Introduction to Game Theory and Applications

Stefano MORETTI and Fioravante PATRONE

LAMSADE (CNRS), Paris Dauphine
and
DIPTEM, University of Genoa

Paris, Telecom ParisTech, 2010

Summary

Games with incomplete information

- An example
- Bayesian games
- Consistent case

Implementation, auctions

- Implementation: setting the problem
- Auctions
- Happy to pay taxes

ESS

The situation to be analyzed

There are two players. But *II* can be of two different "types".
We can summarize with a couple of tables (see Myerson, page 128):

$I \backslash II.1$	y_1	y_2
x_1	1 2	0 1
x_2	0 4	1 3

$I \backslash II.2$	y_1	y_2
x_1	1 3	0 4
x_2	0 1	1 2

Notice that y_1 is strongly dominant for player *II* of type 1, and y_2 is strongly dominant for player *II* of type 2. So, if we had common knowledge, it would be a very easy game to "solve".

But we assume that *I* does not know which "type" of player *II* he is facing.

Beliefs

The additional assumption that allows us to "close" the model is that each player has beliefs about the possible types of the other player(s) (in our small example, *II* knows which player he is facing, so we need only the beliefs of player *I* w.r.t. the two possible types of player *II*).

What is a "belief"? It is the assignment of a (subjective?) probability to each of the types.

So, in our simple model we need only to know the probability that *I* attaches to the two types of player *II*: let us say that p_1 is the probability he assigns to type 1 (otherwise stated, to *II*.1), and p_2 to *II*.2.

Assume, for example, that they are respectively 0.6 and 0.4.

Solving the example

Now we can easily "solve" the game.
Player I know that the strategy played by $II.1$ is y_1, and by $II.2$ is y_2. So, I knows that his opponent will play y_1 with probability 0.6 and y_2 with probability 0.4.

From now on things are easy (obvious): player I is able to evaluate his expected payoffs:
- playing x_1 his payoff is $0.6 \cdot 1 + 0.4 \cdot 0 = 0.6$
- playing x_2 his payoff is $0.6 \cdot 0 + 0.4 \cdot 1 = 0.4$

So, I chooses x_1, while $II.1$ (we already know that) chooses y_1 and $II.2$ chooses y_2.

Bayesian game

Generalizing the example, we have the following formal model to describe games with *incomplete information* (not anymore CK of everything!).

A Bayesian game is:
$G^b = ((A_I, A_{II}), (T_I, T_{II}), (p^I, p^{II}), (f, g))$
Where:
A_I is the set of actions available to each of the types of player I (similarly for A_{II})
T_I is the set of types of player I (similarly for T_{II})
And where beliefs of the different types of players are described by:
- $p^I : T_I \longrightarrow \Delta(T_{II})$
- $p^{II} : T_{II} \longrightarrow \Delta(T_I)$

and we have the payoffs of all of the types of both players described as follows:
- $f : A \times T_I \longrightarrow \mathbb{R}$
- $g : A \times T_{II} \longrightarrow \mathbb{R}$

(for sake of notational simplicity we could also say that $f, g : A \times T \longrightarrow \mathbb{R}$)
Of course, $A = A_I \times A_{II}$ and similarly $T = T_I \times T_{II}$

Bayesian game

What is a strategy for a player in a Bayesian game?

Given G^b, a strategy for player I is (simply) a map $s_I : T_I \longrightarrow A_I$, and similarly (as usual... really boring) for player II.

We could also give now the definition of Bayesian Nash equilibrium for a Bayesian game, but it will be deferred, when we shall have available another way of describing Bayesian games, at least in the "consistent" case.

Consistent case

Is it possible to find p on T (i.e.: $T = T_I \times T_{II}$) s.t.
- conditioning p to $t_I \in T_I$ we get precisely $p^I(t_I)$ (which is an element of $\Delta(T_{II})$
- conditioning p to $t_{II} \in T_{II}$...

?

Answer: not always.

Exercise: find an example for which it is not possible.
Exercise: generalize the idea of consistency to the case of a finite number of players.

Consistent case

Let's see an example:

$I \backslash II$	$II.1$	$II.2$	$II.3$
$I.1$	0.2	0.3	0.1
$I.2$	0.1	0.1	0.2

The matrix above describes a probability distribution p on $T = T_I \times T_{II}$. Clearly, knowing p, we get the probability on T_{II}, conditioned to the fact that player I is type 1 and similarly when he is of type 2:

$p(II.1|I.1) = \frac{0.2}{0.6}$ $\quad p(II.2|I.1) = \frac{0.3}{0.6}$ $\quad p(II.3|I.1) = \frac{0.1}{0.6}$

$p(II.1|I.2) = \frac{0.1}{0.4}$ $\quad p(II.2|I.2) = \frac{0.1}{0.4}$ $\quad p(II.3|I.2) = \frac{0.2}{0.4}$

The same thing can be done inverting the role of the players. So, from the knowledge of p on T we get:

$p^I : T_I \longrightarrow \Delta(T_{II})$

$p^{II} : T_{II} \longrightarrow \Delta(T_I)$

Consistent case: back to the classical setting

Just see the tree here:

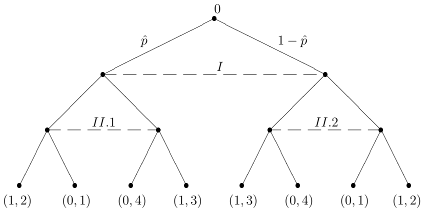

The main message is: nothing new. Also as far as the solution is concerned. Just write down the strategic form and compute Nash equilibrium.

Remark: no subgame! So subgame perfectness has no bite. Use instead weak Bayesian Nash equilibrium (or perfect Bayesian Nash equilibrium, or some other similar animal), or Sequential equilibrium (Kreps and Wilson, 1982). Have a look also at Mas-Colell, Green and Whinston (1995).

Suggested readings

- Chapter 10 of Osborne - Rubinstein.

- "The Theory of Implementation of Social Choice Rules", by Roberto Serrano (SIAM Review, 2004)

- "Mechanism Theory" (The Encyclopedia of Life Support Systems, 2001(?)), or "A crash course in implementation theory" (Social Choice and Welfare, 2001), by Matthew O. Jackson.

- Krishna, "Auction Theory", 2002.

- (in Italian): "implementazione_formalizzazione", heavily based on Osborne - Rubinstein, in my web page: `www.dri.diptem.unige.it`. Includes also "all" about King Solomon's dilemma.

- (in Italian): "aste_appunti_sbrigativi", heavily based on Krishna in my web page.

The implementation problem

- There is a set of alternatives A
- There is a set of individuals N
- There is a *social planner* (or the like) who is characterized by a social choice correspondence (or maybe function) $\phi : \mathcal{P} \rightrightarrows A$, where:
- - $\mathcal{P}$ is a *subset* of the set of all possible preference profiles $(\sqsupseteq_i)_{i \in N}$, where $\sqsupseteq_i$ is a total preorder (preferences of individual i on the set of alternatives A).

More to follow, but first an example

Example of implementation problem

I have two candies, one with lemon flavor, the other mint.
Two kids: $N = \{Giovanna, Paolo\}$.

- Assume that the alternatives available for me are (I cannot give two candies to one kid...):

- a_1: lemon to Giovanna, mint to Paolo
- a_2: lemon to Paolo, mint to Giovanna

- Assume that are possible only two profiles of preferences: $\sqsupseteq^1$ and $\sqsupseteq^2$, where:
- - $(\sqsupseteq^1_{Giovanna}, \sqsupseteq^1_{Paolo})$ means that Giovanna strictly prefers lemon to mint, and vice versa for Paolo.
- - $(\sqsupseteq^2_{Giovanna}, \sqsupseteq^2_{Paolo})$ means that Giovanna strictly prefers mint to lemon, and vice versa for Paolo.

- Assume that I like efficiency. This means that my social choice correspondence (a function, in this case) is:
$\phi(\sqsupseteq^1) = a_1, \phi(\sqsupseteq^2) = a_2$

A "small" aside: social choice theory

A nontrivial question is "from where" my social choice correspondence comes.

My choice could be the end-product of a rule that, given the preferences of the individuals, determines a "social preference". This rule is usually called a "social welfare function". And Ken Arrow key question was precisely about social welfare functions: is there a "reasonable" social welfare function? That is, one that satisfies reasonable conditions (axioms, properties)? His answer was: NO!

His result is known as an "impossibility" theorem...

Key question and key tool in implementation

Of course, if the social planner knows the preferences of individuals, things are easy. The problem is given by the fact hat he does not know these preferences profiles (or at least has just some partial information on them).

Examples:
- I want that you pay taxes depending on your (yearly) income.
- I ask for a recommendation letter as info for hiring a researcher.
- I have an indivisible good (a nice diamond necklace) that I want to allocate to the one who prefers it most.
- How can I convince these PhD students, sitting in front of me, to work hard on the course that I am giving?

I need to choose an appropriate game form (or mechanism) to implement my social choice correspondence.

Auctions: background

I am the (legitimate...) owner of an indivisible object z (a nice watch, a painting, a glove that belonged to Michael Jackson...).
There is a finite set N of individuals.
I want to allocate z to one of these individuals.
I assume that the individuals have preferences with a simple structure, so that they can be described just by a real number: their valuation of z.
I will call v_i the valuation of individual $i \in N$.

Essential feature: I do not know their valuation of z (I don't know the v_i's).

Auctions: which are the goals?

Which is my goal? It depends. I may want-to-apply/have different social choice correspondences. For example:

- allocate z to the (an) individual that values it most (my social choice correspondence is the "Pareto" social choice correspondence: allocate z to the individual that has the highest valuation of the object means that I am interested in an efficient result).

- allocate z selling it to one of the individuals in such a way to maximize my (expected?) revenue.

Game forms, mechanism design, implementation

I can allocate z in many different ways:

- choose $i \in N$ at random and just give z to him for free (gift).

- choose $i \in N$ (maybe via some random device, or in some other way) and bargain with him.

- posted price: I say I am willing to sell z at a price decided (arbitrarily?) by me.

- auction: I choose one of the many auction formats available to allocate (sell!) z.

- Etc.

Game forms, mechanism design, implementation

From each of the methods described before I get a different game form (quite trivial in the first case!), for which the set of players will be N (some could be very "quiet", as in the first two examples), the set of strategies and how is h will be determined by the specific game form that I choose, while the set of outcomes E is the set A.

The main goal of mechanism design is to find the "best" game form (or mechanism) for implementing my social choice correspondence (or to come as close as possible to it).

Example: auctions only

Four well known types of auctions:

- sealed bid, first price

- sealed bid, second price (or Vickrey auction, from Vickrey, 1961)

- English (or oral ascending) ["auction" from augere, "to increase", in Latin]

- Dutch

Four different game forms (mechanisms) that individuals are "obliged" (by you, social planner) to play to get z.

You choose the best game form to implement your social choice correspondence.

RET

RET: Revenue Equivalence Theorem (Vickrey, 1961):

Under reasonable assumptions (a restrictive one is that individuals involved are indifferent to risk), the expected revenue is the same for all of the four kinds of auctions that have been described.

Trust me

Consider the following game (strategic form to the right):

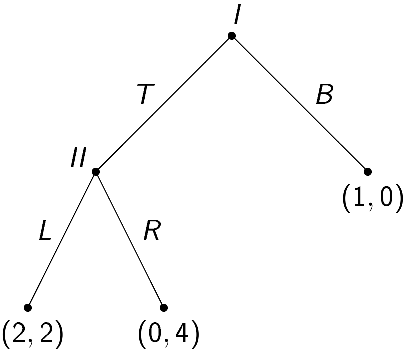

I \ II	L	R
T	2 2	0 4
B	1 0	1 0

Trust me, I will keep my promise...

Assume that we are in a non-cooperative setting, i.e., binding agreements are not allowed.
So, as a "solution" for this game we can consider the idea of equilibrium. This game has a unique Nash equilibrium (which is also subgame perfect): (B, R).

Problem is that the outcome is not efficient. Both players would prefer (T, L).

As an example of a situation described by such a game, consider the case in which player II has the opportunity of making a good investment, but he needs some money from player I. It is in the interest of both players to collaborate, and of course player II could try to convince I that he will "play" L. But this promise is not so much credible (as said, binding agreements are not available).

Where is the State?

Both players would like to have some external authority (the State?) that allows to make binding agreements (if you do not respect the contract, you are sent to jail!).

The problem is that the State has costs (jails, police, judges, etc.). But both player would be happy to pay some costs (only condition: that these costs are not too high) to be able to reach the "collaborative" outcome.

Let $2a$ be the costs of the State, and assume that these costs are equally charged to the players. So, the game is modified as in the next slide (where I am assuming, for sake of simplicity, that payoffs coincide with money).

Provided that the State is not too expensive!

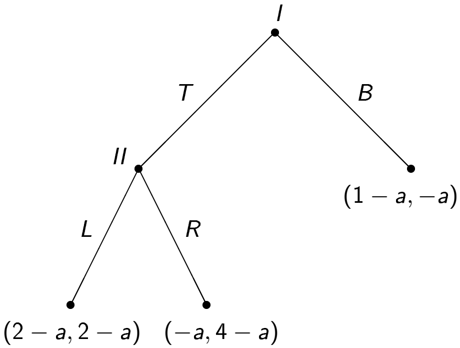

For the new game, the outcome from (T, L) is still better for both, compared with the original outcome from (B, R), provided that $a \leq 1$.

A possible mechanism: fines

One possibility to achieve efficiency is to modify the game form, obliging player *II* to pay a fine in case he does not respect the agreement, that is, in case he plays R. One should be sure that the fine is s.t. player *II* prefers the outcome from B to the outcome from L.

If this is feasible, players will be very happy to pay the taxes needed to have the State authority available, because the State will allow them to reach an outcome preferred by them. Look also at Hobbes' Leviathan...

This example is discussed (in Italian) with more details here: `http://dri.diptem.unige.it/altro_materiale/implementazione_legge_sui_prestiti.pdf`.

Evolutionary Stable Strategies

Another solution concept: ESS (Evolutionary Stable Strategy), from Maynard Smith.
Needed a symmetric game in *strategic form*.
$x^* \in X$ is an ESS if, for every $x \in X$ different from x^*, there exists $\bar{\varepsilon} > 0$ s.t. the following condition is true for all $\varepsilon > 0$ s.t. $\varepsilon < \bar{\varepsilon}$:

$$(1-\varepsilon)f(x^*, x^*) + \varepsilon f(x^*, x) > (1-\varepsilon)f(x, x^*) + \varepsilon f(x, x)$$

Equivalent formulation (to see that being a couple of ESS is a condition slightly stronger than Nash equilibrium):

$f(x^*, x^*) \geq f(x, x^*)$ for all $x \in X$

and

$[f(x^*, x^*) = f(x, x^*) \Rightarrow f(x^*, x) > f(x, x)]$ for all $x \in X$, $x \neq x^*$