

ECON106P: Pricing and Strategy

Yangbo Song

Economics Department, UCLA

June 30, 2014

Game theory is a methodology used to analyze *strategic* situations in economics, politics, psychology, etc.

- There are multiple (> 1) agents, each needs to make a choice
- Choices may be made simultaneously or sequentially
- “Strategic” means that each agent is *selfish*, i.e. he/she is trying to maximize his/her own payoff
- The choice of one agent may affect the payoffs of others

Example of a simultaneous game: beauty contest

- There are 100 persons, each required to write down a real number between 0 and 100
- No one can see the number chosen by the others
- After everybody writes down their number, the average of the 100 numbers are calculated, and the person whose number is closest to $\frac{1}{2}$ of this average will win \$1,000,000 (in the case of a tie, the winners share the prize evenly)
- What number should a contestant choose?

Example of a sequential game: pirates of the Caribbean

- 100 pirates need to divide 5 gold coins among themselves. They agree to the following rule:
 - Starting with pirate 1, each pirate proposes a plan, then the 100 pirates take a vote (for/against)
 - If at least half of the (remaining) pirates vote for the plan, then the gold coins are divided accordingly; otherwise, the proposer is executed
- Each pirate values his life most
- Being alive, he prefers more gold coins than less
- Being alive and having the same number of gold coins, he prefers to watch others die
- How many pirates would die in the end?

Essential elements of a game:

- Players/Agents: who are playing the game
- Strategies: what the available choices for each player are (a *complete* description of game plan)
 - Finite strategies: strategies can be counted as $1, 2, 3, \dots$
 - Infinite strategies: strategies lie in a continuous interval
- Payoffs: given each combination of choices, what each player gets

Example: simultaneous game with finite strategies – the prisoners' dilemma

- Two outlaws are caught and are given the choice of confession or denial
- If both confess, both get 2 years of jail time
- If one confesses and the other denies, the former walks free, and the latter will do 5 years
- If both deny, both get 1 year

- Players: outlaws 1 and 2
- Strategies: for each player: C and D
- Payoffs: $(C, C) \rightarrow (-2, -2)$; $(C, D) \rightarrow (0, -5)$; $(D, C) \rightarrow (-5, 0)$; $(D, D) \rightarrow (-1, -1)$
- This game can be represented by a *matrix* (drawn in class).

Example: sequential game with finite strategies – coordinated attacks

- Two generals are deciding which spot (L or R) to attack in a military operation
- General 1 gives his order first, and General 2 can see 1's order
- If they attack the same spot, the operation will succeed; otherwise, it will fail

- Players: generals 1 and 2
- Strategies:
 - For 1: L and R
 - For 2: $(L|L, L|R)$, $(L|L, R|R)$, $(R|L, L|R)$ and $(R|L, R|R)$
- Payoffs: $(L, (L|L, \cdot)) \rightarrow (1, 1)$; $(L, (R|L, \cdot)) \rightarrow (-1, -1)$;
 $R, (\cdot, L|R) \rightarrow (-1, -1)$; $R, (\cdot, R|R) \rightarrow (1, 1)$
- This game can be represented by a *tree diagram* (drawn in class).

Example: simultaneous game with infinite strategies – Cournot competition

- Two firms are competing in a market with demand function $P(q) = a - bq$
- Quantity sold in the market would be the sum of quantities produced by the two firms
- They cannot see how much the competitor has produced

- Players: firms 1 and 2
- Strategies: for player i : $q_i \in [0, \infty)$
- Payoffs:
 $(q_1, q_2) \rightarrow ((a - b(q_1 + q_2))q_1 - C_1(q_1), (a - b(q_1 + q_2))q_2 - C_2(q_2))$
- There is no straight-forward graphical representation of this game.

Example: sequential game with infinite strategies – Stackelberg competition

- Two firms are competing in a market with demand function $P(q) = a - bq$
- Quantity sold in the market would be the sum of quantities produced by the two firms
- Firm 2 can see the quantity produced by firm 1 before its own production

- Players: firms 1 and 2
- Strategies:
 - For 1: q_1
 - For 2: any *function* $q_2(q_1)$
- Payoffs: $(q_1, q_2(q_1)) \rightarrow ((a - b(q_1 + q_2(q_1)))q_1 - C_1(q_1), (a - b(q_1 + q_2(q_1)))q_2 - C_2(q_2(q_1)))$
- This game can be represented by a *tree diagram* (drawn in class).

Consider a player i . Let a_i denote one of i 's strategies and let a_{-i} denote one combination of the other players' strategies.

We say that a_i is i 's **best response** to a_{-i} if given a_{-i} , a_i yields the highest payoff for player i .

- Never say “ a_i is i 's best response” – should be “ a_i is i 's best response to a_{-i} ”
- Different a_{-i} 's may point to different best responses for i
- Even for one a_{-i} , there can be multiple best responses for i

Dominant and dominated strategy

We say that a_i is i 's:

- **weakly dominant strategy** if a_i is one of the best responses to *any* a_{-i} , and the only best response to *some* a_{-i}
- **strictly dominant strategy** if a_i is the only best response to *any* a_{-i}
- **weakly dominated strategy** if there exists some other strategy a'_i , such that a'_i is weakly better than a_i for *any* a_{-i} , and strictly better than a_i for *some* a_{-i}
- **strictly dominated strategy** if there exists some other strategy a'_i , such that a'_i is strictly better than a_i for *any* a_{-i}
- Does the existence of a dominant strategy imply the existence of a dominated strategy? How about the reverse?

In the monopoly's problem we solved for the profit-maximizing action.
What do we solve for in a strategic game?

- Each player wants to maximize his/her own payoff
- The solution should be “stable”, i.e. given the others' choice, each player should be happy to stay with his own choice
- The above suggests mutual best response

Simultaneous game: Nash Equilibrium

A **Nash Equilibrium** is a strategy profile $a = (a_1, \dots, a_N)$ such that for any i , a_i is a best response to a_{-i} .

- A strictly dominated strategy will never be played in any NE (how about a weakly dominated one?)
- NE may not be unique
- NE may not be the efficient outcome

Sequential game: Subgame Perfect Nash Equilibrium

In a sequential game, a **subgame** is a smaller game that can be obtained from the whole game, i.e. a “mini tree” that can be obtained from the whole tree diagram.

A **Subgame Perfect Nash Equilibrium** is a strategy profile $a = (a_1, \dots, a_N)$ such that it is a Nash Equilibrium in any subgame.

- A NE of the whole game may not be a SPNE
- To solve for a SPNE, we use *backward induction*

Cournot Competition

Now we take a closer look at a general version of Cournot competition.

- There are N firms in the market: $1, 2, \dots, N$
- The market demand function is $P = a - bQ$
- Firm i has a cost function $C_i(q_i)$
- Firms choose their quantities *simultaneously*

Cournot Competition

Case 1: identical firms: $C_i(q_i) = C(q_i)$

For any i , its profit can be written as

$$\pi_i = (a - b \sum_j q_j) q_i - C(q_i)$$

Thus firm i 's best response to $q_{-i} = (q_1, \dots, q_{i-1}, q_{i+1}, \dots, q_N)$ is the solution to $\max_{q_i} \pi_i$. FOC:

$$a - b \sum_{j \neq i} q_j - 2bq_i - C'(q_i) = 0$$

Cournot Competition

NE is the solution to the following equation system:

$$a - b \sum_{j \neq 1} q_j - 2bq_1 - C'(q_1) = 0$$

...

$$a - b \sum_{j \neq N} q_N - 2bq_N - C'(q_N) = 0$$

However, notice that the firms are *identical*, thus intuitively in NE each firm should produce the same quantity: $q_1 = \dots = q_N = q$. Then each of the above equations becomes

$$a - b(N-1)q - 2bq - C'(q) = 0$$

We can then solve for the equilibrium quantity q^* .

Cournot Competition

Example:

- $N=3$
- Market demand: $P = 12 - Q$
- Identical costs: $C(q_i) = 4q_i$

The previous equation becomes

$$12 - 2q - 2q - 4 = 0$$

Therefore $q^* = 2$, i.e. the NE is $q_1 = q_2 = q_3 = 2$.

Cournot Competition

Case 2: heterogeneous firms

We have a similar FOC:

$$a - b \sum_{j \neq i} q_j - 2bq_i - C_i(q_i) = 0$$

And a similar system of equations:

$$a - b \sum_{j \neq 1} q_j - 2bq_1 - C'_1(q_1) = 0$$

...

$$a - b \sum_{j \neq N} q_N - 2bq_N - C'_N(q_N) = 0$$

But now we cannot assume that firms produce the same quantity.

Cournot Competition

Example:

- $N=3$
- Market demand: $P = 12 - Q$
- Heterogeneous costs: $C_1(q_1) = q_1$, $C_2(q_2) = 2q_2$, $C_3(q_3) = 3q_3$

The previous system of equations becomes

$$12 - (q_2 + q_3) - 2q_1 - 1 = 0$$

$$12 - (q_1 + q_3) - 2q_2 - 2 = 0$$

$$12 - (q_1 + q_2) - 2q_3 - 3 = 0$$

We have $q_1 = \frac{7}{2}$, $q_2 = \frac{5}{2}$, $q_3 = \frac{3}{2}$.

Bertrand Competition

In Cournot competition, firms choose their quantities. Now we consider a model where firms choose prices.

- There are N firms in the market: $1, 2, \dots, N$
- Firm i has a (linear) cost function $C_i(q_i) = c_i q_i$
- Firms choose their prices *simultaneously*

Bertrand Competition

Case 1: firms produce identical goods (this is *always* the case in Cournot competition).

- Consumers will buy from the firm whose price is the lowest
- Suppose that $c_1 < c_2 < \dots < c_N$
- In NE, can the market price (i.e. the lowest price offered) be higher than c_N ?
- In NE, can the market price be between c_{N-1} and c_N ? How about between c_{N-2} and c_{N-1} ?

Bertrand Competition

Example:

- $N=3$
- Market demand: $P = 12 - Q$
- Heterogeneous costs: $C_1(q_1) = q_1$, $C_2(q_2) = 2q_2$, $C_3(q_3) = 3q_3$

Suppose that $\min\{P_1, P_2, P_3\} > 3$. Now each firm has an incentive to choose a price just below this minimum, to capture the whole market.

Thus this cannot happen in any NE.

Suppose that $\min\{P_1, P_2, P_3\} \in (2, 3]$. Now firm 1 and firm 2 both have an incentive to choose a price just below this minimum, to capture the whole market. Thus this cannot happen in any NE.

Following this argument, the NE is that firm 1 chooses $P_1 \in [1, 2]$, and none of the other firms are selling anything.

The previous argument can be generalized.

- If there is a single firm having the lowest marginal cost, then the market price would be between the lowest and the second-lowest marginal cost
- Otherwise, the market price would be equal to the lowest marginal cost

Case 2: firms produce heterogeneous goods

- Firm i faces the demand function $q_i = a_i - b_i P_i + \sum_{j \neq i} d_{ij} P_j$
- Implication: prices set by competitors affect own demand
- If $d_{ij} > 0$, then goods i and j are *substitutes*; otherwise, they are *complements*

Bertrand Competition

Firm i 's problem:

$$\max_{P_i} P_i(a_i - b_i P_i + \sum_{j \neq i} d_{ij} P_j) - c_i(a_i - b_i P_i + \sum_{j \neq i} d_{ij} P_j)$$

FOC:

$$a_i + \sum_{j \neq i} d_{ij} P_j - 2b_i P_i + b_i c_i = 0$$

As in Cournot competition, we solve a system of equations.

Bertrand Competition

Example:

- $N=2$
- Demand: $q_i = 10 - P_i + P_j$ for $i = 1, 2, j \neq i$
- Cost: $c_1 = c_2 = 2$

The system of equations becomes

$$10 + P_2 - 2P_1 + 2 = 0$$

$$10 + P_1 - 2P_2 + 2 = 0$$

We have $P_1 = P_2 = 12$. Remark: we can assume $P_1 = P_2$ to simplify the problem, but we need *symmetric demand functions* and *identical cost functions*.

ECON106P: Pricing and Strategy

Yangbo Song

Economics Department, UCLA

July 23, 2014

Principal-Agent Model

Recall the basic setting of a principal-agent model:

- One principal, possibly multiple agents
- The principal proposes a mechanism/game (or specifies parameters in a given game)
- The agents observe their private information
- The agents choose their actions in the game proposed

Before thinking about the objective of the principal, we ask: what is a “reasonable” mechanism?

Principal-Agent Model

Two conditions need to be satisfied:

- Participation constraint: the agents must be willing to take part in the game
- Incentive compatibility constraint: if an option is designed for a certain type of agent, that agent must have the incentive to choose it but not something else

Essentially, these two conditions describe a best response: it's best to

- Play the game rather than walk away
- Choose what is meant for the agent himself in the game

Principal-Agent Model

Example: education

- A firm has two job openings, A and B
- There are two possible types of workers, a and b ; type is a worker's private information
- The firm wants a type a worker to do job A , and a type b worker to do job B
- The firm can set two different wages, w_A and w_B , for the two jobs, then let a worker choose the job he wants

How can the firm achieve its goal?

Naturally, a worker will choose the job with the higher wage, so the firm cannot separate the two types for different jobs. However, assume the following:

- A worker can choose to receive education and obtain a degree
- For a type a worker, education costs c_a ; for a type b worker, education costs c_b (“cost” also reflects the intellectual difficulty of finishing a degree)
- Now, a firm can offer contracts of jobs differentiated by requirement of a degree

Principal-Agent Model

To make the degree requirement helpful, it has to be that one job requires a degree while the other does not. So we may as well assume that job A requires a degree. How should the firm set the wages?

- Participation constraint:

$$w_A - c_a \geq 0$$

$$w_B \geq 0$$

- Incentive compatibility constraint:

$$w_A - c_a \geq w_B$$

$$w_B \geq w_A - c_b$$

Principal-Agent Model

In general, there are two types of principal-agent model:

- Adverse selection: the principal cannot observe the agents' *types*, as in the previous examples
- Moral hazard: the principal cannot observe the agents' *actions*, as in classes (that's why we need exams and grades)

We will discuss each type using a detailed example.

Let's go back to the 2nd-degree price discrimination story.

- A zero-cost firm is facing a consumer whose type is private information
- The consumer is of type A with probability $\frac{2}{3}$, and of type B with probability $\frac{1}{3}$
- A type A consumer has demand function $P = 2(2 - q)$, and a type B consumer has demand function $P = 2 - q$
- The firm can offer two plans, (q_A, T_A) and (q_B, T_B) , where q is the quantity provided to the consumer and T is the consumer's payment
- The firm's ultimate goal: maximize expected profit using a "reasonable" mechanism

As before, we first ask what a “reasonable” mechanism is. We need to compute the consumer’s surplus given his type:

- A type A consumer’s surplus is a function

$$v_A(q) = q(4 - q)$$

- A type B consumer’s surplus is a function

$$v_B(q) = q(2 - \frac{q}{2})$$

Next, we write down the relevant constraints for each type of consumer.

- Participation constraint:

$$q_A(4 - q_A) - T_A \geq 0 \quad (1)$$

$$q_B(2 - \frac{q_B}{2}) - T_B \geq 0 \quad (2)$$

- Incentive compatibility constraint:

$$q_A(4 - q_A) - T_A \geq q_B(4 - q_B) - T_B \quad (3)$$

$$q_B(2 - \frac{q_B}{2}) - T_B \geq q_A(2 - \frac{q_A}{2}) - T_A \quad (4)$$

Adverse Selection

We can simplify the constraints by the following steps (work for all adverse selection problems!)

Step 1: note that (1) is satisfied once (2)-(4) are satisfied.

- $q_B(4 - q_B) = 2(q_B(2 - \frac{q_B}{2})) > q_B(2 - \frac{q_B}{2})$
- Hence, from (2) and (3) we get (1)

Step 2: ignore (4) at the moment.

- Presumably, the cost of a plan aiming a higher demand (type A) should cost more
- If choosing plan A , a type B consumer does not have as much benefit as a type A consumer

Now we have a *relaxed* problem with only constraints (2) and (3). Since the firm wants to maximize expected profit, we can do further simplification:

Step 3: when the firm is maximizing profit, both (2) and (3) must be “=”.

- If (2) is not “=”, can raise T_B to increase profit (decreasing LHS of (2) and RHS of (3))
- If (3) is not “=”, can raise T_A to increase profit (decreasing LHS of (3))

Now, we have a clean representation of the firm's profit.

$$\mathbb{E}[\pi] = \frac{2}{3}T_A + \frac{1}{3}T_B$$

where

$$\begin{aligned}T_B &= q_B(2 - \frac{q_B}{2}) \\T_A &= q_A(4 - q_A) + T_B - q_B(4 - q_B) \\&= q_A(4 - q_A) - q_B(2 - \frac{q_B}{2})\end{aligned}$$

The firm's problem then becomes

$$\max_{q_A, q_B} \frac{2}{3}q_A(4 - q_A) - \frac{1}{3}q_B(2 - \frac{q_B}{2})$$

This problem is equivalent to

$$\begin{aligned} \max_{q_A} q_A(4 - q_A) \\ \min_{q_B} q_B(2 - \frac{q_B}{2}) \end{aligned}$$

Note that q_A and q_B must lie in $[0, 2]$.

Hence, the profit-maximizing plans in this *relaxed* problem are $(q_A, T_A) = (2, 4)$ and $(q_B, T_B) = (0, 0)$. The final step is to check the omitted constraint (4):

$$q_B(2 - \frac{q_B}{2}) - T_B \geq q_A(2 - \frac{q_A}{2}) - T_A$$

The LHS is equal to zero while the RHS is negative. Hence this constraint is satisfied, and we have found the profit-maximizing plans.

Important facts in the profit-maximizing plan:

- Both types of consumer have zero surplus
- q_A is efficient while q_B is not
- The firm earns an expected profit of $\frac{8}{3}$

If demand functions are allowed to be non-linear, the general result is

- $q_A > q_B, T_A > T_B$
- q_A is efficient while q_B is not
- Type B consumer always has zero surplus
- Type A consumer may have some surplus

What if there is no asymmetric information, i.e. the firm can observe the type of the consumer?

- Essentially, the firm is applying 1st-degree price discrimination on both types
- The firm's problem becomes

$$\max_{q_A, q_B} \frac{2}{3}q_A(4 - q_A) + \frac{1}{3}q_B(2 - \frac{q_B}{2})$$

- The profit-maximizing plans are $(q_A, T_A) = (2, 4)$ and $(q_B, T_B) = (2, 2)$

Important facts in the profit-maximizing plan:

- Both types of consumer have zero surplus
- q_A and q_B are both efficient
- The firm earns an expected profit of $\frac{10}{3}$

If demand functions are allowed to be non-linear, the general result is

- Both types of consumer have zero surplus
- q_A and q_B are both efficient
- Profit is higher than the case with asymmetric information

Without asymmetric information, the outcome is called *first-best*.

- Efficient allocation
- Zero surplus for consumer
- Maximum profit for firm

With asymmetric information, the outcome is called *second-best*.

- Efficient allocation for higher-demand consumer
Inefficient allocation for lower-demand consumer
- Zero surplus for lower-demand consumer and some surplus for higher-demand consumer
- Lesser profit for firm

Moral hazard describes the situation where the principal cannot observe the agent's *action*.

- If the principal and the agent's interests are aligned (i.e. they want exactly the same thing), then there is no "hazard"
- Otherwise, without an appropriate mechanism, the agent may do something that is both inefficient and harmful to the principal
- As in adverse selection, we may expect in the profit-maximizing outcome:
 - Not efficient allocation
 - Less profit than with observable action

Consider the following model:

- A firm is hiring a manager
- The firm's profit partially depends on the manager's effort e :
 $\pi = e + \epsilon$, where ϵ is a random variable with expectation 0 and variance σ^2
- The manager has a cost of $\frac{1}{2}e^2$ of exerting effort
- The firm cannot see the effort level e , but can see the realized profit π

The firm offers a contract – a wage function – $s(\pi)$ to the manager

- The firm's payoff: $\mathbb{E}[\pi - s(\pi)]$
- The manager's payoff: $\mathbb{E}[s(\pi)] - \frac{1}{2}Var(s(\pi)) - \frac{1}{2}e^2$
- The firm is risk-neutral: it only cares about expectation
- The manager is risk-averse: it also prefers high expectation but dislikes fluctuation

Some remarks about risk aversion:

- If the manager is also risk-neutral, then the firm could just set the wage to be “profit minus some constant”
- The manager will then exert effort to maximize the expected surplus, and this surplus will be taken by the firm
- However, when the manager is risk-averse, setting such a wage function results in a large variance
- It may cause the manager unwilling to accept the contract
- Hence, with risk aversion, the firm will not be able to obtain the largest possible surplus

We restrain our discussion on linear contracts: $s(\pi) = a + b\pi$.

- The firm's payoff becomes $(1 - b)e - a$
- The manager's payoff becomes $a + be - \frac{b^2\sigma^2}{2} - \frac{1}{2}e^2$

The game can then be described as follows:

- Given a contract, the manager will act for his own benefit, according to the participation and incentive compatibility constraints
- The firm will choose a and b to maximize its profit

How do we represent the constraints in this setting?

- Incentive compatibility constraint: the manager is choosing the best effort level for himself
 - The manager solves

$$\max_e a + be - \frac{b^2\sigma^2}{2} - \frac{1}{2}e^2$$

- FOC implies that $e = b$
- Participation constraint: the maximum payoff from the contract is non-negative

$$a + \frac{b^2}{2} - \frac{b^2\sigma^2}{2} \geq 0$$

Thus we have the firm's problem:

$$\max_{a,b} (1-b)e - a$$

subject to

$$\begin{aligned} e &= b \\ a + \frac{b^2}{2} - \frac{b^2\sigma^2}{2} &\geq 0 \end{aligned}$$

The first constraint can immediately be eliminated:

$$\begin{aligned} &\max_{a,b} (1-b)b - a \\ \text{s.t. } &a + \frac{b^2}{2} - \frac{b^2\sigma^2}{2} \geq 0 \end{aligned}$$

Next, note that in the profit-maximizing contract, it must be the case that the second constraint is binding:

$$a + \frac{b^2}{2} - \frac{b^2\sigma^2}{2} = 0$$

Otherwise, the firm can always reduce a to save cost, i.e. increase profit. Hence, the problem becomes

$$\max_b (1-b)b - \frac{(\sigma^2 - 1)b^2}{2}$$

The maximand can be simplified as $b - \frac{(\sigma^2 + 1)b^2}{2}$. FOC would imply that

$$b = \frac{1}{1 + \sigma^2}$$

And from the binding constraint,

$$a = \frac{\sigma^2 - 1}{2(1 + \sigma^2)^2}$$

The firm's profit is equal to $\frac{1}{2(1 + \sigma^2)}$.

We now have the second-best outcome. What is the first-best one?

- The firm can observe, and hence dictate, the manager's effort level
- The contract is like “do e and get the wage function, or get nothing”)
- Only the participation constraint needs to be satisfied (and it's binding at optimum)
- The firm's problem now:

$$\begin{aligned} & \max_{e,a,b} (1-b)e - a \\ \text{s.t. } & a + be - \frac{b^2\sigma^2}{2} - \frac{1}{2}e^2 = 0 \end{aligned}$$

Re-write the problem as

$$\begin{aligned} \max_{e,a,b} \quad & e - (a + be) \\ \text{s.t.} \quad & a + be = \frac{b^2\sigma^2}{2} + \frac{1}{2}e^2 \end{aligned}$$

- For any e , the cost of the firm is at least $\frac{1}{2}e^2$
- This lower bound is reached at $b = 0$ and $a = \frac{1}{2}e^2$
- Hence, the problem can be simplified as $\max_e e - \frac{1}{2}e^2$, with solution $e = 1$

Compare the first-best outcome with the second-best outcome:

- The former is efficient while the latter is not
 - The efficient outcome is (a, b, e) that maximizes

$$(1 - b)e - a + a + be - \frac{b^2\sigma^2}{2} - \frac{1}{2}e^2 = e - \frac{1}{2}e^2$$

- Hence the solution is $b = 0, e = 1, a = \frac{1}{2}$
- The firm's profit is higher in the former case ($= \frac{1}{2}$)
- The manager earns zero in expectation in both cases

What if there is still asymmetric information, but the manager is risk-neutral?

- The manager's payoff now becomes $a + be - \frac{1}{2}e^2$
- The incentive compatibility constraint still implies that $e = b$
- The (binding) participation constraint becomes $a + \frac{1}{2}b^2 = 0$
- The firm's problem becomes

$$\max_{a,b} (1-b)b - a$$

$$\text{s.t. } a + \frac{b^2}{2} = 0$$

As before, we can eliminate the binding constraint and get

$$\max_b b - \frac{1}{2}b^2$$

The solution is $b = 1$ and $a = -\frac{1}{2}$. It is identical to the first-best (efficient) case in terms of

- Effort level
- Manager's payoff
- Firm's profit

To summarize:

- First-best: the principal knows everything, and hence can maximize the surplus and take all that surplus
- Hence the first-best outcome must be efficient, and must yield zero to the agent
- Second best: the agent still gets zero in expectation
- When the agent and the principal have aligned interests (e.g. both of them only care about expectation), the second-best outcome “coincides” with the first-best one
- Otherwise, inefficiency and less profit result from asymmetric information and conflict of interests

ECON106P: Midterm Exam

Yangbo Song

July 9, 2014

1 Multiple Layers of Vertical Separation (6 pts.)

Consider a retail market with demand function $P = a - bq$. There is one manufacturer, denoted M , and n distributors, denoted $1, 2, \dots, n$. The distributors are vertically aligned, i.e. distributor n buys from $n - 1$ and sells to the retail market, distributor $n - 1$ buys from $n - 2$ and sells to n , \dots , and distributor 1 buys from M and sells to 2 . The manufacturer's cost function is $C(q) = cq$. Parameters a, b, c are positive constants which satisfy $a > c$.

Each firm is a monopoly in their own market. The manufacturer sets a wholesale price P_M , distributor $i = 1, 2, \dots, n - 1$ each sets an intermediate price P_i (i.e. firm $i + 1$ buys at price P_i), and distributor n sets a retail price P .

Answer the following questions:

1. What is firm n 's demand function (i.e. the demand function faced by firm $n - 1$)? How about $n - 1$'s? How about $n - 2$'s? (Hint: use backward induction.) (2 pts.)
2. What is firm 1 's demand function (i.e. the demand function faced by firm M)? (1 pt.)
3. How many units will firm M produce? What will be the retail price? (2 pts.)
4. When n becomes very large, what are the approximate price and quantity in the retail market? (1 pt.)

2 Cournot Competition with Non-Linear Cost (6 pts.)

Consider a market with demand function $P = 100 - 2Q$. There are n firms in the market as Cournot competitors, each with the same cost function $C(q) = q^2 + q$.

The following information is given: the way to produce quantity Q at the smallest possible cost is to make each firm produce $\frac{Q}{n}$.

Answer the following questions:

1. What is the socially optimal cost function $C^e(Q)$ (i.e. the smallest cost of producing Q)? (1 pt.)
2. Given this cost function $C^e(Q)$, what is the efficient quantity Q^* ? (2 pts.)
3. When n becomes very large, what is the approximate Q^* ? (1 pt.)
4. Show that when n becomes very large, the total quantity produced in the (symmetric) Nash equilibrium of the Cournot game converges to your answer to 3. (2 pts.)

3 Merger among Bertrand Competitors (6 pts.)

Consider a market with demand function $P = 10 - Q$. There are 3 firms in the market, denoted 1, 2 and 3, as Bertrand competitors. The cost functions are $C_1(q_1) = 2q_1$, $C_2(q_2) = 2q_2$ and $C_3(q_3) = 4q_3$.

For simplicity, assume the following: if there are multiple firms charging the same price and that price is the lowest in the market (i.e. it is the market price), then consumers go to the firm with the lowest marginal cost; if the firms charging the market price have the same marginal cost, then consumers are split evenly among the firms. Also, assume that firms cannot refuse to produce if consumers want to buy from it at the market price.

Answer the following questions:

1. Show that the market price must be 2 in any Nash equilibrium. (2 pts.)
2. If firm 1 has an opportunity to merge with firm 2 (the merged firm M's cost function is still $C_M(q_M) = 2q_M$), what is the maximum amount that it is willing to offer to firm 2, in order to promote this merger? Will the merger succeed? Assume that firm 3 will charge $P_3 = 4$ after the merger. (2 pts.)
3. If firm 1 has an opportunity to merge with both firm 2 and firm 3 (the merged firm M's cost function is still $C_M(q_M) = 2q_M$), what is the maximum amount that it is willing to offer to firm 2 and 3, in order to promote this merger? Will the merger succeed? (2 pts.)

4 Hotelling's Beach (6 pts.)

Consider a beach with length 1 (think about this as the interval $[0, 1]$ on the real line). Two ice-cream vendors, 1 and 2, are simultaneously picking their locations L_1 and L_2 on this beach. By choosing a location, each vendor wants to attract as many consumers as it can.

The consumers are uniformly distributed on the beach (think of any point on the interval $[0, 1]$ as one consumer), whose preference is as follows: each consumer will definitely buy one ice-cream, but the cost to a consumer depends on the price and the distance. In other words, suppose that vendor i charges P_i and the distance between a consumer and vendor i is d_i , then the consumer's cost of an ice-cream will be $P_i + d_i$. Each consumer will buy from the vendor that leads to a lower cost. If both vendors induce the same cost, then consumers are split evenly between the two. For example, if two vendors charge the same price and choose the same location, then they each get $\frac{1}{2}$ of the consumers.

Answer the following questions:

1. Assume that the vendors charge the same price. If vendor 1 chooses a location on $[0, \frac{1}{2})$, what would be vendor 2's best response? If vendor 1 chooses a location on $(\frac{1}{2}, 1]$, what would be vendor 2's best response? (2 pts.)
2. Using your answer in 1, show that the only Nash equilibrium is $L_1 = L_2 = \frac{1}{2}$. (Hint: first show that this is a NE; then show that there is no other NE.) (2 pts.)
3. Assume that the price of vendor 1 is higher than vendor 2's by 25 cents. Is $L_1 = L_2 = \frac{1}{2}$ still a Nash equilibrium? Briefly explain. (2 pts.)

5 Three-Firm Stackelberg Game (6 pts.)

Consider a market with demand function $P = 10 - Q$. Three firms, denoted 1, 2 and 3, are competing in the following Stackelberg game: first firm 1 chooses its quantity q_1 ; then firm 2 observes q_1 and chooses its quantity q_2 ; then firm 3 observes q_1 and q_2 , and chooses its quantity q_3 . The firms have identical costs: $C_i(q_i) = 2q_i$ for $i = 1, 2, 3$.

Answer the following questions:

1. What is firm 3's best response function? (2 pts.)
2. Given your answer to 1, what is firm 2's best response function? (2 pts.)
3. Given your answers to 1 and 2, what is firm 1's equilibrium quantity? (2 pts.)

ECON106P: Final Exam

Yangbo Song

July 30, 2014

1 Infinitely Repeated Bertrand Competition (8 pts.)

A total of n firms are Bertrand competitors in the market. The market demand function is given by $P = 20 - Q$, and each firm is producing at zero cost. They play the Bertrand game for infinitely many periods. Answer the following questions:

1. Suppose that they collude and operate as a monopoly, sharing production and profits equally. What profit would each firm make in each period? (2 pts.)
2. If one firm deviates from collusion, how much profit can it make at most (approximately) in the current period? (2 pts.)
3. Suppose that they do not collude and play the stage-game Nash equilibrium. What profit would each firm make in each period? (1 pt.)
4. Consider the grim trigger strategy. What is the smallest discount factor δ such that the grim trigger is a SPNE? As n goes to infinity, what is the limit of this lower bound on δ ? (3 pts.)

2 Partnership Problem (10 pts.)

Persons 1 and 2 are forming a firm. The value of their relationship depends on the effort that each expends simultaneously. Suppose that for $i = 1, 2$, person i 's utility from the relationship is $x_j^2 + x_j - x_i x_j$, where x_i is person i 's effort and x_j is the effort of the other person. Assume that $x_1, x_2 \geq 0$. Answer the following questions:

1. What is the symmetric Nash equilibrium of this game? (2 pt.)
2. Is the symmetric Nash equilibrium efficient? If yes, provide a brief argument; if no, give a counter example. (2 pts.)
3. Consider an infinitely repeated version of this game, and consider the following grim trigger for player i : start with $x_i = k > 0$ in the first period; in any other period, if both players have always expended effort k before, then continue to play $x_i = k$; otherwise, play $x_i = 0$ forever after. Let δ be the discount factor. Under what condition can the players sustain effort level k forever by the grim trigger in SPNE? (4 pt.)
4. Suppose that $\delta = \frac{1}{2}$. What is the highest level of effort that can be sustained with the grim trigger in SPNE? (2 pts.)

3 Reserve Price in Sealed-Bid Auction (6 pts.)

Consider the following sealed-bid auction between two bidders 1 and 2: the auctioneer announces a reserve price P . Each agent can only choose to quit the auction or bid weakly higher than P . If both bidders quit, the object still belongs to the auctioneer; if only one quits, the other gets the object and pays P ; if neither quits, the bidder with the higher bid gets the object and pays P . In the case of a tie, the winner is determined randomly, and the winner still pays P .

Assume that the bidders' valuations are i.i.d. with uniform distribution on $[0, 1]$. Also, assume that bidders can only choose their bids on the interval $[0, 1]$. Answer the following questions:

1. Assume that $P \in (0, 1)$. Find a weakly dominant strategy for each bidder. (Hint: the form of the strategy should be "quit if \dots , and bid \dots if \dots ". Keep in mind that a bidder only knows his private valuation; he cannot observe the valuation or action of his opponent.) (3 pts.)
2. Recall that in a first-price auction or a second-price auction with the same valuation structure, the auctioneer's expected revenue is $\frac{1}{3}$. Find a P (it does not have to be the optimal P for the auctioneer) such that the auctioneer has a higher expected revenue in this auction. (3 pts.)

4 First-Price Auction with Three Bidders (10 pts.)

Consider a first-price auction with three bidders 1, 2 and 3, whose valuations are i.i.d. with uniform distribution on the interval $[0, 1]$. Answer the following questions:

1. Suppose that player 2 is using the bidding function $b_2(v_2) = v_2$, and player 3 is using the bidding function $b_3(v_3) = v_3$. If player 1's valuation is v_1 and he bids b_1 , what is his expected payoff? (3 pts.)
2. Given your answer above, what is player 1's best response as a function of v_1 ? (3 pts.)
3. Disregard the assumptions above. If agents always use linear bidding functions, i.e. $b(v_i) = av_i$ where a is some constant, what is the bidding function in a symmetric Bayesian Nash equilibrium? (4 pts.)

5 Adverse Selection among Producers (8 pts.)

Consider the following principal-agent model: the principal is a consumer with utility function $u(q) = 2q^{\frac{1}{2}}$. The agent is a producer which can be of type A with probability $\frac{1}{2}$ and type B with probability $\frac{1}{2}$. A type A agent has a production function $C_A(q) = q$, and a type B agent has a production function $C_B(q) = 2q$. The principal can offer two contracts (q_A, T_A) and (q_B, T_B) , where q_A and q_B are quantities and T_A and T_B are payments. The principal's payoff is equal to utility minus payment; the agent's payoff is equal to payment minus cost.

Answer the following questions:

1. Write down the participation constraint and incentive compatibility constraint for each type of agent. (2 pts.)
2. Which of the participation constraints can be ignored? Provide a brief argument. (2 pts.)

3. After eliminating the participation constraint derived above, and ignoring the incentive compatibility constraint for type B , what is the principal's maximization problem (note that the principal is maximizing his expected payoff)? In your answer, rewrite the original constraints as binding constraints (i.e. constraints that are equations rather than inequalities), and provide a brief argument explaining why they should be binding. (2 pts.)
4. What is the optimal contract? (2 pts.)

6 Moral Hazard with Discrete Actions (8 pts.)

A firm is hiring a manager to work on a project. The outcome of the project depends on the manager's effort (either high or low) and a random factor. In particular, if the manager exerts high effort, which costs him 1, the project will succeed with probability $\frac{1}{2}$ and fail with probability $\frac{1}{2}$; if the manager exerts low effort, which costs him nothing, the project will fail for sure.

The firm can offer the following contract: pay the manager w for sure and in addition, if the project succeeds, pay him a bonus amount of b . The manager's utility from money is x^α where x is the amount of money received and $\alpha \in (0, 1)$ is some constant.

Answer the following questions:

1. Suppose that the firm wants the manager to exert high effort. What would be the manager's participation constraint and incentive compatibility constraint? (Hint: the manager is maximizing his expected payoff.) (3 pts.)
2. Suppose that the firm wants the manager to exert high effort. What is the firm's expected payment to the manager in terms of w and b ? (2 pts.)
3. Suppose that the firm wants the manager to exert high effort, and assume that $\alpha = \frac{1}{2}$. Given that the two constraints are binding at optimum, what is the contract that incurs the smallest expected payment for the firm? (3 pts.)