Appendix F
Game Theory

This appendix provides some elementary facts about game theory and equilibrium concepts and should serve as a refresher for those readers with some background in game theory. However, for a proper and more complete explanation of the theory, the reader should refer to Fudenberg and Tirole [195], Myerson [399] or Mas-Colell, Whinston, and Green [365] (from which the material in this appendix is obtained). Also, Tirole [513] provides a User's Manual on Game Theory.

The Normal Form of Games

A game consists of a set of *players,* the *actions* that they can take (or in other words, the rules of the game), and the *information* that each player possesses at the time he takes his action. For each possible set of actions, the game defines a set of *outcomes* and *payoffs* for each player (such as how much profit or utility each player gets).

For instance, in the Bertrand pricing game (Section 8.4.1.4), there are n players (firms in the oligopoly). Their action space is the prices they set. Each possesses information that all the demand goes to the lowest-priced firm, and all have the same marginal costs. The outcome is that the demand goes to the lowest-priced firm. The payoffs are the revenues minus costs (profits).

Formally, let there be n players, let $\mathcal{H}_i$ be the collection of player i's information sets and $C(H) \subset \mathcal{A}$ be the set of actions possible for player i with information set *H.*

A *(pure) strategy* for player i is a function $s_i : \mathcal{H}_i \to \mathcal{A}$,—that is, the player has a mapping from each possible information set to a unique action. Moreover, the actions have to be feasible, so we assume that the strategy map is such that $s_i(H) \in C(H)$ for all $H \in \mathcal{H}_i$. Each player, given a set of pure strategies, can also randomize over these strategies (his strategy is to choose one of his pure strategies with a certain probability). This creates what are called *mixed strategies.*

A game's actions, outcomes, and payoffs can be defined by an extensive form or a normal form. Here we concentrate on the normal form. The normal form of the game is a specification of a set of possible strategies S_i for player i, and a payoff function $u_i(s_1, \ldots, s_N)$ if each player plays strategy $s_i \in S_i, i = 1, \ldots, N$. The game Γ is defined as the triple $\Gamma = [N, \{S_i\}, \{u_i(\cdot)\}]$.

For example, in the Bertrand pricing game with two players, the strategy space for player i is $(0, \infty)$, and the payoff if player 1 plays $p_1 \in (0, \infty)$ is $p_1 d_1(p_1, p_2)$, where

$$d_1(p_1, p_2) = \begin{cases} d(p_1) & \text{if} \quad p_1 < p_2 \\ d(p_1)/2 & \text{if} \quad p_1 = p_2 \\ 0 & \text{if} \quad p_1 > p_2 \end{cases} \tag{F.1}$$

and $d(\cdot)$ is the market-demand function.

Simultaneous Move Static Games

In the simultaneous move static game, all players move exactly once and make their moves simultaneously. Hence, no player knows what the other players' moves are going to be, nor do they have any information on past moves of their opponents (as it is a one-move game).

These are rather restrictive assumptions. Nevertheless, such games are applicable in some situations (for instance, a sealed-bid auction), and they serve as the basis for the study of more complicated repeated games.

Game theory is concerned with predicting the outcomes of a game assuming the players are rational (utility-maximizing players). To this end, we define the concept of *equilibrium,* essentially a prediction of the possible outcomes of the game. There are many equilibrium concepts, depending on the nature of the information, and the assumptions on players' behavior.

We assume that players have *complete information* about the game. Each player knows the strategy sets, utility functions, and any other relevant parameters for all other players, and they also know that all the other players are rational and, like themselves, have complete information.

Dominant Strategies

A strategy $s_i \in S_i$ is a *dominant strategy* for player i if his payoff from playing s_i is no less than that from playing any other of his strategies, for all possible strategies of the other players.

Formally, let the the vector $\mathbf{s}_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_N)$ represent strategies of all players other than i, and S_{-i}, the set of all vectors of all possible strategies of all the players other than i.

Then, for a game $\Gamma = [N, \{S_i\}, \{u_i(\cdot)\}]$, $s_i \in S_i$ is a dominant strategy if for all $s_i' \neq s_i$,

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}), \quad \forall s_{-i} \in S_{-i}.$$

If rational player i has a strictly dominant strategy, it is reasonable to predict he would always play that strategy. There are very few games, however, where such dominant strategies exist.

Nash Equilibrium

Perhaps the most important and widely accepted notion on the outcome of games with rational players is the Nash equilibrium.

Nash equilibrium, definition A strategy vector $\mathbf{s}$ is a *(pure strategy) Nash equilibrium* for the game $\Gamma = [N, \{S_i\}, \{u_i(\cdot)\}]$ if for every player $i = 1, \ldots, N$, given the strategies of the other players $\mathbf{s}_{-i}$, his strategy s_i is optimal, that is,

$$u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i}), \quad \forall s_i' \in S_i, \; i = 1, \ldots, N.$$

If we allow players to randomize over their strategies, then a vector of mixed strategies, $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_N)$ is a mixed-strategy Nash equilibrium if for every player i, given a profile of mixed strategies of the other players $\boldsymbol{\sigma}_{-i}$, player i's mixed strategy σ_i is optimal. A game can have no Nash equilibrium, a unique Nash equilibrium, or many equilibria (pure-strategy or mixed).

Here are two basic results on the existence of pure and mixed-strategy Nash equilibria.

PROPOSITION F.16 *Every game with finite strategy sets for all the players has a mixed-strategy Nash equilibrium.*

PROPOSITION F.17 *If the strategy sets for the players S_i are nonempty, convex and compact subsets of $\Re^m$, $u_i(\mathbf{s})$ is continuous in* $\mathbf{s}$ *and quasiconcave in s_i for all i, then the game has a pure-strategy Nash equilibrium.*

However, in either case, there is no guarantee that the equilibrium is unique.

Bayesian Nash Equilibrium

Games with *incomplete information* model situations where the players do not know with certainty what the other players' strategy sets, parameters, and utility functions are. Each forms a probabilistic view of the other players' private information (akin to a Bayesian prior; this probabilistic view may be updated in a repeated game as the game reveals more information to the players).

The model is as follows: Player i's payoff function is now given by $u_i(\mathbf{s}, \theta_i)$, where $\theta_i \in \Theta_i$ is a random variable whose realization is observed only by player i. Let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ and $\boldsymbol{\Theta} = \Theta_1 \times \Theta_2 \times \cdots \times \Theta_N$. However, the joint probability distribution of $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, $F(\boldsymbol{\theta})$ is common knowledge among the players. The Bayesian game is then $\Gamma = [N, \{S_i\}, \{u_i(\cdot)\}, \boldsymbol{\Theta}, F(\cdot)]$.

A pure-strategy for player i is in this case a decision rule $s_i(\theta_i)$. His strategy is a function of the realization of his θ_i. Given a vector of pure strategies for all the players $(s_1(), \ldots, s_N(\cdot)$, player i's payoff is given by the expectation over the θ's:

$$\tilde{u}_i(s_i(\cdot), \mathbf{s}_{-i}(\cdot)) = E_{\boldsymbol{\theta}}[u_i(s_i(\cdot), \mathbf{s}_{-i}(\cdot))].$$

The extension to the Nash equilibrium concept is then as follows. A *(pure-strategy) Bayesian Nash equilibrium* is a vector of decision rules $\mathbf{s}(\cdot) = (s_1(\cdot), \ldots, s_N(\cdot)$ such that

$$\tilde{u}_i(s_i(\cdot), \mathbf{s}_{-i}(\cdot)) \geq \tilde{u}_i(s_i'(\cdot), \mathbf{s}_{-i}(\cdot)).$$

Repeated Games

Finite repeated games are one-shot games that are repeated over a number of periods. At the beginning of each period, firms are aware of the others' past moves and make their decisions simultaneously and noncooperatively for that period.

For example, a repeated Bertrand game would have firms setting prices simultaneously at the beginning of each period, and a repeated Cournot game would have firms deciding how much to produce at the beginning of each period. For instance, when all the firms in the market post their prices on a centralized industrywide reservation system every day, the repeated game's period is one day, and at the beginning of each day firms set prices simultaneously without knowing how the other firms will choose their prices that day.

Models where moves don't occur simultaneously but have a leader-follower structure are called *Stackelberg* games. It can be shown that the first mover in a Stackelberg game has an advantage under certain scenarios [195].

To analyze repeated games, we need a refinement of the Nash equilibrium concept known as a *subgame-perfect equilibrium.* Roughly, a subgame-perfect equilibrium is one in which the initial equilibrium is simultaneously a Nash equilibrium for any subgame (the game from any subsequent stage assuming all the information on actions from the previous stages) of the initial game.

The idea is best illustrated by an example. Consider a T period, two-player, Bertrand pricing game where prices are the strategic variables. Then $[(p^1(1), p^2(1)), \dots, (p^1(T), p^2(T))]$ is subgame-perfect equilibrium if (i) it is a Nash-equilibrium and (ii) for all $1 < t \leq T$, the decisions $[(p^1(t), p^2(t)), \dots, (p^1(T), p^2(T))]$ is a Nash equilibrium for the subgame starting from period t to period T.

The subgame-perfect equilibrium refinement allows one to restrict attention to strategies that only contain credible threats or promises. For instance, in a two-player, two-period Bertrand pricing game, suppose firm one adopts a strategy of pricing high in period one and promises to continue to price high in period two provided the other firm does not undercut its price in period one. While this may result in a Nash equilibrium with both firms pricing high in each period, it does not constitute a subgame-perfect equilibrium because once the firms reach period two, it is in firm 1's interest to deviate from it's announced strategy and undercut its rival's price. Thus, the promise to continue to price high is not credible.

Infinitely repeated games (called *supergames*) provide a richer set of results than do finite repeated games. The assumption of infinite interaction may seem excessive, but in situations where there are many opportunities for frequent interactions or when the end of a game is uncertain, it is a reasonable modeling assumption.