Diploma Lecture

The idea of this short catch up lecture is to recap some of the things Economics students will have covered in year 1 of their degree which are important for the EC202 course, but as diploma students, you might not have met yet. There are two main topics here that I cover: Section 1 on consumer theory, with subsections on budget constraint, preferences and utility maximisation. Section 2 is on Game Theory with subsections dominance, Nash Equilibrium and Cournot.

1 Consumer Theory

The framework here is that we have a consumer with a set of goods they can purchase and we aim to analyse what decisions they will take. In Section 1.1, the budget constraint defines which bundles of goods are affordable. In Section 1.2 we talk about preferences and in Section 1.3 we combine the previous two Sections to ask what actions a rational consumer would take. In all 3 sections, we pay particular focus to the 2 goods case. In reality there are often many goods, but many of the principles for more general analysis can also be seen in the 2 good case.

1.1 Budget constraint

Our framework is that we have J goods[1] which a consumer may want to consume and we denote a particular proposed bundle by $x = (x_1, x_2, \ldots, x_J)$ where for each $j \in J$, we let x_j denotes how many units of good j the consumer receives or buys. Typically, we also assume a price vector $p = (p_1, p_2, \ldots, p_J)$ where for each $j \in J$, we let p_j denote the price per unit of good j. We denote the amount of money the consumer has to spend, often called income or wealth by m. With this framework, we can define the budget constraint and some associated terms:

Definition 1.1. The budget constraint is

$$p_1x_1 + p_2x_2 + \cdots + p_jx_j \leq m$$

The logic for this is simple: p_jx_j is the amount of money spent on good j and summing across the j goods, we get the left hand side of the equation which is the amount of money the consumer spends. We require this to be less than or equal to the amount of money available, m, for the bundle to be affordable. Also note that there is no requirement that x_j be integer valued. That is we allow fractions of units of goods to be purchased. The primary reason for this is to make the maths easier.

Definition 1.2. The budget set is

$$\left\{x \in \mathbb{R}^J_{\geq 0} \mid p_1x_1 + p_2x_2 + \cdots + p_jx_j \leq m\right\}$$

In words, the budget set is the set of bundles that are affordable.

Definition 1.3. The budget line is

$$\left\{x \in \mathbb{R}^J_{\geq 0} \mid p_1x_1 + p_2x_2 + \cdots + p_jx_j = m\right\}$$

[1] In a slight abuse of notation but conforming to tradition, we let J be both the set of goods $J = \{1, 2, \cdots, J\}$ and also the set of goods.

In words, the budget set is the set of bundles that are just affordable (with no money to spare).

Exercise 1.1. Let $J = 2$. Draw a diagram of the budget set with x_1 on the horizontal axis and x_2 on the vertical. Use this to:

1. Verify that the budget line is a boundary of the budget set. What are the other boundaries?

2. Find the corners of the budget set and verify that the slope of the budget line is $-\frac{p_1}{p_2}$.

Throughout EC202, we will be making the simplifying assumptions made above that the price per unit of a good is constant regardless of how many units are purchased and that non-integer amounts of goods can be purchased. Although when one goes to a shop, they might often see this not being the case. We can still depict budget sets in these cases although it is more messy:

Exercise 1.2. You go to a shop selling apples and bananas at prices $p_a = 20$ and $p_b = 10$. You have an income of $m = 100$.

a) Allowing for the purchase of non-integer amounts, draw the budget set.

b) Stipulating only integer amounts of goods can be purchased, draw the budget set.

c) Still stipulating only integer amounts of goods can be purchased, suppose there is a buy one get one free offer on apples. Draw the budget set.

d) Can this budget set be approximated well by a straight line? How about as m increases?

1.2 Preferences

Key to understanding how consumers will act is to model their preferences - ie which bundles do they prefer to others? Given some set X of possible[2]

[2] normally we let $X = \mathbb{R}^J_{\geq 0}$ for some integer J and often set $J = 2$.

bundles of goods, we want to use a preference relation $\succeq$ to describe which we prefer to which others.

Definition 1.4. The preference relation $\succeq$ defines *weak preference* over bundles. For any two bundles of goods in X, $\hat{\mathbf{x}} = (\hat{x}_1, \dots \hat{x}_J)$ and $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_J)$, $\hat{\mathbf{x}} \succeq \bar{\mathbf{x}}$ translates as $\hat{\mathbf{x}}$ is weakly preferred to $\bar{\mathbf{x}}$.

We can then derive the strict preference relation and indifference relation from the weak preference relation:

Definition 1.5. Given $\succeq$ we define:

The *strict preference* relation: $\hat{\mathbf{x}} \succ \bar{\mathbf{x}} \iff \hat{\mathbf{x}} \succeq \bar{\mathbf{x}}$ and not $\hat{\mathbf{x}} \preceq \bar{\mathbf{x}}$.

The *indifference* relation: $\hat{\mathbf{x}} \sim \bar{\mathbf{x}} \iff \hat{\mathbf{x}} \succeq \bar{\mathbf{x}}$ and $\hat{\mathbf{x}} \preceq \bar{\mathbf{x}}$.

Two fundamental axioms on preferences are completeness and transitivity:

Definition 1.6. *Completeness*: For any two bundles of goods in X, $\hat{\mathbf{x}} = (\hat{x}_1, \dots \hat{x}_J)$ and $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_J)$, at least one of $\hat{\mathbf{x}} \succeq \bar{\mathbf{x}}$ or $\bar{\mathbf{x}} \succeq \hat{\mathbf{x}}$ must hold. (If both hold then $\bar{\mathbf{x}} \sim \hat{\mathbf{x}}$)

This says that for any two bundles in our consumption set, we can compare them and say which is better. In reality, this is not always easy for consumers to do. But for us economists, wanting to build a model we can analyse, this axiom is a must.

Definition 1.7. *Transitivity*: If we prefer bundle 1 to bundle 2 and bundle 2 to bundle 3, then we should also prefer bundle 1 to bundle 3. If $\mathbf{x}^{(\mathbf{1})} \succeq \mathbf{x}^{(\mathbf{2})}$ and $\mathbf{x}^{(\mathbf{2})} \succeq \mathbf{x}^{(\mathbf{3})}$ then $\mathbf{x}^{(\mathbf{1})} \succeq \mathbf{x}^{(\mathbf{3})}$.

If we have completeness and transitivity then we are able to order/rank the alternative bundles in X from best to worst. Note the necessity of both assumptions: without transitivity we could get cycles appear and so no clear best bundle, while without completeness, we might not be able to place some alternatives in our order/ranking. Typically in Economics, we call a consumer *rational* if they satisfy Completeness and Transitivity.

We usually make a further simplifying assumption: that preferences can be represented by a utility function:

Definition 1.8. The preference relation $\succeq$ can be represented by a utility function $u : X \to \mathbb{R}$ if for every pair of bundles $\bar{\mathbf{x}}, \hat{\mathbf{x}} \in X$, $\hat{\mathbf{x}} \succeq \bar{\mathbf{x}} \iff u(\hat{\mathbf{x}}) \geq u(\bar{\mathbf{x}})$.

This is not a completely inoccuous assumption: while every utility function yields rational preferences[3], the converse is not true. There are rational preferences which cannot be represented by any preference relation - as you will see in Lecture 1. The assumption though is often made for analytical convenience and there is a wide range of interesting preferences which are representable. Another two assumptions that are often made are: i) a monotonicity requirement meaning that for all goods, more is better and ii) a convexity requirement meaning that averages are better than extremes, but you will see this in more detail in Lecture 1.

Now we look at some common preferences and how we would represent them by a utility function (for simplicity we keep $J = 2$):

Exercise 1.3. For each of the following descriptions of preferences, find a utility function which represents them and draw indifference curves:

a) We only like good 1 and don't care at all about how much good 2 we get. The more good 1 we have the better.

b) We only like good 2 and don't care at all about how much good 1 we get. The more good 2 we have the better.

c) Perfect substitutes (equal ratio): we like both goods equally and can trade a unit of one for a unit of another and be equally well off.

d) Perfect substitutes (unequal ratio): we like both goods but good 1 more than good 2. For every unit of good 1 we sacrifice, we need two additional

[3] This is trivial as $\geq$ defines a complete and transitive ordering of $\mathbb{R}$

units of good 2 to be equally well off.

e) Perfect complements (equal ratio): one good is only useful to us in the presence of the other. We want to consume the two goods in equal quantities. Any surplus of either good makes us no better off. (Eg left and right shoes)

f) Perfect complements (unequal ratio): For every unit of good 1 we have, we need two units of good 2 to go with it. Any surplus of either good makes us no better off. (Eg coffee and sugar).

g) One commodity is a "bad": Commodity 1 makes us better off but commodity makes us worse off. For every additional unit of commodity 2 we get, we need an additional unit of commodity 1 to be equally well off.

Remark 1.1. Whenever preferences are representable by a utility function, there are infinitely many utility functions which represent these preferences (any increasing transformation).

There is one class of utility functions of particular interest:

Definition 1.9. Cobb-Douglas preferences are $u : X \to \mathbb{R}$ defined by $u(x_1, x_2) = x_1^{\alpha} x_2^{\beta}$. This is often written more simply as $u(x_1, x_2) = x_1^{\gamma} x_2^{1-\gamma}$, which we can obtain by taking a monotonic transformation (taking to the power of $\frac{1}{\alpha+\beta}$). Or by taking the natural logarithm, $u(x_1, x_2) = \alpha \ln x_1 + \beta \ln x_2$ where again often $\alpha + \beta = 1$. Note that $u(x_1, x_2) = x_1^{\alpha} x_2^{\beta}$ is defined over $X = \mathbb{R}^2_{\geq 0}$ while $u(x_1, x_2) = \alpha \ln x_1 + \beta \ln x_2$ is defined only over $X = \mathbb{R}^2_{>0}$. To deal with this, some authors might take preferences

$$u(x_1, x_2) = \begin{cases} \alpha \ln x_1 + \beta \ln x_2 & x_1 x_2 > 0 \\ -\infty & x_1 x_2 = 0 \end{cases}$$

Exercise 1.4. Consider Cobb-Douglas preferences with $\alpha = \beta = \frac{1}{2}$ and plot some indifference curves.

For differentiable utility functions like Cobb-Douglas we have the following:

Definition 1.10. The Marginal Utility of good i (MU_i) measures how much extra utility the consumer gets per unit increase of good i.

The Marginal Rate of Substitution (MRS) is the slope of the indifference curve. This measures how how many units of good 2 would be needed per unit decrease in good 1 to keep the consumer equally well off.

This can be calculated by

$$MRS_{1,2} = -\frac{\text{MU}_1}{\text{MU}_2} \text{ where } \text{MU}_i = \frac{\partial u_i}{\partial x_i}$$

Exercise 1.5. Find the MRS at the bundle $(x_1, x_2) = (2, 3)$ for the following utility functions:

a) $u(x_1, x_2) = 2x_1 + x_2$.
b) $u(x_1, x_2) = x_2$.
c) $u(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$.

1.3 Utility maximisation

Our consumer would like to find his most preferred bundle over the set of bundles that are affordable to him. From the above two sections, we know how to draw budget sets and indifference curves. So the aim here is to find the highest indifference curve that the conumer can reach subject to staying within his budget set.

Definition 1.11. The **optimal bundle(s)** is (are) the most preferred bundle within the consumer's budget set. When preferences are representable by a utility function $u : X \to \mathbb{R}$, this means maximising utility subject to budget constraint. This is known as the utility maximisation problem (UMP).

Our method will be to draw both the budget set and indifference curves and try to read off our diagram which bundle will be on the highest indifference curve while remaining within budget. Sometimes our diagram will show us the answer immediately; sometimes it will inform us of what mathematics we need to do to find the optimal bundle. We have lots of examples to practice with:

Exercise 1.6. Let budget be given by$p_1 = 20$ and $p_2 = 10$, $m = 100$. For each of the following preferences, draw indifference curves and find the optimal bundle(s).

a) $u(x_1, x_2) = x_1$.

b) $u(x_1, x_2) = x_2$.

c) $u(x_1, x_2) = x_1 + x_2$.

d) $u(x_1, x_2) = 2x_1 + x_2$.

e) $u(x_1, x_2) = \min\{x_1, x_2\}$.

f) $u(x_1, x_2) = \min\{2x_1, x_2\}$.

g) $u(x_1, x_2) = x_1 - x_2$.

For other utility functions like Cobb-Douglas, we might need slightly more complicated mathematics. Recall that the slope of the budget line is $-\frac{p_1}{p_2}$ and that the slope of the indifference curve can be found by calculating MRS.

Exercise 1.7. Let budget be given by$p_1 = 20$ and $p_2 = 10$, $m = 100$ and our consumer have the utility function $u(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$.

a) Sketch budget set and indifference curves. Use your diagram to argue that the optimal bundle is where the slope of the indifference curve equals the slope of the budget line.

b) Find the optimal bundle.

c) How much income does the consumer spend on each good?

d) How would your answer change if the utility function changed to $u(x_1, x_2) = x_1^{\frac{1}{3}} x_2^{\frac{2}{3}}$?

2 Game Theory

This knowledge will be useful in the latter parts of term 1 and in term 2.

In consumer theory we had an agent taking decisions, trying to maximise their utility without any outside interference. Game Theory, on the other hand, allows us to model situations where there are multiple decision makers and each person's utility is influenced both by their own actions and by the actions of others. As a simple example, in football, imagine a striker taking a penalty against a goalkeeper. The striker can kick the ball left or right and the goalkeeper can dive left or right. Now, how likely the striker is to score doesn't just depend on his own action, it also depends on the goalkeeper's action too.

Matrix representation

For 2 player games where each player has a small finite number of actions available, we can illustrate them in a matrix or table form, where we let player 1's actions be the rows and player 2's actions be the columns. Then each box of the table has 2 numbers: the left is player 1's utility and the right is player 2's utility:

Example 2.1. Player 1 is the striker and can aim left or right. If he aims left his shot will be on target 100% of the time, while if he aims right his shot will only be target 80% of the time. Meanwhile the goalkeeper, player 2, can dive to either the striker's left or right. If the goalkeeper dives the same way as the striker hit the ball then he saves it half the time, but if he dives the wrong way he never saves it. The striker's utility is the percentage probability that the penalty results in a goal; while the goalkeeper's utility is the percentage probability that it doesn't. We can represent this game by the table below:

		Player 2	
		Left	Right
Player 1	Left	50,50	100,0
	Right	20,80	40,60

In writing the game in this form, we typically assume that the two players move simultaneously, without knowing what the other one is doing. We can use game theory tools to analyse games where one player moves first but that is beyond the scope of this short lecture.

Utilities

The other thing to note is that utilities/payoffs represent all information about a player's preferences. So each player only cares about their own payoff and wants to maximise it. They are not interested in the payoff of their opponent. Perhaps in some situations people are altruistic and like to see the other player happy, or are jealous if an opponent gets a higher utility than themselves. If this is the case then these factors will have already been incorporated into the payoffs that you see written into the table. Also, we assume utilities are in "Bernouilli payoffs" meaning that any considerations of risk have already been accounted for in the utility function and when mixed strategies happen players only care about their expected payoff. This will make more sense when Christian Soegaard teaches you about risk aversion in term 2.

Our aim in game theory is, with this setup to analyse what rational, utility maximising players will do.

2.1 Dominance

Our first attempt at asnswering how rational players will play is dominance. The idea of dominance is to see if one strategy is unambiguously better than another, regardless of what other players do. Often what your best action is will depend upon what your opponent does, but in games where there is clearly a best action, it seems logical to think players will prefer to take that action. This seems a good starting point to predicting how games should be solved.

Definition 2.1. Strategy x **strictly dominates** strategy y if (for any opponent strategy) strategy x always gives a strictly higher payoff than strategy y.

In notation: for player i, for each strategy profile s_{-i} that other players could do, $u_i(x, s_{-i}) > u_i(y, s_{-i})$.

Note that we assume our player who chooses between x and y only cares about the effect on their own utility, not that on anybody else's. There is another weaker but still quite compelling form of dominance:

Definition 2.2. Strategy x **weakly dominates** strategy y if (for any opponent strategy) strategy x always gives a weakly higher payoff than strategy y and sometimes gives a strictly higher payoff.

In notation: for player i, for each strategy profile s_{-i} that other players could do, we have $u_i(x, s_{-i}) \geq u_i(y, s_{-i})$ and there exists some s'_{-i} for which $u_i\left(x, s'_{-i}\right) > u_i\left(y, s'_{-i}\right)$.

Remark 2.1. If x strictly dominates y then x weakly dominates y. However the converse is not true.

Example 2.2. Consider the following game:

		Player 2		
		Left	Middle	Right
Player 1	Top	2,2	3,3	3,2
	Bottom	1,4	2,-1	2,0

For player 1, Top strictly dominates Bottom. For player 2 Left weakly dominates Right. However there is no dominance between Left and Middle or Middle and Right.

Exercise 2.1. Consider the game below where we only look at player 1's payoff.

		Player 2	
		Left	Right
Player 1	Top	a,	b,
	Bottom	c,	d,

Complete the following sentences:

Top strictly dominates Bottom if and only if ...

Top weakly dominates Bottom if and only if ...

Next we move on to the idea of a dominant strategy. This is one that dominates all other strategies.

Definition 2.3. Strategy x is a **strictly dominant** strategy for a player if it strictly dominates all other strategies that player has.

Strategy x is a **weakly dominant** strategy for a player if it weakly dominates all other strategies that player has.

Example 2.3. Consider the following game:

		Player 2		
		Left	Middle	Right
Player 1	Top	2,2	3,3	3,3
	Bottom	1,4	2,5	2,0

For player 1 Top is strictly dominant as it strictly dominates the only other pure strategy for player 1.

For player 2 Middle is weakly dominant as it weakly dominates Left and Right (in fact it strictly dominates Left).

Allowing for mixing

The above analysis only considered what we call *pure strategies*. That is when a player does one action all the time. But we could also ask what if players mix between two or more actions. Recalling our comment that players only care about maximising their expected utility, we can apply the same defintions and allow for mixed strategies:

Example 2.4. Consider the game below where we only care about the payoffs of player 2:

		Player 2		
		Left	Middle	Right
Player 1	Top	,a	,b	,c
	Bottom	,d	,e	,f

The mixed strategy $pL+(1-p)\,M$ strictly dominates R iff $pa+(1-p)\,b > c$ and $pd+(1-p)\,e > f$.

The mixed strategy $pL+(1-p)\,M$ weakly dominates R iff $pa+(1-p)\,b \geq c$ and $pd+(1-p)\,e \geq f$, with at least one of the two inequalities being strict.

IESDS

IESDS stands for "iterative elimination of strictly dominated strategies". The idea here is that when a strategy is strictly dominated, that player will never play it. Furthermore, if both players are rational, then the opponent knows it will never be played. Given this, we can remove it from our analysis. And once we remove one strictly dominated strategy, it might allow us to remove more as the next example shows

Exercise 2.2. Consider the game below and argue that

		Player 2		
		X	Y	Z
Player 1	A	5,1	3,0	1,5
	B	1,2	2,3	0,0
	C	2,6	2,6	0,6

Apply IESDS to eliminate C, then X, then B, then Y, to leave (A,Z). At each stage write down the strictly dominating strategy you used.

When a game is solvable by IESDS, then the IESDS outcome seems a very good prediction for how the game should be played. However in many games IESDS fails to give us a unique outcome. Also you may be tempted to consider iteratively eliminating weakly dominated strategies, but this is generally considered not a good idea.[4]

2.2 Nash Equilibrium

Nash Equilibrium is the main solution concept in Game Theory upon which other more advanced solution concepts are built. The idea of Nash Equilibrium is to give an outcome which is stable in the sense that no player has an incentive to unilaterally deviate from it. Before defining Nash Equilibrium, we need to put some other things in place first:

A natural question for players to ask is:

> Given how the other player(s) are playing, what is the best strategy for me?

This is the idea of best response:

Definition 2.4. The **best response correspondence** of a player gives, for each possible strategy profile of other players, the set of strategies that maximise that player's utility.

Example 2.5. Consider the game below

		Player 2	
		Left	Right
Player 1	Top	2,5	1,3
	Bottom	4,1	4,2

Best responses for player 1: B against L. B against R.

Best reponses for player 2: L against T. R against B.

These are the best responses against the pure strategies. We could also find the best response to mixed strategies. For example the best response of player 1 against $\frac{1}{2}L + \frac{1}{2}R$ would be B.

[4] The reasons are: i) there are examples where you can eliminate weakly dominated strategies in one order to get one outcome and in a different order to get another outcome; ii) we might eliminate some Nash Equilibria.

Before defining Nash Equilibrium, we should explain what is meant by a strategy profile. A strategy profile is a description of what everyone is doing. So for each player, it prescribes a strategy. Often we use the notation σ_i for the strategy of player i. So in Example 2.5, an example of a strategy profile would be $(\sigma_1, \sigma_2) = (T, L)$. Or allowing for mixed strategies, a strategy profile would be something of the form

$$(\sigma_1, \sigma_2) = (pT + (1-p)\,B, qL + (1-q)\,R) \text{ where } p, q \in [0,1]$$

Definition 2.5. A **Nash Equilibrium** is a strategy profile such that no player has an incentive to change their strategy.

This means that each player is playing a best response to the strategy of the other player(s). In a 2 player game like in Example 2.5, this means that σ_1 is a best response to σ_2 and σ_2 is a best response to σ_1.

Example 2.6. Recall the game in example 2.5:

		Player 2	
		Left	Right
Player 1	Top	2,5	1,3
	Bottom	4,1	4,2

$(\sigma_1, \sigma_2) = (B, R)$ is a Nash Equilibrium. Player 1 has no incentive to deviate $(1 < 4)$ and player 2 has no incentive to deviate $(1 < 2)$.

$(\sigma_1, \sigma_2) = (T, L)$ is not a Nash Equilibrium. At this strategy profile player 2 best responds to player 1 $(3 < 5)$, but player 1 is not best responding to player 2 $(4 > 2)$.

Exercise 2.3. Consider the general 2 player game below:

		Player 2	
		Left	Right
Player 1	Top	a,b	c,d
	Bottom	e,f	g,h

Complete the following sentences:

$(\sigma_1, \sigma_2) = (T, L)$ is a Nash Equilibrium iff ...

$(\sigma_1, \sigma_2) = (T, L)$ is not a Nash Equilibrium iff ...

We can complete similar phrases for other strategy profiles too.

Finding pure strategy NE

When we have a 2 player game in matrix form, an easy way to find all the pure strategy Nash Equilibria is the underlining method.

Definition 2.6. The **underlining method** (for 2 player games) is as follows:

Go through each column, comparing the player 1 payoffs. For each column, underline the highest number (or numbers if there is a tie). This is the pure strategy best responses of player 1.

Go through each row, comparing the player 2 payoffs. For each row, underline the highest number (or numbers if there is a tie). This is the pure strategy best responses of player 2.

Then the pure strategy Nash Equilibria occur where both numbers are underlined.

Exercise 2.4. Consider the game below

		Player 2			
		W	X	Y	Z
Player 1	A	4,3	7,2	6,2	5,1
	B	3,4	6,6	5,9	2,4
	C	2,5	5,8	8,9	11,9
	D	1,5	3,0	6,1	7,7

Apply the underlining method to justify that there are 3 pure strategy Nash Equilibria: (A, W), (C, Y) and (C, Z).

Finding mixed strategy NE

I give an example to show how to do this in a simple 2 player, 2 action game. This gives the idea which can then be extended to more complex games.

Exercise 2.5. Consider the game below and let player 1 mix between his actions with probability p and player 2 mix with probability q as shown below:

		Player 2	
		Left (q)	Right ($1-q$)
Player 1	Top (p)	2,3	0,1
	Bottom ($1-p$)	1,0	3,2

Complete the following algebra for player 1:

$$u_1(T) \geq u_1(B) \iff$$
$$\iff$$
$$\iff$$

From this we can write the best response of player 1:

$$\mathrm{BR}_1(q) = \begin{cases} & q < \\ & q = \\ & q > \end{cases}$$

Complete the following algebra for player 2:

$$u_2(L) \geq u_2(R) \iff$$
$$\iff$$
$$\iff$$

From this we can write the best response of player 1:

$$\mathrm{BR}_2(p) = \begin{cases} & p < \\ & p = \\ & p > \end{cases}$$

From this we can deduce there are 3 NE in total:
i)
ii)
iii)

2.3 Cournot

In all of the games above, players have finitely many actions. Although this need not be the case. We can define and analyse games where the action space is infinite. A much studied example of this is the Cournot duopoly game:

Cournot setup

We have 2 firms in the market, producing an identical good to each other. Firm 1 has a marginal cost of c_1 per unit and firm 2 a marginal cost of c_2 per unit. There is a linear market demand meaning that we get price of $P = a - bQ$, for some constants $a, b \in \mathbb{R}$ when Q units are produced. We let the firms outputs be given by q_1 and q_2 so that $Q = q_1 + q_2$. Now each firm chooses their quantity trying to maximise the profit that they make. (Utility equals profit). I will solve this model for the simplest case: where $a = b = 1$ and $c_1 = c_2 = 0$

Monopoly analysis

Before solving this game, I first analyse what would happen if we only had one firm instead of two, or equivalently that one firm is forced to produce $q_2 = 0$. In the notation I use, we only have firm 1. Now this is a standard decision theory problem akin to what we studied in Section 1 as there is only one decision make and so no game theoretic aspect.

The firm wants to maximise profit and so solves:

$$\max_{q_1 \geq 0} \pi_1 = (P - c_1)\, q_1 = (1 - q_1)\, q_1$$

To solve this, we could use symmetry or differentiate and we find solution $q = \frac{1}{2}$, $P = \frac{1}{2}$, $\pi = \frac{1}{4}$.

Cournot dupoly analysis

Now we solve our model as intended with 2 firms. Note that this is a more complex situation as now for each firm, their best action will depend on what the other does. But the ideas are the same as we have seen above. For a given level of q_2, the objective of firm 1 is:

$$\max_{q_1 \geq 0} \pi_1 = (P - c_1)\, q_1 = (1 - q_1 - q_2)\, q_1$$

We solve by finding the first order condition:

$$\frac{d\pi_1}{dq_1} = 0 \Longleftrightarrow 1 - 2q_1 - q_2 = 0$$

From this we can write down the best response of firm 1 as a function of q_2 as:

$$\mathrm{BR}_1 : \quad q_1 = \begin{cases} \frac{1-q_2}{2} & q_2 \leq 1 \\ 0 & q_2 > 1 \end{cases}$$

Using a symmetric for firm 2, we can write down the best response of firm 2 as a function of q_1 as:

$$\mathrm{BR}_2: \quad q_2 = \begin{cases} \frac{1-q_1}{2} & q_1 \leq 1 \\ 0 & q_1 > 1 \end{cases}$$

The Nash Equilibrium occurs when each firm is best responding to the other. By inspection we see this cannot happen when either q_1 or $q_2 = 0$. Substituting BR_2 into BR_1 we get:

$$q_1 = \frac{1}{2} - \frac{1}{2}\left(\frac{1-q_1}{2}\right) \Longleftrightarrow q_1 = \frac{1}{3}$$

Substituting this back into BR_2 and our expressions for Price and profit we get:

$$q_1 = q_2 = \frac{1}{3}. \quad P = \frac{1}{3}. \quad \pi_1 = \pi_2 = \frac{1}{9}$$

Exercise 2.6. (For extra practice) Solve the model for different values of $a, b, c > 0$. Solve the general model, for arbitrary $a, b, c > 0$.