Lecture 1: Decisions, Games, and Solution Concepts

ECON 2141/2142/8053

Semester 2, 2009

24 July 2009

- Decision theory – models of decision making for single individuals

  Examples: consumer choice, . . .

- Game theory – models of multiple interacting decision makers with potentially conflicting objectives

  Examples . . .

Models

What are models?

Why?

- identify the relevant aspects of a strategic situation that will help us predict a possible outcome
- derive general principles that apply to various strategic situations

Examples . . .

Models in game theory

- players
- available actions and order in which players move (sequential vs. simultaneous)
- what is the players' knowledge/information when they have to make a move (complete vs. incomplete, perfect vs. imperfect)
- what are the outcomes and how are they determined by the *joint* actions
- what are the players' preferences over outcomes/payoffs resulting from outcomes

⇒ Different classes of games and corresponding solution concepts

Rational choice theory

Decision making in microeconomics – apples vs. oranges

- a single agent
- a set of choices
- a fixed environment

How does the agent make a choice?

Imagine you are an economist and are trying to model the agent's decisions?

- Do you know the agent's utility function?
- Do you know his preferences over consumption bundles?
- What if you can observe choices made under different circumstances (incomes, prices, etc.)?

Revealed preference theory – If the agent's choice behavior reveals that the agent has "nice" preferences over his choice set, then the agent can be modeled as if he maximises some utility function. Assuming that the choice behavior is stable over time, such a utility function may be used to predict future choices.

choices $\Rightarrow$ preferences $\Rightarrow$ utility $\Rightarrow$ future choices

"Nice" preferences

Consider an agent who has to make a choice from a (finite) set A. A preference relation $\succeq$ over A is "nice" if

(i) it is *complete*, i.e., if for any two elements of A, a and b, either $a \succeq b$, $b \succeq a$, or both, and

(ii) it is *consistent* (or *transitive*), i.e., if for every $a, b, c \in A$, $a \succeq b$ and $b \succeq c$ imply that $a \succeq c$.

If both $a \succeq b$ and $b \succeq a$, then the decision maker is indifferent between a and b, denoted by $a \sim b$.

If $a \succeq b$ and $b \not\succeq a$, then the decision maker strictly prefers a to b, denoted by $a \succ b$.

Utilities

We say that a preference relation $\succeq$ is represented by a (real-valued) utility function u if

$$u(a) \geq u(b) \text{ if and only if } a \succeq b.$$

If an agent's preferences can be represented by a utility function, then his choices can be modeled as if he is maximising this utility function.

Every "nice" preference relation can be represented by a (non-unique) utility function.

Example ...

We will assume that all players in a game have preferences over outcomes that can be represented by a utility (payoff) function.

Solution concepts

Interpretations:

(1) Descriptive (positive): How do players play in a given game?

(2) Prescriptive (normative): How "should" players play in a given game?

(3) Theoretical: What outcomes are consistent with certain assumptions regarding "reasonable" or "rational" behavior on the part of the players?

$\Rightarrow$ (1) and (2)?

Example: All-pay auction

- All players can make sequential bids for a payout of \$10.
- Each new bid must be at least \$1 higher than the last bid.
- The game ends when no additional player makes any bid.
- The highest bidder gets the \$10 and pays his last bid.
- *All other players must also pay their last bid!*

What should you do given the other players' actions?

What are your beliefs about the other players' (future) actions?

...

One more example: Guess $\frac{2}{3}$ of the average ... (Osborne 34.1)

Strategic games

Definition

A *strategic game (with ordinal preferences)* consists of

(1) a (finite) set of players, with individual players denoted by $i, j, \ldots,$

(2) for each player i, a set of actions A_i, with a particular player i action denoted by a_i,

(3) for each player i, preferences over the set of action profiles, which can be represented by a utility function u_i.

Note:

- An action profile specifies a particular action for each player. For example, in a two-player game, an action profile a must specify an action a_1 for player 1 and an action a_2 for player 2, so $a = (a_1, a_2)$.

- In a game with an arbitrary number of players we often denote the actions of all players other than a given player i by a_{-i}. We can then write down an action profile a from the point of view of player i as $a = (a_i, a_{-i})$.
- Since outcomes in a game are determined by the actions of all players, i.e., by a given action profile, each player i's utility is a function of (a_i, a_{-i}), so we can write $u_i(a_i, a_{-i})$.

Assumptions:

(i) Players choose actions simultaneously, without knowing the other players' actions.

(ii) The description of the game is known to all players. (Common knowledge!)

(iii) All players are rational (i.e., they have "nice" preferences and choose actions that yield most preferred outcomes).

(iv) Players care only about instantaneous payoffs and not about future effects of current actions. (No inter-temporal strategic links.)

Examples

	Quiet	*Fink*
Quiet	2, 2	−1, 3
Fink	3, −1	0, 0

Prisoner's dilemma

	B	*S*
B	2, 1	0, 0
S	0, 0	1, 2

Battle of the sexes

	Stag	*Hare*
Stag	2, 2	0, 1
Hare	1, 0	1, 1

Stag hunt

	H	*T*
H	1, −1	−1, 1
T	−1, 1	1, −1

Matching pennies

Nash Equilibrium

Idea: No player should have an incentive to change his action given the actions of the other players. (steady state, stability property)

Definition

An action profile a^* is a *Nash Equilbrium* (NE) if for all players i and all $a_i \in A_i$,

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*).$$

But players choose actions simultaneously, so how does a player know what the other players will do?

Requirements for Nash Equilibrium:

(1) No player has an incentive to deviate given (his beliefs about) the other players' actions; and

(2) all players have correct beliefs about their opponents' actions!

Why should we assume that players know the actions of their opponents?

(i) previous experience playing similar games (learning)
(ii) pre-play communication (a NE is a self-enforcing agreement)
(iii) social norms and conventions
(iv) focal points ("obvious" organising principles)
(v) evolution (only "high payoff" players survive in the long run)

How to find Nash equilibria?

Consider every action profile and check whether any player has a profitable deviation. ($\rightarrow$ 30.1)

Best response functions

Find the "best response" of a player for every fixed action profile of his opponents.

	L	C	R
T	$1, 5$	$2, 5$	$2, 1$
M	$2, 9$	$5, 2$	$3, 0$
B	$0, 1$	$4, 3$	$4, 2$

Definition

Player i's *best response function* is defined as

$$B_i(a_{-i}) := \{a_i \in A_i \mid u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i}) \text{ for all } a'_i \in A_i\}.$$

The best response function is set-valued!

Proposition

An action profile a^* is a NE iff *(if and only if) $a^*_i \in B_i(a^*_{-i})$ for all i.*

Tutorial problems for next week: Osborne 30.1, 42.2, 44.1