case goes to trial, the plaintiff wins damages $W = 1$ if the defendant is Liable and $W = 0$ otherwise

Remark

- The Png game is a game with asymmetric information. One player has private information of his own type which is instead unknown to his opponent.
- So far we have assumed that players knew all the relevant information about each other
- We consider now games in which players have incomplete information over the other players' preferences (e.g. two firms do not know each other cost).

Taxonomy

Games of

1. perfect information: history is known by each player
2. complete information: Nature does not move first, or her move is observed by both players
3. symmetric information: no player has information different from other players when he moves, or at the end nodes
4. certain information: Nature does not move after any player moves

The Harsanyi transformation

- The Png game is a game of incomplete (and thus imperfect), asymmetric, but certain information.
- By means of the *Harsanyi transformation*, we turn games where information regarding the rules or the type of the players is incomplete into games of imperfect information.
- We introduce a fictitious player, Nature, that chooses other player's type.
- Each player's preferences are determined by realization of a random variable. Each player observes the realization of her own random variable, while the probability distribution is assumed common knowledge.
- This transformation avoids to define player's believes and, mainly, believes about other players' believes.

Bayesian games (i.e. static games with incomplete information)

Let us focus on games in normal form where players have private information regarding their own type.
These games are known as Bayesian games

$$\Gamma_B = [I, \{A_i\}, \{u_i\}, \Theta, F(.)]$$

and are define by

- A_i the set of "actions" available to each player $i \in I$
- $u_i(a_i, a_{-i}, \theta_i)$ payoff functions which depend on player's own type θ_i
- $\theta_i \in \Theta_i$ is player i's "type", i.e. the realization of a random variable
- $F : \Theta \rightarrow [0, 1]$ is the joint probability distribution over $\Theta_1 \times ... \times \Theta_I$

Decision Rule

A strategy in a Bayesian game is called "decision rule"

Definition

A decision rule is a function $s_i : \Theta_i \longrightarrow A_i$ which prescribe an action for each player's type.

The set of all decisions rules is S_i

In a Bayesian game each player has a payoff function $u_i(s_i, s_{-i}, \theta_i)$ where $\theta_i \in \Theta$ is a random variable chosen by nature and observed by player i only.

Definition

Player $i's$ expected payoff is

$$E_\theta\left[u_i(s_1(\theta_1), ..., s_I(\theta_I), \theta_i)\right] = \tilde{u}_i\left(s_1\left(\cdot\right), ..., s_I\left(\cdot\right)\right)$$

where the expectation is taken over the type-profiles.

- Note: From each Bayesian game we can define a normal form game $\Gamma_N = [I, \{S_i\}, \{\widetilde{u}_i\}]$

Bayesian Nash Equilibrium (BNE)

Definition

A Bayesian Nash equilibrium in pure strategy is a strategy profile $(s_1(\cdot), ..., s_n(\cdot))$ that constitutes a Nash equilibrium of the game $\Gamma_N = [I, \{S_i\}, \{\widetilde{u}_i\}]$ such that for all $i \in I$

$$\tilde{u}_i(s_i(\cdot), s_{-i}(\cdot)) \geq \tilde{u}_i(s_i'(\cdot), s_{-i}(\cdot)) \text{ for all } s_i' \in S_i$$

Theorem

A strategy profile $(s_1(\cdot), \ldots, s_n(\cdot))$ *is a BNE if and only if for all* i *and for all* $\theta_i \in \Theta_i$ *occurring with positive probability*

$$E_{\theta_{-i}}\left[u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \,|\theta_i\,)\right] \geq E_{\theta_{-i}}\left[u_i(s_i'(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \,|\theta_i\,)\right]$$

for all $s_i' \in S_i$ *, where the expectation is taken over* θ_{-i} *and conditional to* θ_i*.*

Note: we can think of each type of player i as being a separate player who maximizes his payoff given his conditional probability distribution over the strategies of his rivals.

Application of Bayesian Nash equilibrium concept

- prisoner's dilemma with private information
- battle of sexes with private information
- Cournot competition
- Joint R&D

Png game in normal form

	OO	OR	RO	RR
G..	0;0,0	0;0,0	0;0,0	0;0,0
SST	**0.15;-0.15,-0.15**	-0.085;-0.15,-0.2	0.247;-1.2,-0.15	0.03;-1.2,-0.2
SSD	0.15;-0.15,-0.15	0.019;-0.15,0	0.130;0,-0.15	0;0,0
SRT	0.03;-1.2,-0.2	0.03;-1,2,-0.2	0.03;-1.2,-0.2	**0.03;-1.2,-0.2**
SRD	0.03;-1.2,-0.2	0.117;-1.2,0	-0.087;0,-0.2	0;0,0

row: Plaintiff; cols: Defendant
SST = (sue,agree to settle,try it)
SSD = (sue,agree to settle, drop it)
SRT = (sue,refuse settlement, try it)
SRD = (sue,refuse settlement, drop it)
OR = (offer a settlement, resist)
etc.