

1. Basics of Noncooperative Games

Introduction

Microeconomics studies the behavior of individual economic agents and their interactions. Game theory plays a central role in modeling the interactions between economic agents. The type of games we shall study are noncooperative games, in the sense that each economic agent in the game, who acts in her self interest, is the unit of analysis. (There is another branch of game theory, cooperative game theory, that treats groups or subgroups of economic agents as the unit of analysis and assumes they can achieve certain outcomes among themselves through binding cooperative agreements.) When we say game theory, we usually mean noncooperative game theory. In the last 20 years or so, developments in game theory and its applications in economics represent some of the most important advances in microeconomics.

What is a game? A game is a decision situation with multiple decision makers where each person's welfare depends on her own as well as other individuals' actions. That is, a game is a decision situation with strategic interactions among all decision makers. To describe a game, we need to know four things:

1. The players: who is involved?
2. The rules: how do players move? what do they know when they move? what can they do?
3. The outcomes: for each possible set of actions by the players, what is the outcome of the game?
4. The payoffs: what are the players' preferences (utility functions) over the outcomes.

Let's see two simple examples:

The Matching Pennies Game. There are two players, denoted 1 and 2. Each player simultaneously puts a penny down, either heads up or down. If the two pennies match, player 1 pays \$1 to player 2; otherwise player 2 pays \$1 to player 1. The amount of money a player wins or loses is his payoff.

The Capacity Game. Suppose that a firm, A , needs to decide whether to build high capacity (denoted as H) or low capacity (denoted as L) in an industry. There is another firm, B , needs to decide whether to enter the market (denoted as E) or stay out the market (denoted as NE) after observing A 's capacity level. If A chooses H and B chooses E , A will have 1 unit of profit and B will have two units of losses (1,-2); If A chooses H and B chooses NE , A will have 5 units of profit and B will have 0 unit of profit (5,0); if A chooses L and B chooses E , A will have 2 units of profit and B will have two units of profits (2,2); if A chooses L and B chooses NE , A will have 8 units of profit and B will have 0 unit of profit (8,0). The amount of profit or losses a player receives is his payoff.

The matching pennies game is called a zero-sum game because what one player wins is exactly what the other player loses.

Although a game can be fully described as in the above examples, it is useful to represent a game in some particular ways.

The Extensive Form Representation of a Game

One way to represent a game is called its extensive form. Let's first see the extensive form of the Capacity Game. The game starts at an initial decision node (represented by an open circle), where player A makes his move, deciding whether to have H capacity or L capacity. Each of the possible choices by player A is represented by a branch from this initial decision node. At the end of each branch is another decision node (represented by a solid dot), at which player B can choose between two actions, E or NE , after seeing A 's choices. The initial decision node is referred to as player

A 's decision node, and the latter two as player B 's decision nodes. After B 's move, we reach the end of the game, represented by terminal nodes. At each terminal node, we list the players' payoffs arising from the sequence of moves leading to that terminal node. This representation is also called a game tree.

A game is one of perfect information if each player, when it is her turn to move, knows all the previous moves of other players in the game. The Capacity Game is a game of perfect information.

A subset of a player's decision nodes is called an information set of this player if, when the player reaches one of the decision nodes in the subset, she cannot tell which of these nodes she is actually at. An information set may contain only one decision node, in which case the information set is a singleton. In the Capacity Game above, there are three information sets that are all singletons.

Now let's see the extensive form of the Matching pennies game. When player 2 moves, she does not know which decision node in the circle she is at, since she has not observed player 1's move. Thus player 2's two decision nodes are in an information set. This is how a game with players moving simultaneously is represented in the extensive form.

There are two restrictions on an information set. First, at every node within a given information set, a player must have the same set of possible actions. Second, there is perfect recall, which means that a player does not forget what she once knew, including her own actions. Two games not satisfying perfect recall.

Definition: A game is one of perfect information if each information set contains a single decision node. Otherwise, it is a game of imperfect information.

In some situations, the outcome of a game may depend on some random factors. This can be captured in the extensive form by including random moves of nature. We can illustrate this through a variation of the matching pennies game. Suppose a variation of the matching pennies game has the following form. Instead of simultaneous moves, the two players make their choices of heads or tails sequentially, and the two players first flip a coin to see who moves first.

It is generally assumed in game theory that the structure of the game is common knowledge.

Formally, a game represented in extensive form is specified by the collection:

$$\Gamma_E = \{\chi, A, I, p(\cdot), \alpha(\cdot), \mathcal{I}, H(\cdot), \iota(\cdot), \rho(\cdot), u\}$$

where: χ is a finite set of nodes; A a finite set of possible actions; I is a finite set of players $\{1, 2, \dots, I\}$; $p(\cdot)$ is a function $p : \chi \rightarrow \{\chi \cup \emptyset\}$ that specifies a single immediate predecessor of each node x , and $p(\cdot)$ is nonempty except for the initial node x_0 ; $\alpha(\cdot)$ is a function $\alpha : \chi \setminus \{x_0\} \rightarrow A$ that gives the action leading to any noninitial node x from its immediate predecessor $p(x)$; \mathcal{I} is a collection of information sets and $H : \chi \rightarrow \mathcal{I}$ assigns each decision node x to an information set $H(x) \in \mathcal{I}$; $\iota(\cdot)$ assigns each information set in \mathcal{I} to the player who moves at the decision nodes in the set; $\rho(\cdot)$ assigns probabilities to actions at information sets where nature moves; u is a collection of payoff functions assigning utilities to the players for each terminal node that can be reached.

Strategies and the Normal Form Representation of a Game

A player's strategy is a complete contingent plan that specifies how the player will act in every possible distinguishable circumstance in which she might be called upon

to move. Since each possible distinguishable circumstance that a player might be called upon to move is represented by an information set for the player, a player's strategy is a specification of how she plans to move at each of her information sets.

Definition. Let \mathcal{I}_i denote the collection of player i 's information sets, A the set of possible actions in the game, and $C(H) \subset A$ the set of actions possible at information set H . A strategy for player i is a function $s_i : \mathcal{I}_i \rightarrow A$ such that $s_i(H) \in C(H)$ for all $H \in \mathcal{I}_i$.

In words, a strategy of player i is a function s_i that maps from the collection of player i 's information sets to the action set such that at each information set player i specifies an action among the possible actions at this information set.

Example. Strategies in the Capacity Game:

Player A has two possible strategies: H or L.

Player B has two information sets, and a strategy of his needs to specify what he will do at each of the two information sets. Thus he has four possible strategies:

Play E if A plays H; Play E if A plays L. (E if H; E if L)

Play E if A plays H; Play NE if A plays L.

Play NE if A plays H; Play E if A plays L.

Play NE if A plays H; Play NE if A plays L.

Example. Strategies in the Matching Pennies Game.

Each player has two possible strategies: play heads (H) or tails (L).

It is very important to remember that a strategy is a complete specification of what a player will do at each of the entire collection of her information sets, even an information set is not reached in the actual play by this player or other players' strategies. Consider a modified Capacity Game, where if player B has chosen to enter the market (E), Player A can then choose to fight (F) or not to fight (NF).

Denote the A's decision node after it has chosen H and B has chosen E by x_1 , and A's decision node after it has chosen L and B has chosen E by x_2 . Then A's strategies are:

H, then F if at x_1 , and F if at x_2 .

H, then F if at x_1 , and NF if at x_2 .

H, then NF if at x_1 , and F if at x_2 .

H, then NF if at x_1 , and NF if at x_2 .

L, then F if at x_1 , and F if at x_2 .

L, then F if at x_1 , and NF if at x_2 .

L, then NF if at x_1 , and F if at x_2 .

L, then NF if at x_1 , and NF if at x_2 .

A profile of strategies in an I -player game can be denoted by a vector $s = (s_1, s_2, \dots, s_I)$, where s_i is the strategy of player i . Sometimes we also write s as (s_i, s_{-i}) , where s_{-i} is the $(I - 1)$ vector of strategies for players other than i .

Once we know how to write all the strategies for each player in a game, and the associated payoffs for each strategy profiles, we can represent a game in a form known as the normal form or the strategic form.

Example. The Normal form of the Matching Pennies Game.

		Player 2	
		H	T
	H	-1,1	1,-1
Player 1			
	T	1,-1	-1,1

Example. The Normal Form of the capacity Game.

		B			
		E if H; E if L	E if H; NE if L	NE if H; E if L	NE if H; NE if L
	H	1, -2	1, -2	5, 0	5, 0
A					
	L	2, 2	8, 0	2, 2	8, 0

Definition. For a game with I players, its normal form representation Γ_N specifies for each player i a set of strategies S_i with $(s_i \in S_i)$ and a payoff function $u_i(s_1, \dots, s_I)$ giving the v.N-M utility levels associated with the (possibly random) outcomes arising from (s_1, \dots, s_I) . It is written as $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$.

Relations between the extensive form representation of a game and its normal form: For any extensive form representation of a game, there is a unique normal form representation. But a normal form representation of a game could correspond to more than one extensive forms. We can think of the normal form as a more condensed way to represent a game than the extensive form, and the normal form generally omits some of the details present in the extensive form. But if in a game players all move at the same time (i.e., a simultaneous-move game), then the extensive form and the normal form representations of a game are always equivalent.

Randomized Choices

So far we have assumed that players make their choices with certainty. We call the deterministic strategies players so choose pure strategies. In many situations,

however, a player might want to randomize her choices. That is, she might choose randomly among her pure strategies. In the matching pennies game, for instance, a player may choose H or T each with probability 1/2. This is called a mixed strategy. Formally,

Definition. Given player i 's pure strategy set S_i , a mixed strategy for player i , $\sigma_i : S_i \rightarrow [0, 1]$, assigns to each pure strategy s_i a probability $\sigma_i(s_i) \geq 0$ that it will be played, where $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.

Suppose that player i has M pure strategies in $S_i = (s_{1i}, \dots, s_{Mi})$. The set of player i 's mixed strategies is simplex

$$\Delta(S_i) = \{(\sigma_{1i}, \dots, \sigma_{Mi}) \in R^M, \sigma_{mi} \geq 0 \text{ for all } m = 1, \dots, M \text{ and } \sum_{m=1}^M \sigma_{mi} = 1\}.$$

This simplex is called the mixed extension of S_i . Note that a pure strategy can be view as a special case of a mixed strategy where the probability distribution over the elements of S_i is degenerate.

Recall that player i 's payoff under a pure strategy profile $s = (s_1, \dots, s_I)$ is $u_i(s)$. Given a profile of mixed strategy $\sigma = (\sigma_1, \dots, \sigma_I)$, player i 's payoff (which is the von Neumann-Morgenstern expected utility) is $\sum_{s \in S} [\sigma_1(s_1)\sigma_2(s_2)\dots\sigma_I(s_I)]u_i(s)$, which we shall denote as $u_i(\sigma)$.

When mixed strategies are allowed, the normal form representation of a game is denoted as $\Gamma_N(I, \{\Delta(S_i)\}, \{u_i(\cdot)\})$.

If we use the extensive form representation of a game, player i can choose to randomize separately over the possible actions at each of her information sets $H \subset \mathcal{I}_i$. This way of randomizing is called a behavior strategy.

Definition. Given an extensive form game Γ_E , a behavior strategy for player i specifies, for every information set $H \subset \mathcal{I}_i$, and every action $a \in C(H)$, a probability $\lambda_i(a, H) \geq 0$, with $\sum_{a \in C(H)} \lambda_i(a, H) = 1$ for all $H \subset \mathcal{I}_i$.

For games with perfect recall, these two types of randomization are equivalent.

That is, for any behavior strategy of player i , there is a mixed strategy for that player that yields exactly the same distribution over outcomes for any strategies that might be played by i 's rivals, and vice versa. For convenience, we usually use behavior strategies when analyzing extensive form games, and use mixed strategies when analyzing normal form games. We often call both types of randomized strategies mixed strategies.