

EC201 — Summary of Each Topic

Christoph Walsh

August 8, 2017

1 The Budget Constraint

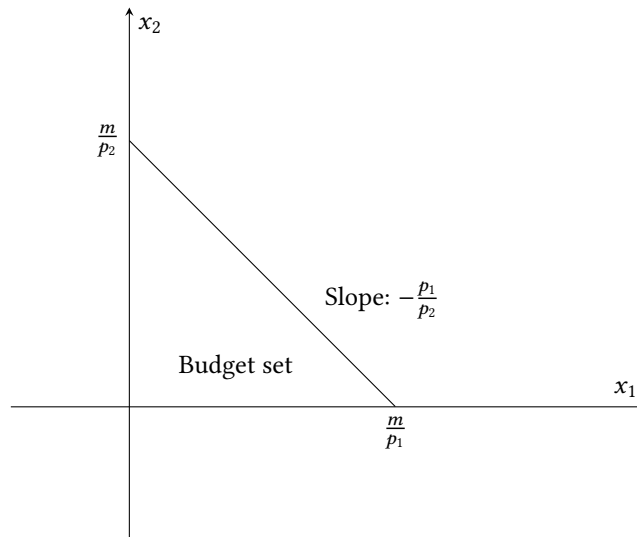
The *budget constraint* is that you cannot spend more on goods x_1 and x_2 than your income m :

$$p_1x_1 + p_2x_2 \leq m$$

The *budget line* is when you exhaust your entire income:

$$p_1x_1 + p_2x_2 = m$$

The vertical intercept of the budget line is at $\frac{m}{p_2}$ and the horizontal intercept is that $\frac{m}{p_1}$. Below is a graph of the budget line. The area below the budget line is the budget set, which is every bundle (x_1, x_2) the consumer can afford. The budget line represents the bundles the consumer can exactly afford (there is no money remaining).



The slope of the budget line is $-\frac{p_1}{p_2}$. This is the *opportunity cost* of consuming good 1. To consume one more unit of good 1 you need to give up $\frac{p_1}{p_2}$ units of good 2.

- Changes in income result in parallel shifts of the budget line (both the vertical and horizontal intercepts change).
- Changes in one of the prices result in pivots of the budget line (only one of the intercepts change).

2 Preferences

(x_1, x_2) and (y_1, y_2) are two different consumption bundles that differ in the amounts of goods 1 and 2. Consumers rank consumption bundles (x_1, x_2) and (y_1, y_2) .

- If $(x_1, x_2) \succ (y_1, y_2)$, the consumer *strictly prefers* (x_1, x_2) to (y_1, y_2) . The consumer will always choose (x_1, x_2) over (y_1, y_2) if both are available.
- If $(x_1, x_2) \sim (y_1, y_2)$, the consumer is *indifferent* between (x_1, x_2) and (y_1, y_2) .
- If $(x_1, x_2) \succeq (y_1, y_2)$, the consumer *weakly prefers* (x_1, x_2) to (y_1, y_2) . If we observe a consumer choosing (x_1, x_2) over (y_1, y_2) if both were available, then we know that $(x_1, x_2) \succeq (y_1, y_2)$.

Some assumptions we often make about preferences:

COMPLETENESS: Either $(x_1, x_2) \succeq (y_1, y_2)$ or $(y_1, y_2) \succeq (x_1, x_2)$ or both. You can always compare consumption bundles.

REFLEXIVITY: $(x_1, x_2) \succeq (x_1, x_2)$. A bundle is at least as good as itself.

TRANSITIVITY: If $(x_1, x_2) \succeq (y_1, y_2)$ and $(y_1, y_2) \succeq (z_1, z_2)$, then $(x_1, x_2) \succeq (z_1, z_2)$.

MONOTONICITY: If either

- $x_1 \geq y_1$ and $x_2 > y_2$; or
- $x_1 > y_1$ and $x_2 \geq y_2$

then $(x_1, x_2) \succ (y_1, y_2)$.

Under monotonicity, a consumer prefers a bundle that has at least as much of all goods with strictly more of one of the goods.

CONVEXITY: Convexity means that consumers prefer variety. For all weights $t \in [0, 1]$, if $(x_1, x_2) \sim (y_1, y_2)$, then

$$(tx_1 + (1-t)y_1, tx_2 + (1-t)y_2) \succeq (x_1, x_2)$$


Under convexity, a consumer would prefer the average of two extreme bundles to the extreme bundles themselves.

Indifference Curves:

- An *indifference curve* joins all the bundles that are equally preferred to each other.
- Indifference curves cannot cross.
- Indifference curves for perfect substitutes are straight lines.
- Indifference curves for perfect complements are *L-shaped*.

The Marginal Rate of Substitution (MRS):

- The slope of the indifference curve at a point is called the *marginal rate of substitution* (MRS).
- The MRS is how much of x_2 we need to give the consumer to keep them equally happy after taking away a small amount of x_1 .
- In general the MRS is negative so we will often take the absolute value when interpreting it.



- At A in the diagram above, the MRS is very negative. You do not have very much x_1 and have a lot of x_2 . You need to be given a lot of x_2 in order to be compensated for a loss in x_1 .
- At B in the diagram above, the MRS is almost zero. You do have a lot of x_1 do not have very much x_2 . You do not need to be given very much x_2 in order to be compensated for a loss in x_1 .

3 Utility

A *utility function*, u , is a way of assigning numbers to consumption bundles such that for every bundle where $(x_1, x_2) > (y_1, y_2)$ we have that $u(x_1, x_2) > u(y_1, y_2)$.

- The actual numbers the utility function give don't matter, only insofar that the numbers rank bundles consistent with the preferences. This concept is called *ordinal utility*.
- *Cardinal utility* is when the actual numbers associated with different utilities do matter. In this case we can say things like “bundle (x_1, x_2) is preferred *twice as much* to bundle (y_1, y_2) ”.
- An indifference curve is all x_1 and x_2 such that $u(x_1, x_2) = k$, where k is some constant number. For example, if $u(x_1, x_2) = x_1x_2$, then the indifference curve is $k = x_1x_2$, or $x_2 = \frac{k}{x_1}$.

Common utility functions:

Perfect substitutes	$u(x_1, x_2) = ax_1 + bx_2$
Perfect complements	$u(x_1, x_2) = \min\{ax_1, bx_2\}$
Cobb-Douglas	$u(x_1, x_2) = x_1^a x_2^b$

Marginal Utility:

- The *marginal utility* of a good is how much utility changes if we increase the amount of consumption of that good a tiny amount.
- It is the partial derivative of the utility function:

$$MU_1 = \frac{\partial u(x_1, x_2)}{\partial x_1} \quad MU_2 = \frac{\partial u(x_1, x_2)}{\partial x_2}$$

The marginal rate of substitution is equal to the (negative of the) ratio of marginal utilities:

$$MRS = -\frac{MU_1}{MU_2} = -\frac{\frac{\partial u(x_1, x_2)}{\partial x_1}}{\frac{\partial u(x_1, x_2)}{\partial x_2}}$$

To see where this comes from, remember that the *MRS* is the amount of x_2 we need to keep utility constant for a marginal change in x_1 . So two small changes dx_1 and dx_2 that keeps utility constant satisfies:

$$du = \frac{\partial u(x_1, x_2)}{\partial x_1} dx_1 + \frac{\partial u(x_1, x_2)}{\partial x_2} dx_2 = 0$$

Solving yields:

$$\frac{dx_2}{dx_1} = MRS = -\frac{\frac{\partial u(x_1, x_2)}{\partial x_1}}{\frac{\partial u(x_1, x_2)}{\partial x_2}} = -\frac{MU_1}{MU_2}$$

Marginal Utility and *MRS* Example with Cobb-Douglas Preferences:

Suppose $u(x_1, x_2) = x_1^a x_2^b$.

$$MU_1 = \frac{\partial u(x_1, x_2)}{\partial x_1} = ax_1^{a-1} x_2^b$$

$$MU_2 = \frac{\partial u(x_1, x_2)}{\partial x_2} = bx_1^a x_2^{b-1}$$

$$MRS = -\frac{MU_1}{MU_2} = -\frac{ax_1^{a-1} x_2^b}{bx_1^a x_2^{b-1}} = -\frac{ax_1^a x_1^{-1} x_2^b}{bx_1^a x_2^b x_2^{-1}} = -\frac{ax_1^{-1}}{bx_2^{-1}} = -\frac{ax_2}{bx_1}$$

4 Choice and Demand

Optimal Choice

- If a consumer's preferences satisfy monotonicity, the optimal bundle will always lie on the budget line and not beneath it. The consumer will spend all of their money.
- If a consumer's preferences satisfy monotonicity, then indifference curves further away from the origin give higher utility. The consumer will want to choose a bundle that places them on the indifference curve the furthest away that is still affordable. This is the indifference curve that "just touches" the budget line.
- If preferences are smooth and the optimal choice is not on the boundary, the optimal bundle will be where the indifference curve is tangent to the budget line.

- The optimal consumption bundle for the consumer is not necessarily where the budget line is at a tangent to the indifference curve. For example:
 - With perfect substitutes, if the prices are different you only buy the cheaper good. The optimal choice is at the boundary. The budget line will (in general) not line up with the tangent at the indifference curve.
 - With perfect complements, you choose the bundle at the kink of the indifference curves. The tangent is not well defined at the kink.
- With nonconvex preferences, it is possible that the point where the indifference curve is tangent to the budget line is suboptimal.

Demand Functions

- The *demand function* is a function that relates the consumer's optimal choice to prices and incomes. We write demand functions for goods 1 and 2 as:

$$x_1(p_1, p_2, m) \quad x_2(p_1, p_2, m)$$

- If we have perfect substitutes with a utility function $u(x_1, x_2) = x_1 + x_2$, the optimal choice is to buy the cheaper good. If the prices of the two goods are the same, buying either good or a combination of the two goods is optimal. The demand functions are then:

$x_1(p_1, p_2, m) = \begin{cases} \frac{m}{p_1} & \text{if } p_1 < p_2 \\ \text{any number between 0 and } \frac{m}{p_1} & \text{if } p_1 = p_2 \\ 0 & \text{if } p_1 > p_2 \end{cases}$	$x_2(p_1, p_2, m) = \begin{cases} 0 & \text{if } p_1 < p_2 \\ \text{any number between 0 and } \frac{m}{p_2} & \text{if } p_1 = p_2 \\ \frac{m}{p_2} & \text{if } p_1 > p_2 \end{cases}$
--	--

- If we have perfect complements with a utility function $u(x_1, x_2) = \min\{x_1, x_2\}$, the optimal bundle will involve the same amount of goods 1 and 2. This is at the kink of the indifference curve that “just touches” the budget line. Setting $x_1 = x_2$ and using this in the budget constraint:

$$p_1x_1 + p_2x_2 = m \implies p_1x_1 + p_2x_1 = m \implies x_1(p_1 + p_2) = m \implies x_1 = \frac{m}{p_1 + p_2}$$

The demand functions for goods 1 and 2 are then:

$x_1(p_1, p_2, m) = \frac{m}{p_1 + p_2}$	$x_2(p_1, p_2, m) = \frac{m}{p_1 + p_2}$
--	--

- If we have well-behaved preferences (smooth, monotonic, convex) and if the optimal choice is interior (a nonzero amount of all goods) it is where the indifference curve is tangent to the budget line. The slope of the tangent of the indifference curve is the *MRS*. The slope of the budget line is $-\frac{p_1}{p_2}$. Therefore an interior solution to the consumer optimization problem with well-behaved preferences satisfies:

$$\underbrace{MRS}_{\substack{\text{Slope of the} \\ \text{tangent of an} \\ \text{indifference curve} \\ \text{at a point}}} = \underbrace{-\frac{p_1}{p_2}}_{\substack{\text{Slope of the} \\ \text{budget line}}}$$

- With monotonic preferences, the consumer will always exhaust their income when choosing the optimal consumption bundle. Therefore:

$$p_1x_1 + p_2x_2 = m$$

- With these two optimality conditions, the demand functions $x_1(p_1, p_2, m)$ and $x_2(p_1, p_2, m)$ can be found by solving for x_1 and x_2 from these two equations:

$$\boxed{\frac{\frac{\partial u(x_1, x_2)}{\partial x_1}}{\frac{\partial u(x_1, x_2)}{\partial x_2}} = \frac{p_1}{p_2}} \quad \boxed{p_1x_1 + p_2x_2 = m}$$

- With Cobb-Douglas preferences, $u(x_1, x_2) = x_1^a x_2^b$, recall the MRS was $-\frac{ax_2}{bx_1}$. Setting $MRS = -\frac{p_1}{p_2}$ and solving for x_2 :

$$\frac{ax_2}{bx_1} = \frac{p_1}{p_2} \implies x_2 = \frac{bp_1x_1}{ap_2}$$

We can use this in the budget line equation $p_1x_1 + p_2x_2 = m$ to find the demand function for good 1:

$$p_1x_1 + p_2\left(\frac{bp_1x_1}{ap_2}\right) = m \implies p_1x_1 + \frac{bp_1x_1}{a} = m \implies p_1x_1\left(1 + \frac{b}{a}\right) = m \implies p_1x_1\frac{a+b}{a} = m \implies x_1 = \frac{a}{a+b} \frac{m}{p_1}$$

We can use this in the expression for x_2 we found above to find the demand function for good 2:

$$x_2 = \frac{bp_1x_1}{ap_2} = \frac{bp_1\left(\frac{a}{a+b} \frac{m}{p_1}\right)}{ap_2} = \frac{b\left(\frac{a}{a+b}\right) m}{ap_2} = \frac{b\left(\frac{1}{a+b}\right) m}{p_2} = \frac{b}{a+b} \frac{m}{p_2}$$

The demand functions are then:

$$x_1(p_1, p_2, m) = \frac{a}{a+b} \frac{m}{p_1} \quad x_2(p_1, p_2, m) = \frac{b}{a+b} \frac{m}{p_2}$$

Income Changes

- A good is a *normal good* if demand increases as income increases, i.e. $\frac{dx_1}{dm} > 0$.
- A good is an *inferior good* if demand decreases as income increases, i.e. $\frac{dx_1}{dm} < 0$.
- The *income offer curve* joins all the optimal consumption bundles for different levels of income m at a fixed pair of prices p_1 and p_2 .
- The *Engel curve* plots the demand for a good against income, m , at a fixed pair of prices p_1 and p_2 .

Own Price Changes

- If demand increases when the price of a good decreases, we call it an *ordinary good*.
- If demand decreases when the price of a good decreases, we call it a *Giffen good*.
- The *price offer curve* for good 1 joins all the optimal consumption bundles for different p_1 , holding m and p_2 fixed.
- The *demand curve* plots the demand for a good against its price, holding fixed income and the price of the other good.

Changes in the Price of the Other Good


- If the demand for good 1 goes up when the price of good 2 goes up, we say good 1 is a *substitute* for good 2. This happens when $\frac{dx_1}{dp_2} > 0$.
- If the demand for good 1 goes down when the price of good 2 goes up, we say good 1 is a *complement* for good 2. This happens when $\frac{dx_1}{dp_2} < 0$.
- If the demand for good 1 doesn't change when the price of good 2 goes up, good 1 is neither a substitute nor a complement to good 2. This happens when $\frac{dx_1}{dp_2} = 0$.

Income and Substitution Effects

- When the price of a good changes, the optimal consumption bundle often changes. For example, if the price of good 1 decreases, you may buy more of good 1 and less of good 2.
- Part of this change is because the relative prices between the two goods have changed: you substitute towards good 1 because the price of good 1 decreased. This is the *substitution effect*.
- The other part of this change is because with a lower price for good 1 you effectively have more income (your purchasing power increased). This is the *income effect*.

To decompose the income and substitution effects, we consider what choice the consumer would have made had they experienced the relative price change but were at the original purchasing power. The change in the choice they make only facing the relative price change and not the change in purchasing power is the substitution effect. The remaining effect is the income effect.


The substitution effect is always such that the consumer buys more when the price falls and buys less when the price rises. To see why this is true, consider the example below. We are initially at x^A . Suppose we pivoted the budget line around x^A (the blue line). This is what would happen if the price of good 1 fell and we brought the consumer's effective income back to the original bundle. x^C could never be optimal because x^C was available before the pivot while x^A is still available after the pivot. Therefore the optimal choice after the pivot must be to the right of x^A . Therefore for price decreases the substitution effect is always nonnegative ($SE \geq 0$) and for price increases the substitution effect is always nonpositive ($SE \leq 0$)



The income effect, on the other hand, can be positive or negative: it depends on the consumer's preferences. If good 1 normal, the income effect from a price decrease is positive. If good 1 is inferior, the income effect from a price decrease is negative.


Price decrease with a normal good:

We are initially consuming at x^A with the black budget line. After the price of good 1 decreases (from p_1 to p'_1), the budget line pivots outwards to $\frac{m}{p'_1}$ (the blue budget line). The new consumption bundle is x^C . To decompose the income and substitution effects, we draw a new budget line which has the same slope as the new budget line ($-\frac{p'_1}{p_2}$) but passes through the initial consumption plan x^A . This is the red dashed line. If this were the consumer's budget line, x^B would be chosen which involves more of x_1 . The substitution effect is changing from x_1^A to x_1^B and the income effect is changing from x_1^B to x_1^C . The income and substitution effects are both positive when the price decreases.




Price increase with a normal good:

We are initially consuming at x^A with the black budget line. After the price of good 1 increases (from p_1 to p'_1), the budget line pivots inwards to $\frac{m}{p'_1}$ (the blue budget line). The new consumption bundle is x^C . To decompose the income and substitution effects, we draw a new budget line which has the same slope as the new budget line ($-\frac{p'_1}{p_2}$) but passes through the initial consumption plan x^A . This is the red dashed line. If this were the consumer's budget line, x^B would be chosen which involves less of x_1 . The substitution effect is changing from x_1^A to x_1^B and the income effect is changing from x_1^B to x_1^C . The income and substitution effects are both negative when the price increases.




Price decrease with an inferior good (not Giffen):

We are initially consuming at x^A with the black budget line. After the price of good 1 decreases (from p_1 to p'_1), the budget line pivots outwards to $\frac{m}{p'_1}$ (the blue budget line). The new consumption bundle is x^C . To decompose the income and substitution effects, we draw a new budget line which has the same slope as the new budget line ($-\frac{p'_1}{p_2}$) but passes through the initial consumption plan x^A . This is the red dashed line. If this were the consumer's budget line, x^B would be chosen which involves more of x_1 . The substitution effect is changing from x^A to x^B and the income effect is changing from x^B to x^C . The substitution effect is positive and the income effect is negative. The income effect, however, is not enough to outweigh the substitution effect: you still buy more of good 1 after the price decreases.



Price decrease with a Giffen good:

We are initially consuming at x^A with the black budget line. After the price of good 1 decreases (from p_1 to p'_1), the budget line pivots outwards to $\frac{m}{p'_1}$ (the blue budget line). The new consumption bundle is x^C (where $x^C_1 < x^A_1$ as the good is Giffen). To decompose the income and substitution effects, we draw a new budget line which has the same slope as the new budget line ($-\frac{p'_1}{p_2}$) but passes through the initial consumption plan x^A . This is the red dashed line. If this were the consumer's budget line, x^B would be chosen which involves more of x_1 . The substitution effect is changing from x^A to x^B and the income effect is changing from x^B to x^C . The substitution effect is positive and the income effect is negative. The income effect is enough to outweigh the substitution effect such that you consume less of good 1 after the price falls.



5 Intertemporal Choice


This section is a variation on the consumer's choice problem. Instead of having two goods and deciding how much money to spend on each good, there is only one good and the consumer lives for two time periods and is deciding how much to borrow or save in the first period.

Call c_1 the consumption of the good in period 1 and c_2 the consumption of the good in period 2. Your income in period 1 is m_1 and your income in period 2 is m_2 . We call (m_1, m_2) your *endowment*. If you consume $c_1 = m_1$ and $c_2 = m_2$ (no borrowing or saving), we say you consume your endowment.

If there is no bank, you cannot borrow money and your savings do not earn any interest (you are stuffing cash in your mattress). If you save in period 1, your savings are $m_1 - c_1$. Your consumption in period 2 is then $m_2 + m_1 - c_1$. Your budget line is therefore

$$c_2 = \begin{cases} m_2 + m_1 - c_1 & \text{if } m_1 - c_1 > 0 \text{ (you save)} \\ m_2 & \text{if } m_1 - c_1 = 0 \text{ (you don't save)} \end{cases}$$

The budget line is:



Suppose now you can borrow or lend at an interest rate r in period 1. Save $m_1 - c_1$ in period 1 gives you $(1 + r)(m_1 - c_1)$ extra in period 2. Consumption in period 2 is therefore:

$$c_2 = m_2 + (1 + r)(m_1 - c_1)$$

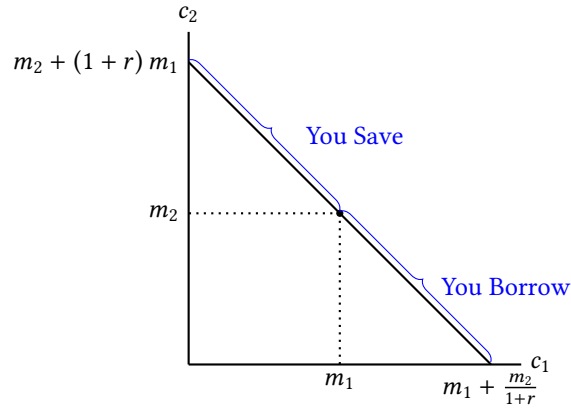
This is the budget constraint.

- If $m_1 > c_1$, you saved in period 1 and earn interest.
- If $m_1 < c_1$ you borrowed in period 1 and you pay interest.


We can rearrange the budget constraint to:

$$c_2 = m_2 + (1 + r)(m_1 - c_1) \quad \iff \quad \underbrace{c_1 + \frac{c_2}{1+r}}_{\text{Present value of lifetime consumption}} = \underbrace{m_1 + \frac{m_2}{1+r}}_{\text{Present value of lifetime income}}$$

Comparing to our old budget constraint, this is like $p_1 = 1$, $p_2 = \frac{1}{1+r}$ and $m = m_1 + \frac{m_2}{1+r}$.



An increase in the interest rate ($r \rightarrow r'$) pivots the budget line around the endowment (m_1, m_2) :



- If you are a *lender-type* person, you are made better off by an increase in the interest rate.
- If you are a *borrower-type* person, you are made worse off by an increase in the interest rate.

Inflation

With an inflation rate π the budget constraint becomes:

$$c_2 = m_2 + \frac{1+r}{1+\pi} (m_1 - c_1)$$

The real interest rate ρ is defined by:

$$1 + \rho = \frac{1+r}{1+\pi}$$

If the inflation rate isn't too high, we can use the approximation: $\rho \approx r - \pi$.

Valuing income streams

- The present value of income is all that matters for your intertemporal budget constraint as long as you are able to borrow and lend at the same interest rate: it doesn't matter if you get all your income in period 1 or 2 as long as the present value of income is the same.

- The present value of an income stream that pays x_1 in period 1 and x_2 in period 2 is $PV = x_1 + \frac{x_2}{1+r}$.
- The present value of an income stream that pays x forever (starting in period 2) is:

$$PV = \frac{x}{1+r} + \frac{x}{(1+r)^2} + \frac{x}{(1+r)^3} + \frac{x}{(1+r)^4} + \dots = \frac{x}{r}$$


6 Uncertainty

This section is another variation on the consumption problem. First we had two goods and the consumer was deciding how much to consume of each good given their budget constraint. Then we had one good and two time periods and the consumer was deciding how much to borrow or save given their income stream and the interest rate. In this section, there will be different states of the world that occur with different possibilities. The consumer is deciding how much insurance to buy to smooth consumption across those states.

Insurance

You have an asset worth A . In state 1 the asset receives a loss in value L . For example, the asset is a car and state 1 is a car crash which results in damage worth L to the car. In state 2, the asset retains its value A (there is no car crash). The probability of state 1 is π_1 and the probability of state 2 is π_2 .

Without insurance your consumption in state 1 is $c_1 = A - L$ and your consumption in state 2 is $c_2 = A$. You can purchase insurance that will pay K if the loss occurs for γK . γ can be interpreted as the cost per dollar of insurance. To fully insure, you set $K = L$. In this case, you get $A - \gamma L$ in both states, where γL is the cost of insurance. This is shown graphically below.



Expected Utility

How much a person will value consumption in a world with risk state depends on both the consumption in each state and the probability that each state will occur. Therefore the utility function will be a function of both the consumption levels and the probabilities: $u(c_1, c_2, \pi_1, \pi_2)$. One way to write this utility function is the weighted average of utility in each state, where the weights are the probabilities that those states occur. This is called *expected utility*:

$$u(c_1, c_2, \pi_1, \pi_2) = \pi_1 v(c_1) + \pi_2 v(c_2)$$

v is a function which gives the utility of consumption within a state. For example, it could be $v(x) = \sqrt{x}$. Expected utility would then be:

$$u(c_1, c_2, \pi_1, \pi_2) = \pi_1 \sqrt{c_1} + \pi_2 \sqrt{c_2}$$

In this context, the marginal rate of substitution of consumption between both states is:

$$MRS = -\frac{\frac{\partial u(c_1, c_2, \pi_1, \pi_2)}{\partial c_1}}{\frac{\partial u(c_1, c_2, \pi_1, \pi_2)}{\partial c_2}} = -\frac{\pi_1 \frac{\partial v(c_1)}{\partial c_1}}{\pi_2 \frac{\partial v(c_2)}{\partial c_2}}$$

Preferences for Risk

Suppose you have \$10 and someone offers you the following gamble based on a coin toss. If the coin lands on heads, you will give them \$5. If the coin lands on tails, they will give you \$5. Your expected value of wealth from not taking this gamble is \$10 (you just keep your money). Your expected value from taking the gamble is:

$$\pi_1 c_1 + \pi_2 c_2 = 0.5 \times \$15 + 0.5 \times \$5 = \$10$$

The “sure thing” is \$10 and the expected value of the gamble is \$10. Whether or not you might take this gamble depends on your preferences for risk. There are three main different types of preferences for risk:

- *Risk-averse* consumers prefer “sure things” to gambles that have the same expected value. Their utility functions are *concave*.
- *Risk-neutral* consumers are indifferent between “sure things” and gambles that have the same expected value. Their utility functions are *linear*.
- *Risk-loving* consumers prefer gambles to “sure things” that have the same expected value. Their utility functions are *convex*.

Going back to the coin toss example:

- An example of a risk-averse utility function is $v(x) = \sqrt{x}$. The expected utility from the “sure thing” is $0.5 \times \sqrt{10} + 0.5 \times \sqrt{10} = \sqrt{10} \approx 3.16$. The expected utility from the gamble is $0.5 \times \sqrt{5} + 0.5 \times \sqrt{15} \approx 3.05$. The risk-averse consumer prefers the sure thing.
- An example of a risk-neutral utility function is $v(x) = x$. The expected utility from the “sure thing” is 10. The expected utility from the gamble is $0.5 \times 5 + 0.5 \times 15 = 10$. The risk-neutral consumer is indifferent between the “sure thing” and the gamble.
- An example of a risk-loving utility function is $v(x) = x^2$. The expected utility from the “sure thing” is $10^2 = 100$. The expected utility from the gamble is $0.5 \times 5^2 + 0.5 \times 15^2 = 125$. The risk-loving consumer prefers the gamble to the “sure thing”.

Optimal Amount of Insurance

If the consumer is risk-averse, the optimal amount of insurance is where the marginal rate of substitution is equal to the slope of the budget line. Here, the slope of the budget line is $-\frac{\gamma}{1-\gamma}$ so:

$$\frac{\pi_1 \frac{\partial v(c_1)}{\partial c_1}}{\pi_2 \frac{\partial v(c_2)}{\partial c_2}} = \frac{\gamma}{1-\gamma}$$

If we assume the insurance company has no costs other than paying out when accidents occur, the profit for the insurance company is $Profits = \gamma K - \pi_1 K$. In a competitive insurance market, the insurance company breaks even ($Profits = 0$) so $\gamma = \pi_1$. When the insurance premium is the same as the probability of the loss occurring, we call this an *actuarially fair* premium. With an actuarially fair premium, the probabilities and γ terms cancel so:

$$\frac{\partial v(c_1)}{\partial c_1} = \frac{\partial v(c_2)}{\partial c_2}$$

In this case¹, $\frac{\partial v(c_1)}{\partial c_1} = \frac{\partial v(c_2)}{\partial c_2}$ implies that $c_1 = c_2$. Therefore if consumers are risk averse and premiums are *actuarially fair*, then they will optimally choose to fully insure.

7 Technology

- A firm's *production function* f shows how factors of production (inputs) are converted into output: $y = f(x_1, x_2)$.
- An *isoquant* shows the various combinations of inputs (x_1, x_2) that will yield the same amount of output y .
- The *monotonicity* assumption is that more inputs will give more output.
- The *convexity* assumption is that the average of two input combinations (x_1, x_2) and (x'_1, x'_2) will yield at least as much output.
- The *marginal product* of a factor is how much output changes in response to a marginal increase in the factor, holding the other factors of production fixed:

$$MP_1 = \frac{\partial f(x_1, x_2)}{\partial x_1} \quad MP_2 = \frac{\partial f(x_1, x_2)}{\partial x_2}$$

- The technical rate of substitution (*TRS*) is how much of factor 2 we need to produce the same amount of output y if we reduce factor 1 a marginal amount:

$$TRS = -\frac{MP_1}{MP_2} = -\frac{\frac{\partial f(x_1, x_2)}{\partial x_1}}{\frac{\partial f(x_1, x_2)}{\partial x_2}}$$

The *TRS* is the slope of an isoquant at a point.

- The *law of diminishing marginal product* is that the marginal product of a factor will fall as we increase the use of that factor, while holding other factors fixed.

Returns to Scale

For some number $t > 1$:

- If $tf(x_1, x_2) = f(tx_1, tx_2)$, then f exhibits constant returns to scale (CRS).
- If $tf(x_1, x_2) < f(tx_1, tx_2)$, then f exhibits increasing returns to scale (IRS).

¹This requires $\frac{\partial v}{\partial c}$ to be a strictly monotonic function. Since the consumer is risk averse, we know v is concave, which ensures that $\frac{\partial v}{\partial c}$ is a strictly decreasing function.

- If $tf(x_1, x_2) > f(tx_1, tx_2)$, then f exhibits decreasing returns to scale (DRS).

For Cobb-Douglas technology, $f(x_1, x_2) = x_1^a x_2^b$, scaling each input by t yields: $f(tx_1, tx_2) = (tx_1)^a (tx_2)^b = t^{a+b} f(x_1, x_2)$. If $a + b = 1$, there is CRS. If $a + b < 1$, there is DRS. If $a + b > 1$, there is IRS.

8 Profit Maximization

Short-Run Profit Maximization

Suppose the production function is $y = f(x_1, x_2)$, the price of output is p and the prices of each input/factor are w_1 and w_2 . The firm is competitive and therefore has no control over input and output prices. Factor 2 is fixed in the short run, $x_2 = \bar{x}_2$. The firm's maximization problem is:

$$\max_{x_1} pf(x_1, \bar{x}_2) - w_1 x_1 - w_2 \bar{x}_2$$

The firm chooses x_1 to maximize profits. This is done by taking the derivative of the above with respect to x_1 and setting the derivative equal to zero:

$$p \underbrace{\frac{\partial f(x_1, \bar{x}_2)}{\partial x_1}}_{=MP_1} - w_1 = 0$$

This gives:

$$\underbrace{pMP_1}_{\substack{\text{Value of the} \\ \text{marginal product} \\ \text{of factor 1}}} = \underbrace{w_1}_{\substack{\text{Price of} \\ \text{factor 1}}}$$

Long-Run Profit Maximization

In the long run, the firm chooses both x_1 and x_2 :

$$\max_{x_1, x_2} pf(x_1, x_2) - w_1 x_1 - w_2 x_2$$

The optimality conditions are:

$$p \underbrace{\frac{\partial f(x_1, x_2)}{\partial x_1}}_{=MP_1} - w_1 = 0$$

$$p \underbrace{\frac{\partial f(x_1, x_2)}{\partial x_2}}_{=MP_2} - w_2 = 0$$

These equations give the factor demand curves $x_1(p, w_1, w_2)$ and $x_2(p, w_1, w_2)$. The factor demand curves give the optimal choice of inputs given the output price and the input prices.

Profit Maximization and Returns to Scale

Suppose there was some choice of x_1 and x_2 that the firm earned a profit:

$$pf(x_1, x_2) - w_1 x_1 - w_2 x_2 > 0$$

If the production function exhibited constant returns to scale and we increased production by a factor of t :

$$pf(tx_1, tx_2) - w_1tx_1 - w_2tx_2 = t[\underbrace{pf(x_1, x_2) - w_1x_1 - w_2x_2}_{>0}]$$

The firm could choose t to be infinitely large and earn infinite profits. Therefore when we observe constant returns to scale, there are zero profits in equilibrium.

9 Cost Curves

- The cost function $c(y)$ gives the minimum cost of producing output level y .
- Costs can be decomposed into *variable costs*, denoted by $c_v(y)$, and *fixed costs*, denoted by FC :

$$c(y) = c_v(y) + FC$$

- *Variable costs* are costs that vary with output. If no output is produced, the variable costs are zero.
- *Fixed costs* are costs that do not depend on output. Fixed costs are paid even when the firm does not produce any output. If a firm does not produce any output, total costs equal fixed costs.
- An example cost function is $c(y) = y^2 + 1$. Here the fixed cost is $FC = 1$ (the constant) and the variable cost is $c_v(y) = y^2$.
- The *average cost function* measures cost per unit of output:


$$AC(y) = \frac{c(y)}{y}$$

- In the example cost function $c(y) = y^2 + 1$, the average cost is $AC(y) = \frac{y^2+1}{y} = y + \frac{1}{y}$.
- We also have average variable costs, $AVC(y)$, and average fixed costs, $AFC(y)$:

$$AC(y) = \frac{c(y)}{y} = \frac{c_v(y) + FC}{y} = \frac{c_v(y)}{y} + \frac{FC}{y} = AVC(y) + AFC(y)$$

Notice that while fixed costs do not depend on y , average fixed costs do depend on y . Average fixed costs always fall as y increases.

- In the example cost function $c(y) = y^2 + 1$, the average variable cost is $AVC(y) = \frac{y^2}{y} = y$ and the average fixed cost is $AFC(y) = \frac{1}{y}$.
- Average variable costs are always below average costs due to the presence of fixed costs (unless fixed costs are zero). The difference between $AC(y)$ and $AVC(y)$ falls as y gets large because $AFC(y)$ gets smaller.



- The *marginal cost* measures the change in costs for a small change in output.
- If we change output by one unit, the marginal cost is $c(y) - c(y - 1)$.
- For small changes, the marginal cost is the derivative of the cost function with respect to output:

$$MC(y) = \frac{dc(y)}{dy}$$

- In the example cost function $c(y) = y^2 + 1$, the marginal cost is $\frac{dc(y)}{dy} = 2y$.
- Since fixed costs do not vary with output, the marginal cost curve is also the derivative of the variable cost curve:

$$c(y) = c_v(y) + FC$$

Taking derivatives of both sides:

$$\frac{dc(y)}{dy} = \frac{dc_v(y)}{dy} + \underbrace{\frac{d(FC)}{dy}}_{=0} = \frac{dc_v(y)}{dy}$$

Therefore $MC(y) = \frac{dc(y)}{dy} = \frac{dc_v(y)}{dy}$

Relationship between Marginal Cost, Average Variable Cost and Average Cost

- If the average variable cost function is decreasing, then then marginal cost must be less than the average variable cost.
 - Suppose you measured the height of everyone in the room and took the average height. Then someone entered the room who was shorter than the average. If you take the average again, the average would be lower as the new person lowered the average.
 - To see this mathematically, we can use the quotient rule to find the derivative of the average variable cost function:

$$\frac{dAVC(y)}{dy} = \frac{d\left(\frac{c_v(y)}{y}\right)}{dy} = \frac{yc'_v(y) - c_v(y)}{y^2}$$


If the average variable cost is decreasing, then $\frac{dAVC(y)}{dy} < 0$. Therefore:

$$\begin{aligned} \frac{yc'_v(y) - c_v(y)}{y^2} &< 0 \\ yc'_v(y) - c_v(y) &< 0 \\ yc'_v(y) &< c_v(y) \\ \underbrace{c'_v(y)}_{MC(y)} &< \underbrace{\frac{c_v(y)}{y}}_{=AVC(y)} \end{aligned}$$

- Similarly, if the average variable cost function is increasing, then the marginal cost function must be greater than the average variable cost.
- The marginal cost curve is equal to the average variable cost curve when the average variable cost is at its minimum, i.e. when $\frac{dAVC(y)}{dy} = 0$. We can use the same steps as above to see this:

$$\begin{aligned} \frac{dAVC(y)}{dy} &= 0 \\ \frac{yc'_v(y) - c_v(y)}{y^2} &= 0 \\ yc'_v(y) - c_v(y) &= 0 \\ yc'_v(y) &= c_v(y) \\ \underbrace{c'_v(y)}_{MC(y)} &= \underbrace{\frac{c_v(y)}{y}}_{=AVC(y)} \end{aligned}$$

- We can also follow the almost the same steps to show that the marginal cost curve is equal to average total cost at its minimum.
- These relationships between marginal cost, average cost and average variable cost can be shown graphically as follows:



Relationship between Marginal Cost, Average Variable Cost and Average Cost

- The variable cost of producing y is the sum of marginal costs up to y .
- If output y is discrete, we can see this as follows:

$$\begin{aligned}
 c_v(y) &= c_v(0) - c_v(0) + c_v(1) - c_v(1) + c_v(2) - c_v(2) + \dots + c_v(y-1) - c_v(y-1) + c_v(y) \\
 &= [c_v(1) - c_v(0)] + [c_v(2) - c_v(1)] + \dots + [c_v(y) - c_v(y-1)] \\
 &= MC(1) + MC(2) + \dots + MC(y-1) + MC(y)
 \end{aligned}$$


- If output y is continuous, we can show this using calculus:

$$MC(y) = \frac{dc_v(y)}{dy}$$

Integrating both sides from 0 to y :

$$\begin{aligned}
 \int_0^y MC(x) dx &= \int_0^y \frac{dc_v(x)}{dx} dx \\
 &= c_v(y) - \underbrace{c_v(0)}_{=0} \\
 &= c_v(y)
 \end{aligned}$$

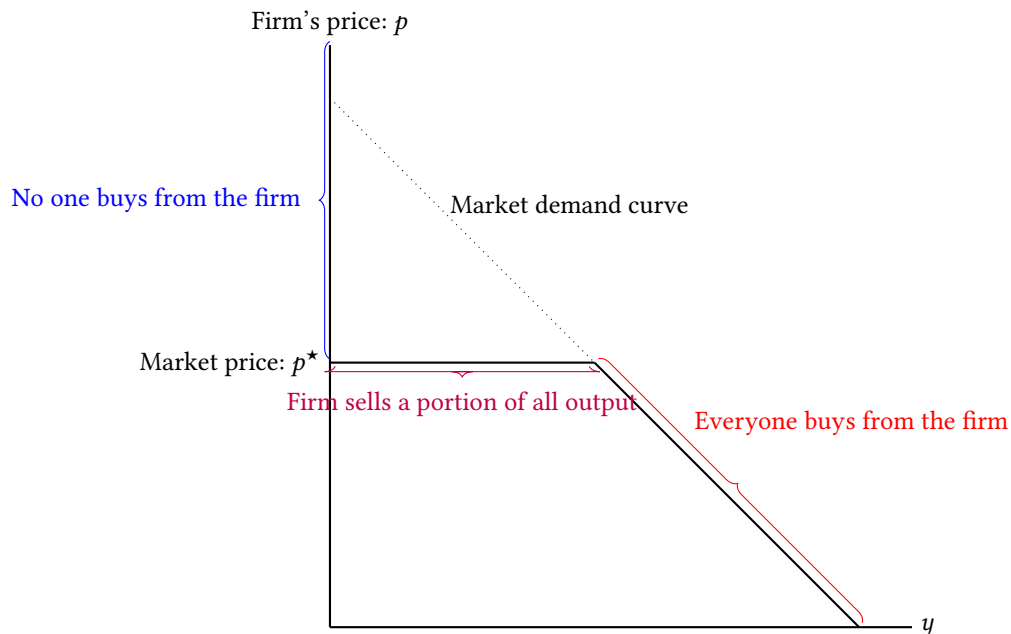
- What this means is that the area under the marginal cost curve left of a point \tilde{y} is the variable cost of producing \tilde{y} :



10 Firm Supply

- A market is *purely competitive* if the output of each individual firm is independent of the market price. Each firm is a *price taker*.
- If a firm charges higher than the market price, the firm will have zero demand for its output. If the firm charges lower than the market price, it will get the entire market demanding its output.

An individual firm's demand curve:



- The firm's maximization problem is:

$$\max_y \underbrace{py}_{\text{Revenue}} - \underbrace{c(y)}_{\text{Costs}} \quad \text{where } y \geq 0$$

- Taking the derivative and setting it equal to zero:

$$p - c'(y^*) = 0 \iff p = c'(y^*) \iff \underbrace{p}_{\text{Marginal revenue}} = \underbrace{MC(y^*)}_{\text{Marginal cost}}$$


- For example, if $c(y) = y^2 + 1$, the marginal cost is $MC(y) = 2y$. The optimal choice is to set $p = MC(y)$ so $p = 2y$ or $y = \frac{p}{2}$.
- The profits from producing nothing is $-F$. Your loss equals the fixed cost.
- If the profits from producing nothing is higher (less negative) than the profits from producing where $p = MC(y^*)$, then the firm should temporarily shut down.
- The firm should shut down when:

$$\begin{aligned}
 -F &> py^* - c_v(y^*) - F \\
 c_v(y^*) &> py^* \\
 \frac{c_v(y^*)}{y^*} &> p \\
 AVC(y^*) &> p
 \end{aligned}$$


- The supply curve of the firm, $S(p)$ gives the amount of output the firm will produce given price. In the short run, this is:

$$S(p) = \begin{cases} y^* \text{ where } p = MC(y^*) & \text{if } p \geq AVC(y^*) \\ 0 & \text{if } p < AVC(y^*) \end{cases}$$

- The supply curve of the firm is the marginal cost curve when the price exceeds the minimum of average variable cost (solid line) and is zero otherwise:



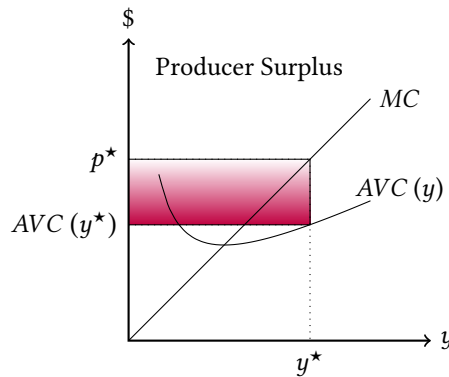
- Firm profits can be written as $py^* - c(y^*) = py^* - \frac{c(y^*)}{y^*}y^* = py^* - AC(y^*)y^* = [p - AC(y^*)]y^*$.
- Graphically, this is the area between the price and the average cost curve, up to the amount produced:



- *Producer surplus* equals revenues minus variable costs, $PS = py - c_v(y)$.
- Producer surplus can be written as:

$$PS = py - c_v(y) = py - AVC(y) \times y = [p - AVC(y)] y$$

- Graphically, this is the area between the price and the average variable cost curve:



- The firm will only produce in the long run if profits are nonnegative: $py - c(y) \geq 0$. This happens when price exceeds average cost:


$$py - c(y) \geq 0 \iff py \geq c(y) \iff p \geq \frac{c(y)}{y} \iff p \geq AC(y)$$

11 Industry Supply


- $S_i(p)$ is the short-run supply curve for firm i . The short-run *industry supply curve* when there are n firms is:

$$S(p) = S_1(p) + S_2(p) + \dots + S_n(p) = \sum_{i=1}^n S_i(p)$$

- The intersection of the industry supply curve and the market demand curve gives the equilibrium market price p^* .
- In the short run, firms set y^* such that $p^* = MC(y^*)$. In the short run, firms could either break even, make profits or suffer losses.



- In the long run, if firms are making losses firms will exit until firms break even. If firms are making profits, other firms will enter until firms break even. This happens when $p^* = \frac{c(y^*)}{y^*}$, where $\frac{c(y^*)}{y^*}$ is the minimum average cost.
- When firms enter (n increases), the industry supply curve moves to the right. When firms exit (n decreases), the industry supply curve moves to the left. In the long run the industry supply curve will intersect the market demand curve $D(p)$ at the point where the price equals the minimum of average total cost:



12 Equilibrium

- The equilibrium price of a good, p^* , is the price where the market supply of the good equals market demand. It is the p^* that solves $D(p^*) = S(p^*)$.

Example

Suppose we have a demand curve $D(p) = a - bp$ and supply curve $S(p) = c + dp$. To find the equilibrium price p^* ,

we set $D(p^*) = S(p^*)$:

$$\begin{aligned} a - bp^* &= c + dp^* \\ a - c &= bp^* + dp^* \\ a - c &= p^*(b + d) \\ p^* &= \frac{a - c}{b + d} \end{aligned}$$

To find the equilibrium quantity, use p^* in either the demand or supply curve. Suppose we used the demand curve:

$$q^* = a - bp^* = a - b \left(\frac{a - c}{b + d} \right) = \frac{a(b + d) - b(a - c)}{b + d} = \frac{ab + ad - ab + bc}{b + d} = \frac{ad + bc}{b + d}$$

- The inverse demand curve, $P_D(q)$, is the price demanders are willing to pay to consume a particular quantity.
- The inverse supply curve, $P_S(q)$, is the price suppliers must receive in order to supply a particular quantity.

Example

To find the inverse demand curve of the above example $D(p) = a - bp$, we simply solve for p :

$$\begin{aligned} q &= a - bp \\ bp &= a - q \\ p &= \frac{a}{b} - \frac{1}{b}q \\ P_D(q) &= \frac{a}{b} - \frac{1}{b}q \end{aligned}$$

We can do the same for supply to find $P_S(q) = \frac{q}{d} - \frac{c}{d}$.

- In equilibrium, $P_S(q^*) = P_D(q^*)$, where q^* is the quantity produced in equilibrium.
- Factors other than price that change demand or supply cause shifts in the demand and supply curves. When price changes we simply move along the curve.
- Changes in technology and input prices are some of the things that cause the supply curve to shift.
- Changes in preferences, prices of other goods and income are some the things that cause the demand curve to shift.

Comparative statics of demand and supply shifts:

If $S \uparrow$ (outward shift)	$p \downarrow$ and $q \uparrow$
If $S \downarrow$ (inward shift)	$p \uparrow$ and $q \downarrow$
If $D \uparrow$ (outward shift)	$p \uparrow$ and $q \uparrow$
If $D \downarrow$ (inward shift)	$p \downarrow$ and $q \downarrow$

Quantity Taxes

- A *quantity tax* is a tax levied per unit of quantity bought or sold. For example, there might be a quantity tax of \$3 on each packet of cigarettes.
- If there is a quantity tax t , then $P_D = P_S + t$, where P_D is the price consumers pay and P_S is the price sellers receive. The difference in the price consumers pay and the price suppliers receive is the quantity tax.

Example

Suppose again that $D(P_D) = a - bP_D$ and $S(P_S) = c + dP_S$. With a quantity tax, $P_D = P_S + t$. Substituting into the demand function give: $D(P_S) = a - b(P_S + t) = a - bP_S - bt$. In equilibrium we will have:

$$\begin{aligned}a - bP_S - bt &= c + dP_S \\-dP_S - bP_S &= bt + c - a \\dP_S + bP_S &= a - c - bt \\P_S(b + d) &= a - c - bt \\P_S &= \frac{a - c - bt}{b + d}\end{aligned}$$

Before the tax we had $p = \frac{a-c}{b+d}$ so P_S is lower now by $\frac{bt}{b+d}$. Sellers receive less for each unit sold. We can also calculate P_D :


$$P_D = P_S + t = \frac{a - c - bt}{b + d} + t = \frac{a - c - bt + t(b + d)}{b + d} = \frac{a - c - bt + bt + dt}{b + d} = \frac{a - c + dt}{b + d}$$

P_D is therefore higher by $\frac{dt}{b+d}$. Consumers pay more for each unit with the tax.

Value Taxes

- A *value tax* is a tax expressed in percentage units. The sales tax that we see is an example of a value tax. If the price is \$1 and the value tax is 5%, then the final price is \$1.05.
- If there is a value tax τ , then $P_D = (1 + \tau) P_S$.

Tax Incidence



- If the pre-tax equilibrium price was p^* , then the portion of the tax that consumers pay is:

$$\frac{P_D - p^*}{t}$$

The portion of the tax that sellers pay is:

$$\frac{p^* - P_S}{t}$$

Example

In the example above with the quantity tax, the portion of the tax that consumers paid is:

$$\frac{P_D - p^*}{t} = \frac{\frac{a-c+dt}{b+d} - \frac{a-c}{b+d}}{t} = \frac{\frac{dt}{b+d}}{t} = \frac{d}{b+d}$$

The portion that sellers paid is:

$$\frac{p^* - P_S}{t} = \frac{\frac{a-c}{b+d} - \frac{a-c-bt}{b+d}}{t} = \frac{\frac{bt}{b+d}}{t} = \frac{b}{b+d}$$


The portions of the tax depend on the slopes of the demand and supply curves. The steeper the slope, the larger portion of the tax you pay.

- The burden of the tax falls on the relatively less elastic side (buyers or sellers).
- If demand is perfectly inelastic, consumers pay all of the tax.
- If supply is perfectly inelastic, suppliers pay all of the tax.
- If demand is perfectly elastic, suppliers pay all of the tax.
- If supply is perfectly elastic, consumers pay all of the tax.

Pareto Efficiency

- An economic situation is *Pareto efficient* if there is no way to make any person better off without hurting anyone else.

- The market equilibrium without taxes is Pareto efficient.
- In the left diagram below, there is no tax. Consumer surplus is shown in blue and producer surplus is shown in red.
- In the right diagram below, there is a tax. The consumer surplus is the smaller blue triangle. The producer surplus is the smaller red triangle. The green rectangle is government revenue which is the tax per unit multiplied by the number of units sold $t \times q_T$. The remaining area (in gray) is the deadweight loss.



13 Demand Elasticities

Elasticity measures the responsiveness of demand. The slope of the demand curve is a potentially a good measure of responsiveness, but it depends on the units. For example, suppose you were interested in the demand for gasoline. If you change the units from gallons to liters, the elasticity would change. We should have a unit-free measure. Therefore we define the own-price elasticity of demand as:

$$\epsilon = \frac{\Delta q/q}{\Delta p/p}$$

If the change in price is small, we can use the derivative of demand with respect to price:

$$\epsilon = \frac{dq}{dp} \frac{p}{q}$$

If we are using an individual's demand function for good 1 $x_1(p_1, p_2, m)$, this is:

$$\epsilon = \frac{dx_1(p_1, p_2, m)}{dp_1} \frac{p_1}{x_1}$$

For most goods (non-Giffen goods), the elasticity of demand with respect to price will be negative. Therefore we will often talk about the elasticity in terms of absolute value, $|\epsilon|$.

If $|\epsilon| > 1$, demand is *elastic*. If $|\epsilon| < 1$, demand is *inelastic*. If $|\epsilon| = 1$, demand is *unit elastic*.

Elasticity can be higher or lower depending on the substitutability of a good. A good with many substitutes will be more elastic than one with fewer.


Elasticity and Revenue

Revenue is equal to price times quantity: $R = pq$. If the price changes from p to $p + \Delta p$, then the quantity demanded will change from q to $q + \Delta q$. Revenue will then be:

$$R + \Delta R = (p + \Delta p)(q + \Delta q) = \underbrace{pq}_{=R} + \Delta p q + p \Delta q + \underbrace{\Delta p \Delta q}_{\approx 0 \text{ if } \Delta p \text{ small}}$$

The change in revenue is therefore approximately:

$$\Delta R = \Delta p q + p \Delta q$$



At p and q , the revenue in the diagram is $C + D$. At the new price and quantity $p + \Delta p$ and $q + \Delta q$, the revenue is $A + C$. The change in revenue is therefore $A - D$. Here, $p \Delta q$ is the negative of box D and $\Delta p q$ is $A + B$, but $B \approx 0$ for small changes in price. So the change in revenue, $\Delta R = \Delta p q + p \Delta q$, is the box A minus the box D .

If we divide the change in revenue formula by Δp :

$$\frac{\Delta R}{\Delta p} = \frac{\Delta p q}{\Delta p} + p \frac{\Delta q}{\Delta p} = q + p \frac{\Delta q}{\Delta p} \frac{p}{p} = q \left(1 + \frac{\Delta q p}{\Delta p q} \right) = q(1 + \varepsilon) = q(1 - |\varepsilon|)$$

If $|\varepsilon| > 1$, then revenue decreases when price increases. If $|\varepsilon| < 1$, then revenue increases when price increases. If $|\varepsilon| = 1$, then revenue doesn't change when price increases.

Marginal Revenue

Marginal revenue is the change in revenue from producing a small increase in quantity. For price-takers (in perfect competition), this was always just the price, as their output didn't affect the market price. More generally, however, if a firm produces more output, this can affect the price of output. Dividing the change in revenue formula by the

change in quantity:

$$\frac{\Delta R}{\Delta q} = \frac{p\Delta q}{\Delta q} + q\frac{\Delta p}{\Delta q} = p + q\frac{\Delta p}{\Delta q} = p + \frac{q}{p}\frac{\Delta p}{\Delta q}p = p + \frac{1}{\varepsilon}p = p\left(1 + \frac{1}{\varepsilon}\right) = p\left(1 - \frac{1}{|\varepsilon|}\right)$$


If demand is elastic, $|\varepsilon| > 1$, then revenue increases when quantity increases. If demand is inelastic, $|\varepsilon| < 1$, then revenue decreases when quantity increases. If demand is unit elastic, $|\varepsilon| = 1$, then revenue doesn't change when quantity increases.

To see the above with derivatives instead of small changes, we can use the product rule on the revenue function expressed in quantities: $R(q) = p(q) \times q$:

$$\frac{dR(q)}{dq} = p(q) \times 1 + \frac{dp}{dq}q = p + \frac{dp}{dq}\frac{q}{p} = p\left(1 + \frac{dp}{dq}\frac{q}{p}\right) = p\left(1 + \frac{1}{\varepsilon}\right) = p\left(1 - \frac{1}{|\varepsilon|}\right)$$

Example with linear demand

- Suppose the inverse demand function is $p(q) = a - bq$.
- The revenue function is $R(q) = p(q) \times q = [a - bq]q = aq - bq^2$.
- Marginal revenue is then $MR(q) = \frac{dR(q)}{dq} = a - 2bq$.
- The inverse demand function meets the vertical axis at a and the horizontal axis at $\frac{a}{b}$.
- The marginal revenue function also meets the vertical axis at a and the horizontal axis at $\frac{a}{2b}$ – exactly half way towards the demand curve.
- $MR(q)$ is positive only when $|\varepsilon| > 1$.



Different Types of Demand Elasticities

So far we have been talking about the *own-price* elasticity of demand. However, there are other types of elasticities:

- The *cross-price elasticity* measures the sensitivity of the demand for good 1 when the price of good 2 changes (or vice versa). It is defined by $\varepsilon_{12} = \frac{dq_1}{dp_2} \frac{p_2}{q_1}$.
 - If $\varepsilon_{12} > 0$, then the goods are substitutes. If $\varepsilon_{12} < 0$, then the goods are complements.
- The *income elasticity* measures the sensitivity of demand with income: $\eta = \frac{dq}{dm} \frac{m}{q}$.
 - If $\eta > 0$, the good is normal. If $\eta < 0$, the good is inferior. If $\eta > 1$, the good is a luxury good.

Suppose we have a Cobb-Douglas utility function $u(x_1, x_2) = x_1^a x_2^b$, where as we have shown before the demand functions are $x_1(p_1, p_2, m) = \frac{a}{a+b} \frac{m}{p_1}$ and $x_2(p_1, p_2, m) = \frac{b}{a+b} \frac{m}{p_2}$. The own-price elasticity of demand for good 1 is:

$$\varepsilon_1 = \frac{\partial x_1}{\partial p_1} \frac{p_1}{x_1} = \underbrace{\left(-\left(\frac{a}{a+b}\right) \frac{m}{p_1^2}\right)}_{\frac{\partial x_1}{\partial p_1}} \frac{p_1}{x_1} = \left(-\left(\frac{a}{a+b}\right) \frac{m}{p_1^2}\right) \frac{p_1}{\frac{a}{a+b} \frac{m}{p_1}} = -1$$

The cross-price elasticity of demand for good 1 with respect to good 2 is zero, as the price of good 2 is not in the demand function for good 1.

The income elasticity of demand for good 1 is:

$$\eta_1 = \frac{\partial x_1}{\partial m} \frac{m}{x_1} = \underbrace{\left(\frac{a}{a+b} \frac{1}{p_1} \right)}_{\frac{\partial x_1}{\partial m}} \frac{m}{x_1} = \left(\frac{a}{a+b} \frac{1}{p_1} \right) \frac{m}{\frac{a}{a+b} \frac{m}{p_1}} = 1$$

14 Monopoly

In monopoly, there is only one firm. Unlike perfect competition, price is no longer given. When the monopolist is choosing the optimal quantity, it knows the effect of the quantity on the price.

The monopolist's problem is: $\max_q R(q) - c(q)$. The optimality condition is then $R'(q) = c'(q)$, or $MR(q) = MC(q)$.

The revenue function is $R(q) = p(q) \times q$. We could write the monopolist's problem as $\max_q p(q)q - c(q)$. At the optimal quantity:

$$p(q) + p'(q)q - c'(q) = 0$$

So $p(q) = c'(q) - p'(q)q$. In perfect competition, the firms set $p = MC(q)$. Here, the monopolist sets price equal to marginal cost *plus* a margin over marginal cost, where the margin is $-p'(q)q$.

From before, we saw that we can write marginal revenue as $p(q) \left[1 - \frac{1}{|\epsilon|} \right]$. So the monopolist charges a markup over marginal cost of:

$$p(q) = \underbrace{\left(\frac{1}{1 - \frac{1}{|\epsilon(q)|}} \right)}_{\text{Mark-up}} MC(q)$$

From the mark-up we can see that the monopolist will always set a price in the elastic part of the demand curve ($|\epsilon(q)| > 1$).

Example with linear inverse demand

The inverse demand function is $p(q) = a - bq$ and the cost function is cq . Marginal cost is therefore c . The profit for the monopolist is:

$$\pi(q) = p(q)q - c(q) = (a - bq)q - cq = aq - bq^2 - cq$$

Maximizing this function:

$$a - 2bq - c = 0 \implies 2bq = a - c \implies q = \frac{a - c}{2b}$$

The price is then

$$p(q) = a - bq = a - b \frac{a - c}{2b} = a - \frac{a - c}{2} = \frac{2a - a + c}{2} = \frac{a + c}{2}$$

The profit for the monopolist is:


$$\pi = pq - cq = [p - c]q = \left[\frac{a + c}{2} - c \right] \frac{a - c}{2b} = \left[\frac{a + c - 2c}{2} \right] \frac{a - c}{2b} = \left[\frac{a - c}{2} \right] \frac{a - c}{2b} = \frac{(a - c)^2}{4b}$$

Let's compare this to what happens under perfect competition. The price each firm charges is $p_{PC} = c$ and each firm makes zero profits, $\pi_{PC} = 0$. The total quantity produced in perfect competition is $q_{PC} = \frac{a-p}{b} = \frac{a-c}{b}$.

	Monopoly	Perfect Comp.
Quantity	$\frac{a-c}{2b}$	$\frac{a-c}{b}$
Price	$\frac{a+c}{2}$	c
Profit	$\frac{(a-c)^2}{4b}$	0

Graphically (on the left), the monopolist chooses the quantity where $MR(q) = MC(q)$. This happens at $\frac{a-c}{2b}$. At this quantity, the market will pay $\frac{a+c}{2}$. The consumer surplus is the area above the price and below the demand curve. The monopolist's profits here is the same as the producer surplus as there are no fixed costs. This is the rectangle of $(p - c)q$. There is a deadweight loss because there are consumers who were willing to buy the good at a price above marginal cost. The monopolist would like to sell to these consumers, but if it did it would have to drop the price for all other consumers.

On the right, we have what happens under perfect competition. The price is the same as marginal cost, $\frac{a-c}{b}$ units are bought, and consumer surplus is the entire red triangle.



Natural Monopolies

A *natural monopoly* occurs when fixed costs are very high and marginal costs are relatively small. Examples are gas delivery pipes, phone networks and subway networks. For industries like this it doesn't make sense to have firms competing with one another and duplicating the fixed cost. It is best to have a monopoly and regulate the price that it is charged (or just have the government run the firm).

For a natural monopoly, operating at $p = MC$ could result in losses to the firm. In this case, we might regulate the natural monopoly to charge at $p = AC$, which is higher than marginal cost so there will be some deadweight loss. Alternatively, we could force the monopolist to charge at $p = MC$ and subsidize the monopolist for the loss.

Monopolies are more likely when the minimum efficient scale (the level of output that minimizes average costs) is similar to the level of market demand. If the minimum efficient scale is smaller than market demand, then competition is more likely to occur.

15 Price Discrimination

For a standard monopolist, if it lowered the price from the optimal price, it could sell additional units. However, it would have to lower the price for all units sold, so it's not worthwhile. But what if the monopolist could charge different prices for each unit? When different people pay different prices, we call this price discrimination. For this to be possible, however, it should be difficult for consumers to sell on goods to each other.


There are three different types of price discrimination:

- *Third-degree price discrimination* is when different people pay different prices but each unit of output sold to a given person sells for the same price (e.g. student discounts for movie tickets).
- *Second-degree price discrimination* is when different units of output sell for different prices but all people face the same prices (e.g. bulk discounts, small/medium/large). This is also known as non-linear pricing.
- *First-degree price discrimination* (also known as *perfect price discrimination*) is when the monopolist sells different units of output for different prices and those prices can differ from person to person. Each person pays their willingness to pay for the products. There is no deadweight loss and no consumer surplus.

First-Degree Price Discrimination

First degree price discrimination is more of a theoretical concept than something that exists in reality. Suppose the marginal cost of producing a good is \$2. One individual may be willing to pay \$5 for the first unit, and \$3 for the second unit and \$1 for the third unit. Another individual may be willing to pay \$3 for the first unit and \$1 for the second unit. What a perfect price discriminating monopolist would do is charge the first person \$8 for 2 units and the second person \$3 for one unit.

Under perfect price discrimination, each consumer pays their willingness to pay for each unit, as long as their willingness to pay exceeds marginal cost. Therefore there is no consumer surplus as the consumer surplus for each unit is the willingness to pay minus the price. The producer surplus, on the other hand, is very large as for each unit the producer surplus is $p - MC$. Therefore the producer surplus is the whole triangle above marginal cost, if demand is linear:



Since there is never a case where the willingness to pay exceeds the marginal cost when the product isn't bought and sold, there is no deadweight loss. A situation with a perfect-price discriminating monopolist is Pareto efficient (although not fair on the consumers). The reason this doesn't exist in reality is because it's very hard to ascertain someone's willingness to pay.

Second-Degree Price Discrimination

Second-degree price discrimination is when different units of output sell for different prices but all people face the same prices (e.g. bulk discounts). The monopolist is not able to give different people different prices, as they cannot tell them apart. An example is business class and economy class in an airline. The monopolist could in principle charge people wearing suits and carrying briefcases more than people wearing shorts and sandals, but after a while the businesspeople will start to wear shorts and pack their suits in their suitcases. What the airline can do, however, is have different packages (high quality and low quality) and get the travelers to *self-select* into the correct package.

Example:

There are two types of flight passengers, equal in number: business (B) and vacationers (V) and there are no costs. The monopolist cannot charge different people different prices but will offer different packages (price-quality pairs) to each customer. The monopolist chooses the price-quality pairs such that the different groups *self-select* into the price-quality pair designed for them.

Each triangle in the diagram is the same size and is equal to one for simplicity. The line V below measures a vacationer's willingness to pay for an additional unit of quality and the line B measures a businessperson's willingness to pay for an additional unit of quality. At q_L below, a vacationer would pay up to $B + E + F = 3$ for a plane ticket. At the same quality, a businessperson would pay $A + B + C + E + F = 5$.

The monopolist can choose between three qualities q_L , q_M and q_H and give each a price. The goal for the monopolist is to choose two qualities and two prices for those qualities to maximize profits. If the monopolist was able to tell the two consumers apart, it would like to sell to the vacationers a medium-quality package q_M at the price of their full willingness to pay for it: $B + E + F + G = 4$. It would sell a high-quality package q_H to the businesspeople at their full willingness to pay for it: $A + B + C + D + E + F + G + H + I = 9$. So it makes 4 from the vacationers and 9 from the businesspeople in this ideal case (13 overall).


But if the monopolist can't tell them apart, what the businesspeople will do is dress in shorts and sandals and buy the medium-quality package q_M for $B + E + F + G = 4$. They will get a positive consumer surplus from this equal to the area below their willingness-to-pay curve minus the price: $A + C + D + H = 4$. The monopolist will only ever sell the q_M package and make $B + E + F + G = 4$ from every consumer (4 from vacationers and 4 from businesspeople).

One strategy the monopolist can do is not to charge so much for q_H . If they left the businesspeople with the same surplus as if they bought the q_M package, they will be indifferent between the two packages (and we can assume that when indifferent they choose the higher-quality package). The monopolist would need to charge $B + E + F + G + I = 5$ instead of the full triangle. This leaves the businesspeople with a surplus of $A + C + D + H = 4$, the same as what they would get from buying q_M . The monopolist gets profits $B + E + F + G = 4$ from the vacationers and $B + E + F + G + H = 5$ from the businesspeople (4 from vacationers and 5 from businesspeople).

But is this the best the monopolist can do? Economy class in airlines has been getting worse and worse so you might guess the optimal thing to do is to sell the vacationers a really crappy package. Suppose you made two packages, q_L and q_H . We want the vacationers to buy q_L and the businesspeople to buy q_H . We will charge the vacationers their full willingness to pay, which is $B + E + F = 3$. If a businessperson bought this package, they would get surplus $A + C = 2$. Therefore when choosing the price for the high quality package q_H , the monopolist needs to ensure that the businesspeople get at least a surplus $A + C = 2$ so they don't buy the low quality package. So it will charge $B + D + E + F + G + H + I = 7$.

The profits for the monopolist is $B + E + F = 3$ from vacationers and $B + D + E + F + G + H + I = 7$ from businesspeople. Therefore with q_L and q_H the monopolist can make $3+7=10$ and with q_M and q_H it was only able to make $4+5=9$. However, this is less than the ideal case where the monopolist can tell consumers apart where it

can make 13. Compared to the perfect price discrimination case, it loses 1 from the vacationers (G) and 2 from the businesspeople (A + C).



So the monopolist is purposefully making economy class really bad to stop the high willingness to pay consumers from being tempted by it, at the detriment of the low willingness to pay consumers.

Third-degree price discrimination

Third-degree price discrimination is when different people pay different prices but each unit of output sold to a given person sells for the same price (e.g. student discounts for movie tickets).

For example, let $p_1(q_1)$ and $p_2(q_2)$ be the inverse demand curves for two types of people (e.g. professionals and students). If the monopolist can price discriminate the monopolist's problem is:

$$\max_{\{q_1, q_2\}} p_1(q_1)q_1 + p_2(q_2)q_2 - c(q_1 + q_2)$$

The first-order conditions are $MR_1(q_1) = MC(q_1 + q_2)$ and $MR_2(q_2) = MC(q_1 + q_2)$. These are two equations with two unknowns (q_1 and q_2). The monopolist should choose quantities to make the marginal revenue from both groups the same: $MR_1(q_1) = MR_2(q_2)$.

Suppose the monopolist optimally chose to charge a higher price to group 1. Then since $MR = p(q) \left[1 - \frac{1}{|\epsilon(q)|} \right]$:

$$p_1(q_1) > p_2(q_2) \implies \left[1 - \frac{1}{|\epsilon(q_1)|} \right] < \left[1 - \frac{1}{|\epsilon(q_2)|} \right] \implies \frac{1}{|\epsilon(q_1)|} < \frac{1}{|\epsilon(q_2)|} \implies |\epsilon(q_2)| > |\epsilon(q_1)|$$

The monopolist will charge a higher price to the *less elastic* group.

Example

$D_1(p_1) = 100 - p_1$ and $D_2(p_2) = 100 - 2p_2$. Marginal cost is 20. The inverse demand curves are $p_1(q_1) = 100 - q_1$ and $p_2(q_2) = 50 - \frac{q_2}{2}$.

The marginal revenue functions are $MR_1(q_1) = 100 - 2q_1$ and $MR_2(q_2) = 50 - q_2$. Setting $MR_1 = MC$ gives $100 - 2q_1 = 20$ so $q_1 = 40$. Setting $MR_2 = MC$ gives $50 - q_2 = 20$ so $q_2 = 30$. The prices are then $p_1 = 60$ and $p_2 = 35$.

Profit is then $60 \times 40 + 35 \times 30 = \$3,450$.

Suppose you couldn't price discriminate? The total demand is then $D(p) = D_1(p) + D_2(p) = 100 - p + 100 - 2p = 200 - 3p$. The inverse demand is then $p(q) = \frac{200}{3} - \frac{q}{3}$. Marginal revenue is $MR(q) = \frac{200}{3} - \frac{2q}{3}$. Setting $MR(q) = MC(q)$ gives $\frac{200}{3} - \frac{2q}{3} = 20$ so $q = 70$. Therefore price is $p = \frac{200}{3} - \frac{70}{3} = \frac{130}{3} = 43\frac{1}{3}$. Profits are then $70 \times 43\frac{1}{3} = \$3,033.33$, smaller than under price discrimination.

Bundling

If different people value different products differently, monopolists are also able to *bundle* products together to increase profits. Suppose there are two types of consumers, journalists and accountants, equal in number. There are two desktop applications, Word and Excel. Journalists and accountants need both software but journalists find Word more useful and accountants find Excel more useful. Their willingness to pay for two desktop applications are:

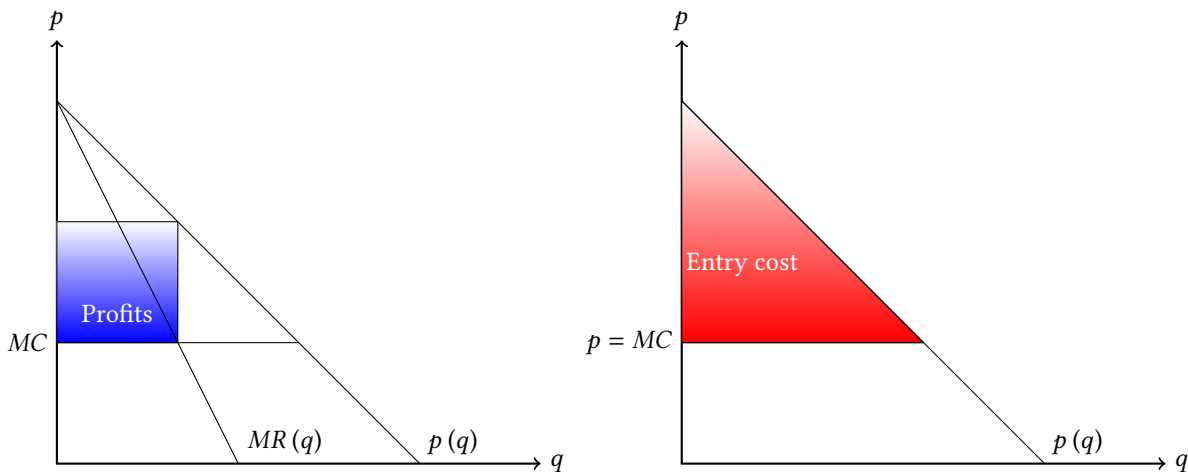
	Word	Excel
Journalist	120	100
Accountant	100	120

The monopolist's marginal cost is zero (in reality, the marginal cost of software is actually zero).

Charging 120 for each would exclude half the market so charging 100 for each is better. However, the monopolist can do even better than that. What if they bundled the two applications into an "Office Suite" for \$220? Both consumers would still be willing to buy the bundle but now they get zero surplus.

Two-Part Tariffs

There is a theme park that charges for entry and then also charges for each ride. Suppose everyone has the same demand curve for rides at the theme park. If the monopolist only charged for rides, it would earn profits as it normally does and the profits is the blue box on the left diagram. However, suppose the monopolist charged a price for rides equal to marginal cost. The consumer surplus would be the whole red triangle on the right diagram. What the monopolist can do is just set the entry price equal to that consumer surplus and steal the entire consumer surplus. If the marginal cost is close to zero, you may see places in real life following this strategy where there is only an entry cost and the price for a ride is zero.



16 Game Theory

Game theory studies the strategic interaction between players

The Normal Form Representation of a Game (Payoff Matrix)

- A *payoff matrix* shows the payoffs for each player given their action and the action of the other player.

		<i>Firm B</i>	
		High price	Low price
<i>Firm A</i>	High price	5, 5	0, 10
	Low price	10, 0	1, 1

- The left number is the payoff for the left player, the right number is the payoff for the top player.

Strategies

- A *strategy* is what a player will do given every action of the other player. Your strategy is like the instructions you would give to a robot playing the game for you.
- A player has a *dominant strategy* if one action is optimal for that player no matter what the other player does. In the example above, playing “Low price” is a *dominant strategy* for both players.
- We say a particular strategy of a player is a *dominated strategy* if one action gives a higher payoff than the dominated strategy, no matter what the other player does. In the example above, “High price” is dominated for both players.

Nash Equilibrium

- We want to be able to predict what the outcome of a game is.
- If both players have a dominant strategy, we would predict that that will be the outcome of the game (in the example above, both firms play “Low price”).
- However, not every game has dominant strategies. Consider the following example where two people had agreed to meet at a classical music concert but both of their phones died. Going alone to a concert gives them zero payoff, but conditional on going together Andy prefers Bach and Betty prefers Stravinsky:

		<i>Betty</i>	
		Bach	Stravinsky
<i>Andy</i>	Bach	2, 1	0, 0
	Stravinsky	0, 0	1, 2

- In the Bach and Stravinsky example, both player’s optimal strategies is to go to the same concert as the other person. For example, Andy’s best response to Betty playing “Bach” is also to play “Bach” and Andy’s best response to Betty playing “Stravinsky” is also to play “Stravinsky”. We can mark each player’s *best response* to the other player’s action with a ★:

		<i>Betty</i>	
		Bach	Stravinsky
<i>Andy</i>	Bach	$2^*, 1^*$	0, 0
	Stravinsky	0, 0	$1^*, 2^*$

- If two players' best responses overlap (one square has two \star s), then we call it a *Nash Equilibrium*. More generally, a pair of strategies is a *Nash Equilibrium* if both players' choices are optimal given the other player's choice. No player has any incentive to unilaterally deviate from a Nash Equilibrium. In the example above (*Bach, Bach*) and (*Stravinsky, Stravinsky*) are two Nash equilibria of the game.
- The best response functions don't always overlap. Consider the following example of a striker and a goalkeeper during a penalty shootout in football. If the goalkeeper jumps in the correct direction, he will save the ball. If not, the striker scores (keep in mind, since the two players are facing each other, left and right are the opposite location for each). The best responses don't overlap in any case here:

		<i>Goalkeeper</i>	
		Jump Left	Jump Right
<i>Striker</i>	Shoot Left	$1^*, -1$	-1, 1^*
	Shoot Right	-1, 1^*	$1^*, -1$

- As an aside, a game in which the total payoffs in each outcome sum to zero is called a *zero-sum game*.

Mixed Strategies

A *pure strategy* is when the agent makes one choice and sticks to it. A *mixed strategy* is when the agent randomizes between two or more strategies.

In the penalty kick example, there was no pure strategy Nash equilibrium (PSNE). However, there is a mixed strategy Nash equilibrium (MSNE).

Suppose the striker played "Left" with probability p and "Right" with probability $1 - p$. Similarly suppose the goalkeeper player "Left" with probability q and "Right" with probability $1 - q$. If the striker is optimally playing a mixed strategy, then it must be the case that given the goalkeeper's randomization, the striker doesn't prefer one strategy over another in expectation. Therefore the expected value from either strategy must be the same, given the other player's randomization. So:

$$\underbrace{1 \times q + (-1) \times (1 - q)}_{\text{Expected payoff from shooting "Left"}} = \underbrace{(-1) \times q + 1 \times (1 - q)}_{\text{Expected payoff from shooting "Right"}}$$

Solving for q :

$$q - (1 - q) = -q + 1 - q \implies 2q - 1 = 1 - 2q \implies 4q = 2 \implies q = \frac{1}{2}$$

So the goalkeeper must be randomizing with probability $\frac{1}{2}$ for the striker to be indifferent between shooting left and right. This makes sense, if the goalkeeper was playing "Left" more often than "Right", then the striker would shoot "Left" all the time and score more often. Similarly, if the goalkeeper was playing "Right" more often than "Left", then the striker would shoot "Right" all the time and score more often.

We can do the same for the goalkeeper:

$$\underbrace{(-1) \times p + 1 \times (1 - p)}_{\text{Expected payoff from jumping "Left"}} = \underbrace{1 \times p + (-1) \times (1 - p)}_{\text{Expected payoff from jumping "Right"}}$$

Solving for p yields $p = \frac{1}{2}$ like above. Both players flip a coin and shoot and jump randomly left and right in equilibrium. Both players get an expected payoff of 0. The MSNE is $p = \frac{1}{2}$ and $q = \frac{1}{2}$.

Bach or Stravinsky example: The Bach or Stravinsky example also has a MSNE. Andy goes to Bach with probability p and Betty goes to Bach with probability q . Andy is indifferent between Bach and Stravinsky if:

$$2q + 0 \times (1 - q) = 0 \times q + 1 \times (1 - q)$$


Solving yields $q = \frac{1}{3}$. Betty is indifferent between Bach and Stravinsky if:

$$1 \times p + 0 \times (1 - p) = 0 \times p + 2(1 - p)$$

Solving yield $p = \frac{2}{3}$. So there is a MSNE with $p = \frac{2}{3}$ and $q = \frac{1}{3}$.

Sequential Games

In the previous examples, both players moved at the same time. The goalkeeper and the striker had to choose which direction to jump and shoot at the same time. We will now consider what happens when players take turns. Consider the following example. Player 1 moves first and can go either L or R . Player 2 can then do either A or B :




This is the *extensive form* of a game. It is a representation of the game showing the order in which the players move. We can also write the game in *normal form* like we did before:

		<i>Player 2</i>	
		A	B
<i>Player 1</i>	L	$1^*, 9^*$	$1, 9^*$
	R	$0, 0$	$2^*, 1^*$

This game has two pure strategy Nash equilibria (L, A) and (R, B) . But does (L, A) make any sense if Player 1 moves first? If Player 1 moves first, he can avoid getting the payoff of 1 by moving R and thereby forcing player 2 to play B (as B gives Player 2 a payoff of 1 and A gives player 2 a payoff of 0). Since the definition of Nash equilibrium can lead us to unreasonable predictions in sequential games we define a new concept of equilibrium called *Subgame Perfect Nash Equilibria* (SPNE).


We can find SPNE of extensive form games using *backward induction*. First you start at the end of the game and mark off what the optimal thing a player would do if the game ended up there. If Player 1 played L , Player 2 is

indifferent between A and B . But in any case, the payoffs are the same. If Player 1 played R , Player 2 would play B . We can therefore delete the node where Player 2 plays A after Player 1 plays R :




Player 1 will now choose R as 2 is greater than 1. The SPNE of this game is (R, B) .


Consider another example:




We can solve it via backward induction. What will Player 1 do at the end of the game? It will play Y at the left node (8 is better than 3) and X at the right node (4 is better than 2). We can highlight those actions and delete the other actions:



What will player 2 do now? At the very left it will choose C over D as 4 is greater than 2. On the right it will choose B over A as 6 is better than 5 (note that we deleted the other possibilities as we know player 1 won't play them).



Now we are at the beginning of the game. What will Player 1 do? If it plays Z it will get 6. If it plays W it will get 4. So it will play Z as 6 is better than 4:



In the SPNE of this game, player 1 plays Z at the beginning and player 2 plays C.

17 Oligopoly

Oligopoly comes from Ancient Greek: *olgos* means “few” and *polein* means sell. Therefore it means “few sellers”. In oligopoly there are only a few firms so the actions of each firm affects the action of every other firm. For example, if one firm decreases price, the other firms might react by also decreasing price. For simplicity, we will only consider the 2-firm case (the algebra gets cumbersome for more firms and the intuition stays the same). We will study 3 different models: simultaneous price setting (Bertrand), simultaneous quantity setting (Cournot) and Staggered quantity setting (Stackelberg). There is another model that studies staggered price setting, but for time reasons we skip that model (it’s in the book if you are interested).

Simultaneous Price Setting (Bertrand)

If each firm has the same constant marginal cost, the equilibrium in Bertrand competition is for both firms to set $p = MC$. If a firm set $p > MC$, the other firm would undercut it. If a firm set $p < MC$ it would make a loss.

Simultaneous Quantity Setting (Cournot)

Each firm's problem is:

$$\max_{q_1} p(q_1 + q_2) q_1 - c(q_1) \qquad \max_{q_2} p(q_1 + q_2) q_2 - c(q_2)$$

The optimal choice of quantity depends on the quantity of the other firm: $q_1^* = q_1(q_2)$ and $q_2^* = q_2(q_1)$. These are the firms' *reaction functions*. These are the best response functions for each firm. In equilibrium, both reaction curves must be satisfied. Let's take a look at an example with linear demand and costs.

With $p(q) = a - bq$ and $c(q) = cq$, the profit for firm 1 is:

$$\pi_1(q_1, q_2) = [a - b(q_1 + q_2)] q_1 - cq_1 = aq_1 - bq_1^2 - bq_1q_2 - cq_1 = (a - c)q_1 - bq_1^2 - bq_1q_2$$

For firm 2:

$$\pi_2(q_1, q_2) = [a - b(q_1 + q_2)] q_2 - cq_2 = aq_2 - bq_2^2 - bq_1q_2 - cq_2 = (a - c)q_2 - bq_2^2 - bq_1q_2$$

Firm 1's reaction curve is q_1 as a function of q_2 where $\frac{d\pi_1(q_1)}{dq_1} = 0$, i.e. profits are maximized:

$$a - c - 2bq_1 - bq_2 = 0 \implies q_1(q_2) = \frac{a - c - bq_2}{2b}$$

Firm 2's reaction curve is q_2 as a function of q_1 where $\frac{d\pi_2(q_2)}{dq_2} = 0$:

$$a - c - 2bq_2 - bq_1 = 0 \implies q_2(q_1) = \frac{a - c - bq_1}{2b}$$

In equilibrium the reaction curves will intersect. We can put $q_2(q_1)$ into $q_1(q_2)$:

$$\begin{aligned} q_1 &= \frac{1}{2b} \left[a - c - \frac{1}{2} (a - c - bq_1) \right] \\ q_1 &= \frac{a - c}{2b} - \frac{a - c}{4b} + \frac{q_1}{4} \\ q_1 - \frac{q_1}{4} &= \frac{a - c}{4b} \\ \frac{3q_1}{4} &= \frac{a - c}{4b} \\ q_1 &= \frac{a - c}{3b} \end{aligned}$$


Similarly, $q_2 = \frac{a-c}{3b}$. The total quantity produced is $q = q_1 + q_2 = \frac{a-c}{3b} + \frac{a-c}{3b} = \frac{2}{3} \frac{a-c}{b}$.

The profit for firm 1:

$$\begin{aligned}
 \pi_1(q_1, q_2) &= (a - c)q_1 - bq_1^2 - bq_1q_2 \\
 &= (a - c)\frac{a - c}{3b} - b\left(\frac{a - c}{3b}\right)^2 - b\left(\frac{a - c}{3b}\right)^2 \\
 &= \frac{(a - c)^2}{3b} - 2b\frac{(a - c)^2}{9b^2} \\
 &= 3\frac{(a - c)^2}{9b} - 2\frac{(a - c)^2}{9b} \\
 &= \frac{(a - c)^2}{9b}
 \end{aligned}$$

The profit for firm 2 is also $\frac{(a-c)^2}{9b}$.

Graphically, the two reaction curves are as follows:



Firm 1's reaction curve $q_1(q_2)$ gives the optimal amount to produce in response to any output of firm 2. If firm 2 produces zero output, firm 1 will optimally react by producing $\frac{a-c}{2b}$, the monopolist output. If firm 2 produces the perfectly competitive output $\frac{a-c}{b}$, then firm 1 will optimally react by producing nothing (as any increase in quantity will lead the price to be lower than marginal cost). If firm 2 produces $\frac{a-c}{3b}$, then firm 1 will optimally react by producing $\frac{a-c}{3b}$ as well. In equilibrium, both firms won't want to react to each other's output, which is at $\frac{a-c}{3b}$. Since neither firm wants to deviate at this out, it is a Nash equilibrium.


A firm's *isoprofit curve* shows the different quantity combinations for both firms (q_1, q_2) that give the same amount of profits. Firm 1's profits are:

$$\pi_1(q_1, q_2) = (a - c)q_1 - bq_1^2 - bq_1q_2$$

At a fixed profit level $\bar{\pi}$, we can solve for q_2 to find the isoprofit function:

$$\begin{aligned}\bar{\pi} &= (a - c) q_1 - b q_1^2 - b q_1 q_2 \\ b q_1 q_2 &= (a - c) q_1 - b q_1^2 - \bar{\pi} \\ q_2 &= \frac{a - c}{b} - q_1 - \frac{\bar{\pi}}{b q_1}\end{aligned}$$

Graphically, the isoprofit curves for firm 1 look as follows:



Isoprofit curves closer to the monopoly profits of $\frac{a-c}{2b}$ give higher profits. For any level of output of firm 2, firm 1 will produce an output to put it on the isoprofit curve with the higher profits. This point is also where the reaction function is: the reaction curve gives the quantity firm 1 should produce to get on the best isoprofit curve, given any output for firm 2. The reaction function joins the maximum of all the isoprofit curves. More specifically, the reaction curves are the loci of maxima of the isoprofit curves for each firm. Firm 2's isoprofit curves are defined in a similar way.

Quantity Leadership (Stackelberg)

Here firm 1 produces its quantity first. Then firm 2 observes this quantity and decides how much to produce. Firm 2 will produce according to the same reaction function as in Cournot, $q_2(q_1)$. When firm 1 decides on a quantity to produce, it will internalize firm 2's reaction in its decision as it knows exactly how much firm 2 will produce in response:

$$\max_{q_1} p(q_1 + q_2(q_1)) q_1 - c(q_1)$$

In this case, the leading firm can take advantage of moving first. Let's return to the linear demand and costs example. Firm 1 knows that firm 2's reaction function is:

$$q_2(q_1) = \frac{a - c - b q_1}{2b}$$

Therefore we can put this directly in firm 1's profit function:

$$\begin{aligned}\pi_1(q_1, q_2(q_1)) &= [a - b(q_1 + q_2(q_1))]q_1 - cq_1 \\ &= \left[a - b \left(q_1 + \frac{a - c - bq_1}{2b} \right) \right] q_1 - cq_1 \\ &= aq_1 - bq_1^2 - \frac{a - c}{2}q_1 + \frac{bq_1^2}{2} - cq_1 \\ &= \frac{a - c}{2}q_1 - \frac{bq_1^2}{2}\end{aligned}$$


Firm 1 will maximize this profit function by taking the derivative with respect to q_1 and setting it equal to zero:

$$\frac{d\pi_1}{dq_1} = \frac{a - c}{2} - bq_1 = 0 \implies q_1 = \frac{a - c}{2b}$$

This turns out to be the monopoly quantity. Firm 2 will respond to this by producing according to its reaction function:

$$q_2(q_1) = \frac{a - c - bq_1}{2b} = \frac{a - c - b\frac{a - c}{2b}}{2b} = \frac{\frac{a - c}{2}}{2b} = \frac{a - c}{4b}$$

Graphically, firm 1 will choose the quantity on firm 2's reaction function which puts it on its highest isoprofit curve. This will be the isoprofit curve that "just touches" firm 2's reaction curve:



The total quantity produced in Stackelberg oligopoly is $\frac{a-c}{2b} + \frac{a-c}{4b} = \frac{3(a-c)}{4b}$, which is higher than under Cournot. We can compare each of the models' total outputs with perfect competition and monopoly:

- In perfect competition, total industry output is $\frac{a-c}{b}$.
- In monopoly, the monopolist produces $\frac{1}{2} \frac{a-c}{b}$.
- In Cournot, each firm produces $\frac{1}{3} \frac{a-c}{b}$ so in total we have $\frac{2}{3} \frac{a-c}{b}$.
- In Stackelberg, the leader produces $\frac{1}{2} \frac{a-c}{b}$ and the follower produces $\frac{1}{4} \frac{a-c}{b}$ so in total we have $\frac{3}{4} \frac{a-c}{b}$.

18 Oligopoly Collusion

Collusion is when competing firms decide to coordinate to maximize their joint profits. If firms in Cournot competition want to collude, they should jointly maximize the sum of profits:

$$\max_{q_1, q_2} p(q_1 + q_2)q_1 - c(q_1) + p(q_1 + q_2)q_2 - c(q_2)$$

Let's use the linear demand and cost example again from before. Here, the firms' joint maximization problem is:

$$\max_{q_1, q_2} [a - b(q_1 + q_2)]q_1 - cq_1 + [a - b(q_1 + q_2)]q_2 - cq_2$$

Using $q = q_1 + q_2$, this is:

$$\pi(q) = (a - bq)q - cq = (a - c)q - bq^2$$

This is the same as the monopoly problem. The firms will jointly produce $q = \frac{a-c}{2b}$ and will jointly earn the monopoly profits. Each will individually produce $q_1 = q_2 = \frac{a-c}{4b}$, half the monopoly output, and each firm will earn profits $\pi_{cartel} = \frac{(a-c)^2}{8b}$, half the monopoly profits.

In Cournot competition, each firm earns profits $\pi_{Cournot} = \frac{(a-c)^2}{9b}$, so profits under collusion are higher.

Cheating on the Cartel

A firm will be tempted to cheat on the cartel because it is producing a quantity that is not on its reaction function. It can achieve a higher profit by producing the quantity according to its reaction function, provided the other firm doesn't change its output.

In the linear demand and costs example, firm 1's reaction function is $q_1(q_2) = \frac{1}{2b}(a - c - bq_2)$. If firm 2 is producing the cartel quantity, $q_2 = \frac{a-c}{4b}$, firm 1 will want to cheat and produce:

$$q_1 = \frac{1}{2b}(a - c - bq_2) = \frac{1}{2b}\left(a - c - b\left(\frac{a-c}{4b}\right)\right) = \frac{1}{2b}\left(3\left(\frac{a-c}{4}\right)\right) = \frac{3}{8}\frac{a-c}{b}$$


This will earn the cheating firm profits:

$$\pi_{cheat} = (a - c)q_1 - bq_1^2 - bq_1q_2 = \frac{3}{8}\frac{(a-c)^2}{b} - \frac{9}{64}\frac{(a-c)^2}{b} - \frac{3}{8} \times \frac{1}{4}\frac{(a-c)^2}{b} = \frac{9}{64}\frac{(a-c)^2}{b}$$

This is provided firm 2 continues producing at the cartel quantity. So:

$$\pi_{cheat} > \pi_{cartel} > \pi_{Cournot}$$

The cartel is not stable as both firms will want to cheat. Graphically, firm 1 will move its reaction function where firm 2 continues to produce the cartel quantity of $\frac{a-c}{4b}$. This is at $\frac{3}{8}\frac{a-c}{b}$. This brings firm 1 onto a better isoprofit curve.



Maintaining a Cartel

The firms might be able to sustain the cartel if the firms decided that if one firm cheated on the cartel, they would both play at the Cournot equilibrium forever after. This would work if the interest rate/discount rate, r , is not too large.

- The present value from staying in the cartel forever (cartel profits today plus cartel profits forever after):

$$\pi_{cartel} + \frac{\pi_{cartel}}{r} = \frac{(1+r)}{r} \pi_{cartel}$$
- The present value from cheating today and getting Cournot profits forever after: $\pi_{cheat} + \frac{\pi_{Cournot}}{r}$
- Provided $\frac{(1+r)}{r} \pi_{cartel} > \pi_{cheat} + \frac{\pi_{Cournot}}{r}$, the cartel will be stable.

With our example of linear demand and costs we have $\pi_{cartel} = \frac{(a-c)^2}{8b}$, $\pi_{Cournot} = \frac{(a-c)^2}{9b}$ and $\pi_{cheat} = \frac{9(a-c)^2}{64b}$. The cartel is stable if:

$$\left(\frac{1+r}{r}\right) \frac{(a-c)^2}{8b} > \frac{9(a-c)^2}{64b} + \frac{(a-c)^2}{9br} \implies \frac{1+r}{8r} > \frac{9}{64} + \frac{1}{9r}$$

The cartel is stable in this case if $r < 0.88$. At a very high interest rate, the firm will want to cheat on the cartel.