| | |
|---|---|
| INSE6441 | Jia Yuan Yu |
| 2: Equilibrium | |
| Concordia | March 28, 2017 |

Last time, we used dominance to predict outcomes of interactions. In this lecture, we predict outcomes by introducing a notion of equilibrium. Instances of these equilibria were first introduced in the context of competition between firms.

**Example 0.1** (Unique bid auction)**.** Bidder with lowest unique bid wins. Clearly, there does not exist a pure dominant strategy (proof by contradiction).

# 1 Nash equilibrium

Many games do not have a solution by iterated strict dominance. We introduce another solution concept that is more general.

**Definition 1.1** (Pure Nash equilibrium)**.** A strategy profile $(s_1^*, \ldots, s_I^*)$ is a Nash equilibrium if for every player $i = 1, \ldots, I$, we have

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \text{for all } s_i \in S_i.$$

Nash equilibria are stable in the sense that no player gains by deviating. They predict the outcome of a game in the sense that if every player predicts that the same equilirium will occur, then no player gains by deviating.

**Definition 1.2** (Best response)**.** For a fixed player $i$, given an opponent strategy profile $s_{-i} \in S_{-i}$, a strategy $s_i* \in S_i$ is a best response if for all $s_i \in S_i$, we have

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}).$$

Let $B_i(s_{-i})$ denote the set of all best responses to $s_{-i}$. A strategy profile $(s_1^*, \ldots, s_I^*)$ is a Nash Equilibrium if

$$s_i^* \in B_i(s_{-i}^*) \quad \text{for every player } i.$$

**Example 1.1** (Exam format)**.** Given that students will write nonsense in order to get partial marks, what's the best response? Penalize wrong answers.

*Remark* 1. Nash equilibria are good predictions of the outcome if players do not make any mistakes and wish to maximize their payoff.

*Remark* 2. If iterated strict dominance yields a single strategy profile $s^*$, then this profile is also a Nash equilibrium. For all $i$ and $s_i$, we have:

$$u_i(s_i, s_{-i}^*) < u_i(s_i^*, s_{-i}^*).$$

**Example 1.2** (Battle of the sexes)**.** Some games have multiple Nash Equilibria. A boy and a girl choose between two venues for Friday night. Here are the payoffs:

| | Kendo | Salsa |
|---|---|---|
| Kendo | 2, 1 | 0, 0 |
| Salsa | 0, 0 | 1, 2 |

**Example 1.3** (Hunter's dilemma[1])**.** Two hunters go hunting. If they collaborate, then they will catch a stag. If one player goes after a hare during the hunt, then the stag gets away, but that player catches the hare and keeps it for him- or herself. A stag is preferable to two hares.

| | stag | hare |
|---|---|---|
| stag | 2, 2 | 0, 1 |
| hare | 1, 0 | 1, 1 |

*Remark* 3. In the case of multiple equilibria, some equilibria may be better than others. In the following example, $(U, L)$ is better, but $(D, R)$ is more robust to errors by an opponent.

| | L | R |
|---|---|---|
| U | 9, 9 | 0, 8 |
| D | 8, 0 | 7, 7 |

The notion of Nash equilibrium was first introduced in the study of competition between firms. We first look at the decision of a single firm.

# 2 Monopoly

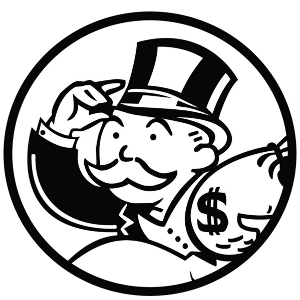

Figure 1: Monopoly Man by Vectorius

Consider a firm that is the only produce of a good (e.g., iPhones, petroleum). For simplicity, we only consider divisible goods in this lecture, so that production quantities and demands are real numbers. Let $p$ denote the price per unit of this good. Let $x : \mathbb{R}_+ \to \mathbb{R}_+$ be a strictly decreasing function, such that $x(p)$ models the demand as a function of the price $p$. Moreover, suppose that there exists $\bar{p}$ such that $x(p) = 0$ for all $p \geq \bar{p}$. Let $\gamma : \mathbb{R}_+ \to \mathbb{R}_+$ denote a strictly increasing production cost function, such that $\gamma(q)$ is the cost for producing $q$ units. We assume that $p$ and $\gamma$ are non-negative, continuous, and twice differentiable. Let $\pi = x^{-1}$ (the inverse of the demand function). The value $\pi(q)$ models the price at which all production is entirely bought in the market when the produced quantity is $q$. Observe that $\pi$ is also strictly decreasing since $x$ is strictly decreasing.

The pricing problem, given $x$ and $\gamma$, is to find the price that maximizes profit:

$$\max_{p \in \mathbb{R}_+} \quad p \cdot x(p) - \gamma(x(p)).$$

[1]Rousseau's *Discours sur l'origine et les fondements de l'inégalité parmi les hommes*

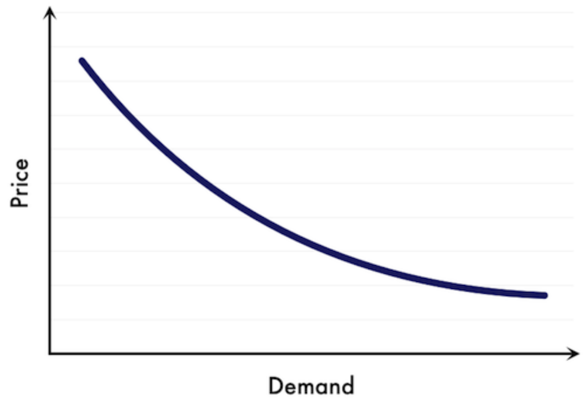


Figure 2: Market-clearing price function (inverse demand function) from `http://floriankugler.com/`.

An equivalent problem is to find the production quantity $q$ that maximizes profit:

$$\max_{q \in \mathbb{R}_+} \quad \pi(q) \cdot q - \gamma(q).$$

By taking derivatives, we find that the optimal production quantity $q^*$ satisfies:

$$\pi'(q^*)q^* + \pi(q^*) = \gamma'(q^*).$$

In other words, the marginal revenue is equal to the marginal cost. Since $\pi'(q) < 0$ (strictly decreasing assumption) for all values of $q$, we have that $\pi(q^*) > \gamma'(q^*)$—the price under monopoly exceeds marginal cost. This is in contrast to the competitive market equilibrium (e.g., the perfect competition case with a large number of competitors), where $\pi(q) = \gamma'(q)$.

In the next sections, we consider the case of multiple firms producing the same good (e.g., Android phones). These cases are called "oligopolies[2]" in general, but we will focus on the special case of duopolies with two firms. Example of oligopolies in supply chains include: OPEC (twelve oil-producing countries), and the mobile phone market in Canada.

# 3 Bertrand duopoly

Consider two firms that produce the same good, and simultaneously set their asking prices for this good. We assume that they can not change their prices afterwards—relaxing this assumption requires different notions and analyses, which is done in the game theory and economics literature. This model was advanced in 1883 by Joseph Louis Francois Bertrand. The demand function $x$ satisfies the same assumptions as the previous section.

Firm 1 and firm 2 simultaneously anounce their prices $p_1$ and $p_2$. Let $x_j(p_j, p_k)$ denote the demand (order size) for Firm $j$—$x_1(p_1, p_2)$ sales for Firm 1 and $x_2(p_2, p_1)$

[2]From the greek ὀλίγος (few) and πωλεῖν (to sell).

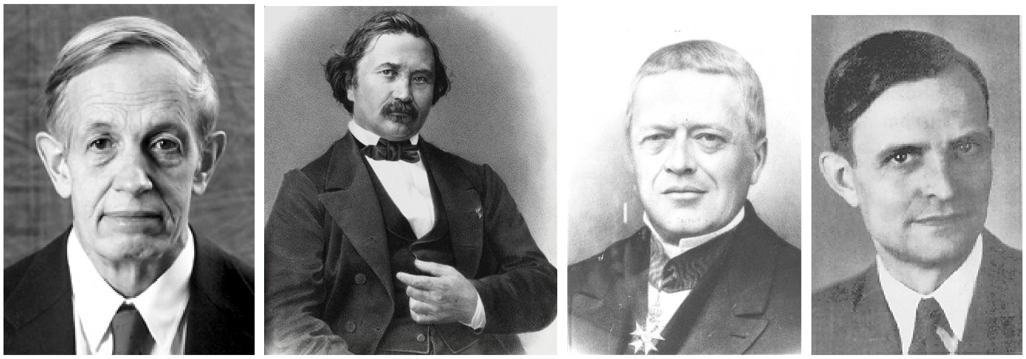

Figure 3: Nash, Bertrand, Cournot, Stackelberg (from `http://www.nobelprize.org/` and Wikipedia).

sales for Firm 2. The competition between firm 1 and firm 2 results in the following outcome (demands, order sizes):

$$x_j(p_j, p_k) = \begin{cases} x(p_j) & \text{if } p_j < p_k \\ x(p_j)/2 & \text{if } p_j = p_k \\ 0 & \text{if } p_j > p_k, \end{cases}$$

for $j = 1, 2$ and $k \neq j$.

We assume that the two firms have the same cost function, which is linear in the production quantity: $\gamma(z) = c \cdot z$. The objective of Firm $j$ is to maximize its profit:

$$u_j(p_1, p_2) = p_j \cdot x_j(p_j, p_k) - c \cdot x_j(p_j, p_k).$$

## 3.1 Equilibrium prices

The notion of equilibrium in the interactions between two decision makers (Firm 1 and Firm 2) is defined as follows. First, let $A$ denote the set of possible decisions of each decision maker; $A$ is the set $\mathbb{R}_+$ of possible prices for Bertrand duopoly. Let $u_j : A \times A \to \mathbb{R}$ for $j = 1, 2$ denote the profit function of Firm $j$.

In the Bertrand duopoly setting, A *Nash equilibrium* is a pair of decisions $(p_1^*, p_2^*) \in A \times A$ such that

$$\begin{aligned} u_1(p_1^*, p_2^*) &\geq u_1(p_1, p_2^*), \quad \text{for all } p_1 \in A, \\ u_2(p_1^*, p_2^*) &\geq u_2(p_1^*, p_2), \quad \text{for all } p_2 \in A. \end{aligned}$$

In other words, when Firm decides $p_2^*$, Firm 1 is best off deciding $p_1^*$ and vice versa.

How can we find a Nash equilibrium? In the Bertrand duopoly case, we will do so by guessing and verifying whether our guess satisfies the defining properties of a Nash equilibrim. Consider the pair of prices $(p_1^*, p_2^*) = (c, c)$, i.e., each firm sells the goods for zero profit at the production cost of the good. Lowering prices results in loses, or negative profit, as opposed to zero profit. Neither firm can gain by raising its price because its sales would drop to zero, resulting in the same profit of zero. Therefore, $(p_1^*, p_2^*) = (c, c)$ is an equilibrium pair of prices.

The equilibrium for the Bertrand duopoly is unique and more desirable than the monopoly outcome, but it does not reflect reality. Can we find an equilibrium that is more realistic, but still more desirable than the monopoly outcome?

# 4 Cournot duopoly

Antoine Augustin Cournot proposed the following model in 1838. Consider two firms who simultaneously compete on their production quantities $q_1 \in \mathbb{R}_+$ and $q_2 \in \mathbb{R}_+$ of the same good instead of prices. Let $\pi$ denote the inverse demand function as in the monopoly case. The price of the good resulting from this competition is

$$\pi(q_1 + q_2),$$

which clears the market by definition of $\pi$.

*Remark* 4. An example of Cournot duopoly is farmers picking up perishable goods from their fields each morning and bringing them to the market. They sell all them items at once in the manner of a Dutch auction: starting with a high price, reducing it until the total of their goods equals the demand.

As in the case of Bertrand duopoly, we assume that the production cost per unit is a constant $c$; but we also assume that $\pi(0) > c$. The function $\pi$ and constant $c$ are given as input to the decision maker; the profit of each Firm $j$ is

$$u_j(q_1, q_2) = \pi(q_j + q_k) \cdot q_j - cq_j.$$

## 4.1 Equilibrium production quantities

Observe that, in contrast to single-decision-maker problems, the optimal decision of each firm is a function of the other firm's decision:

$$\begin{aligned} q_1^{\text{opt}}(q_2) &= \arg\max_{q_1} \quad \pi(q_1 + q_2) \cdot q_1 - cq_1, \\ q_2^{\text{opt}}(q_1) &= \arg\max_{q_2} \quad \pi(q_1 + q_2) \cdot q_2 - cq_2. \end{aligned}$$

These optimal decision functions are called "best-response functions" and are illustrated in Figure 4.1 for the special case of a linear $\pi$ function.

In the case of Cournot duopoly, an equilibrium—if it exists—is a pair of production quantities that we denote $(q_1^*, q_2^*)$. One approach to find an equilibrim by finding the intersection of the best-response functions by visual inspection. An alternative, analytical approach is to set the first derivative ($\frac{\mathrm{d}}{\mathrm{d}q_1}$ and $\frac{\mathrm{d}}{\mathrm{d}q_2}$, respectively) of each firm's profit equal to zero to find the optimal decisions, which is similar to the EOQ solution. We obtain the following necessary conditions for the optimal production

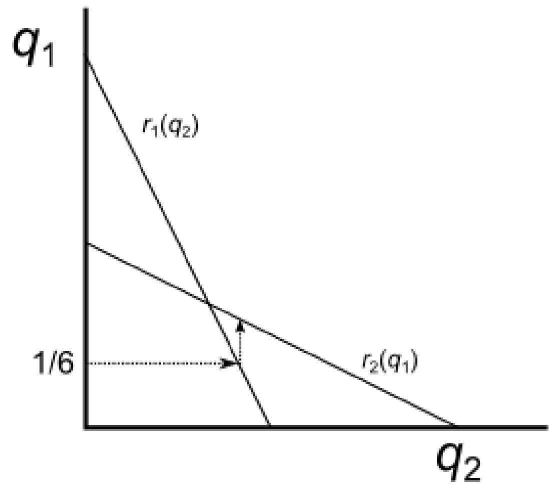


Figure 4: Optimal quantity $q_1^{\text{opt}} = r_1$ as a function of $q_2$ and vice versa. From `http://mindyourdecisions.com/`.

quantities for both firms[3]:

$$\pi'(q_1^* + q_2^*)q_1^* + \pi(q_1^* + q_2^*) = c,$$
$$\pi'(q_1^* + q_2^*)q_2^* + \pi(q_1^* + q_2^*) = c.$$

We can solve this system of two equations in two unknowns by rewriting it as

$$\pi'(2q_1^*)q_1^* + \pi(2q_1^*) = c,$$
$$q_1^* = q_2^*,$$

where we used the fact that $\pi'(q) < 0$ (strictly decreasing assumption) for all $q \geq 0$.

**Example 4.1** (Linear function $\pi$)**.** See Figure 4.1. If $\pi(q) = 1 - q$, then we have $q_1^* = q_2^* = (1 - c)/3$.

How does $q_1^* + q_2^*$ compare with the solution $q^*$ from the monopoly market?

# 5 Mixed strategies

Some games do not have a pure Nash Equilibrium. We review probability and introduce the notion of mixed strategies.

**Example 5.1** (Rock Paper Scissors)**.**

| | R | P | S |
|---|---|---|---|
| R | 0, 0 | -1, 1 | 1, -1 |
| P | 1, -1 | 0, 0 | -1, 1 |
| S | -1, 1 | 1, -1 | 0, 0 |

How do people play this game in real life?

[3]Homework: check that $q_1 = 0$ is not optimal for Firm 1, nor is $q_2 = 0$ for Firm 2.

A mixed strategy $\sigma_i$ is a probability distribution over pure strategies $S_i$. The (expected) payoff to player $i$ corresponding to a profile of mixed strategies $(\sigma_1, \ldots, \sigma_I)$ is

$$\begin{aligned}
\tilde{u}_i(\sigma_i, \sigma_{-i}) &= \sum_{(z_1,\ldots,z_I)\in S_1\times\ldots\times S_I} \underbrace{\sigma_1(z_1)\ldots\sigma_i(z_i)\ldots\sigma_I(z_I)}_{\mathbb{P}(X_1=z_1,\ldots,X_I=z_I)} u_i(z_i, z_{-i}) \\
&= \sum_{z\in S_1\times\ldots\times S_I} \left(\sigma_i(z_i)\prod_{j\neq i}\sigma_j(z_j)\right) u_i(z_i, z_{-i}) \\
&= \sum_{z\in S_1\times\ldots\times S_I} \left(\prod_{j=1}^{I}\sigma_j(z_j)\right) u_i(z_i, z_{-i}).
\end{aligned}$$

**Example 5.2** (Pure strategy strictly dominated by mixed strategy)**.** Consider

| | L | R |
|---|---|---|
| U | 2, 0 | -1, 0 |
| M | 0, 0 | 0, 0 |
| D | -1, 0 | 2, 0 |

Player 1's action M is strictly dominated by the mixed action (1/2, 0, 1/2) regardless of Player 2's mixed or pure strategy.

**Example 5.3** (Mixed strategy strictly dominated)**.** Consider

| | L | R |
|---|---|---|
| U | 1, 3 | -2, 0 |
| M | -2, 0 | 1, 3 |
| D | 0, 1 | 0, 1 |

Player 1's mixed action (1/2, 1/2, 0) is dominated by D regardless of Player 2's mixed or pure strategy. However, neither U nor M, as pure strategies, are dominated.

**Definition 5.1** (Mixed Nash equilibrium)**.** A mixed strategy profile $(\sigma_1^*, \ldots, \sigma_I^*)$ is a Nash equilibrium if for every player $i = 1, \ldots, I$, we have

$$\tilde{u}_i(\sigma_i^*, \sigma_{-i}^*) \geq \tilde{u}_i(z, \sigma_{-i}^*) \quad \text{for all } z \in S_i.$$

Check that in the Rock-Paper-Scissors game, a profile composed of uniform mixed strategies is a Nash equilibrium.

*Remark* 5 (Support of NE mixed strategy). In a Nash equilibrium, a player must be indifferent to all pure strategies with nonzero probability (i.e., in the support of that player's mixed strategy). Otherwise, that player can increase payoff by shifting some probability from one strategy to another.

**Example 5.4** (Inspection game)**.** A worker (PhD student) can work or skip work. A boss (supervisor) can inspect or not. The wage to the worker is $w$, the worker's cost of working is $g$, the boss' cost of inspection is $h$, the boss' profit from the work is $v$. Assume that $w > g > h > 0$.

| | Inspect | Don't inspect |
|---|---|---|
| Skip | $0, -h$ | $w, -w$ |
| Work | $w-g, v-w-h$ | $w-g, v-w$ |

Let $\sigma_1 = (x, 1-x)$ and $\sigma_2 = (y, 1-y)$ be the mixed strategies of the players. Verify that there are no pure Nash equilibrium. Verify that:

$$\begin{aligned}
\tilde{u}_1(\text{Skip}, \sigma_2) &= (1-y)w, \\
\tilde{u}_1(\text{Work}, \sigma_2) &= y(w-g) + (1-y)(w-g), \\
\tilde{u}_2(\sigma_1, \text{Inspect}) &= x(-h) + (1-x)(v-w-h), \\
\tilde{u}_2(\sigma_1, \text{Don't inspect}) &= x(-w) + (1-x)(v-w).
\end{aligned}$$

Let's look for a mixed Nash equilibrium profile of the form $(\sigma_1^*, \sigma_2^*)$. By Remark 5, we must have $\tilde{u}_1(\text{Skip}, \sigma_2) = \tilde{u}_1(\text{Work}, \sigma_2)$ and $\tilde{u}_2(\sigma_1, \text{Inspect}) = \tilde{u}_2(\sigma_1, \text{Don't inspect})$. By algebra, a Nash equilibrium is therefore skipping work with probability $x = h/w$ and inspecting with probability $y = g/w$.

# 6 Reading material

- Chapter 1 of Fudenberg and Tirole.
- Chapter 12 of Microeconomics Theory (Mas-Colell, Whinston, Green).