Intermediate Microeconomics

Greg Leo

These notes are based on my Vanderbilt Economics Course 3012. They are preliminary. If you find any typos or errors in this text, please e-mail me at g.leo@vanderbilt.edu.

Contents

1 Bundles and the Feasible Set

1.1 Bundles

A bundle is a vector (an ordered pair or "tuple" of numbers) representing amounts of things. In this class, our models will often involve two things. Each number in the vector represents an amount of some underlying good. The bundle (x_1, x_2) is a bundle of two goods where x_1 represents the amount of good 1 in the bundle and x_2 represents the amount of good 2.

$$\boxed{\text{Bundle: } x = (x_1, x_2)}$$

To make this concrete, suppose we are building a model about the choice of ice cream bowls. If these bowls can only have two flavors, vanilla and chocolate, then the possible bowls can be written as ordered pairs where x_1 is the amount of vanilla and x_2 is the amount of chocolate.

Here are some possible bundles in this model: $(0, 1)$ one scoop of chocolate. $(2, 0)$ two scoops of vanilla. $(2, 2)$ two scoops of each flavor. Since these bundles are ordered pairs or *vectors*, we can plot them. The ice cream examples are plotted in Figure 1.1.

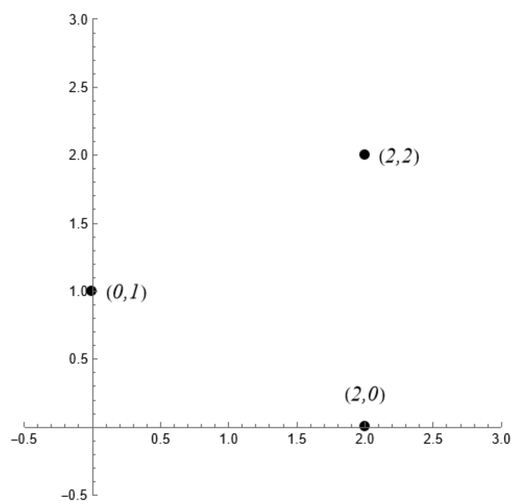

Figure 1.1: The bundles $(0, 1), (2, 2), (2, 0)$.

1.2 Feasible Set

The feasible set is the universe of bundles that might be relevant in a model. The feasible set defines the scope of a model. In our ice cream example, the feasible set X might be the set of all bowls that have a non-negative amount of scoops of vanilla and chocolate.

$$\boxed{\text{Feasible Set: } X \text{ is the "feasible" set of bundles.}}$$

In reality, usually we are limited to choose an integer amount of scoops of ice cream. But allowing only integer choices can cause some complexity in analyzing models. For this reason, we often assume that our feasible sets allow bundles with any real number of each good. In this case, $x = (1.25, 2.387)$ would be a feasible bundle. Perhaps it is best to think of the quantities of each good as something like ounces of ice cream. So this bundle would be 1.25 ounces of vanilla and 2.387 ounces of chocolate ice cream.

2 Budget Set

While the feasible set X is all of the possible bundles, the budget set B is the set of bundles *available* to a particular consumer. The budget set must be a subset of the feasible set.

Budget Set: B

In set notation we write: $B \subseteq X$ which literally says "B is a subset of X". The symbol $\subseteq$ allows the two sets to be equal. B does not have to be strictly smaller than X, it just can't be bigger than X. That is, anything in the budget set has to be a feasible bundle.

Budget sets can be almost anything. For instance, the ice cream shop might give you a coupon that says "This coupon entitles you to either 7 ounces of vanilla ice cream or 3 ounces of vanilla and 4 ounces of chocolate ice cream". This is a weird coupon, but it is perfectly representable with our notation. In this case, your budget set is: $B = \{(7, 0), (3, 4)\}$. Normally, however, our budget sets will be more well-behaved.

2.1 Budget Sets from Prices and Income

Most of the time, we think of "budget" as meaning you have some amount of money you can spend on stuff. Conveniently, this is the usual way we define B. We have prices p_1 and p_2 and an amount of money to spend m.

Prices of good 1 and good 2: p_1, p_2

We usually call m the "income".

Income: m.

To construct the budget set, first we need to calculate the cost of any bundle: $p_1x_1 + p_2x_2$. From here, the set of bundles that a consumer can have is simply all the bundles for which the cost is less than m. Mathematically: $x_1p_1 + x_2p_2 \leq m$. Thus, we can define if formally this way. The budget set: $B = \{x | x \in X \,\&\, x_1p_1 + x_2p_2 \leq m\}$. This set theory notation says that "B is the set of bundles x that are both in the feasible set X and such that such that the price $x_1p_1 + x_2p_2$ of the bundle is less than income m."

We will often want to look at the budget graphically. To do this, first we draw the *budget line*. This is the set of bundles that are "just affordable". That is, they cost exactly m.

Budget Line: $x_1p_1 + x_2p_2 = m$

Now we can plot this on an x_1, x_2 plane. Let's put x_2 of the vertical axis. In this case, it is useful to rewrite the budget line into a form we are more familiar with: $x_2 = \frac{m}{p_2} - \frac{p_1}{p_2}x_1$.

This is now clearly an equation for a line with intercept $\frac{m}{p_2}$ and slope $-\frac{p_1}{p_2}$. Before we plot it, let's interpret it a little. Notice that if $x_1 = 0$ we get $x_2 = \frac{m}{p_2}$. This says "If I were only to buy x_2, I could afford $\frac{m}{p_2}$ units of x_2. Furthermore, for every unit that we increase x_1 by, x_2 goes down by $-\frac{p_1}{p_2}$. This says "Given that I am spending all my money, if I want to buy one more unit of x_1, I have to give up $-\frac{p_1}{p_2}$ units of x_2. This is a very important thing to know about the slope of the budget line. The slope of the budget line represents the trade-off between x_1 and x_2 at the market prices.

Two useful bundles to know about are the endpoints. I recommend always labeling these on a graph of the budget. The two endpoints are $\left(\frac{m}{p_1}, 0\right)$ and $\left(0, \frac{m}{p_2}\right)$. Respectively, these are how much x_1 the consumer can afford if they buy only x_1 and the same for x_2.

Budget Set: $x_1p_1 + x_2p_2 \leq m$

We are now ready to plot the budget set. It is the budget line and all of the bundles "below" the budget line. A budget set formed this way from prices and income is shown in Figure 2.1.

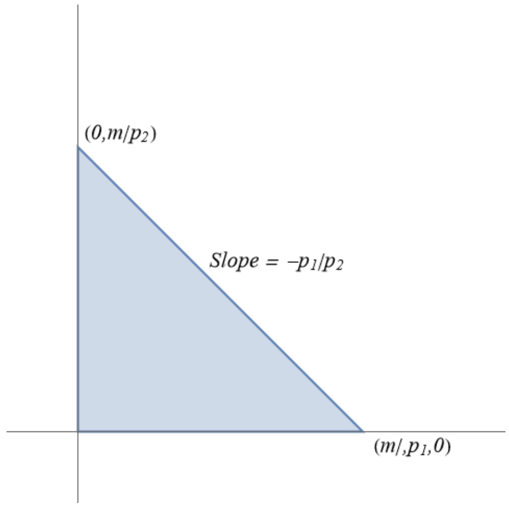

Figure 2.1: Budget set.

2.2 Changing Prices and Income

We are often interested in how the budget set changes when we change one of the budget parameters: p_1, p_2 or m. Since the budget set is just all of the bundles on and below the budget line, we will just focus on what happens to the budget line when we change one of these parameters.

It is easy to determine what happens to the budget line by looking at how a change in one of these parameters affects the three key elements of the budget line: the slope $-\frac{p_1}{p_2}$, the x_1 intercept $\left(\frac{m}{p_1}, 0\right)$, and the x_2 intercept $\left(0, \frac{m}{p_2}\right)$.

Suppose income m increases. Both endpoints therefore change. When m increases, $\frac{m}{p_1}$ (maximum affordable x_1) increases and $\frac{m}{p_2}$ (maximum affordable x_2) increases. The slope does not change. If m decreases, the opposite happens. The case of increasing m is shown in the left panel of Figure 2.2.

Suppose one of the prices changes:

p_1. If p_1 goes up, the slope decreases (becomes more negative). If p_1 goes down, the slope increases. The x_2 intercept stays the same. The case of increasing p_1 is shown in the center panel of Figure 2.2.

p_2. If p_2 goes up, the slope increases. If p_2 goes down, the slope decreases. The x_1 intercept stays the same. The case of increasing p_2 is shown in the right panel of Figure 2.2.

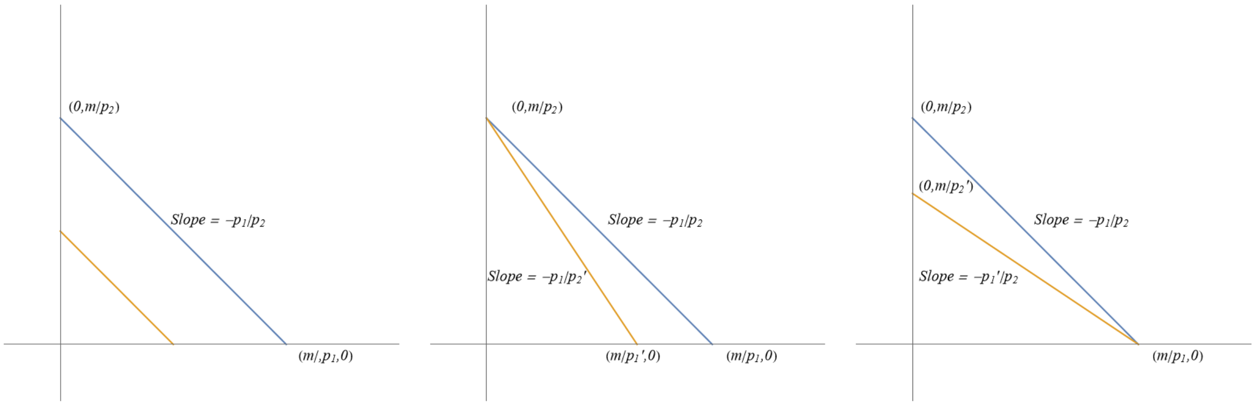

Figure 2.2: How the budget line changes with an increase in: m (left), p_1 (center), p_2 (right). In each case, the blue line is the original budget, and the orange line is the budget after an increase in the relevant parameter.

2.3 Taxes and other scenarios

Budget Line with a quantity tax on x_1: $x_1p_1 + tx_1 + x_2p_2 = m$

Taxes are a familiar method of varying a consumer's budget. There are two common types of taxes: *quantity* and *ad valorem* taxes.

In a *quantity tax*, consumers are charged a fixed amount of money per unit of a good they buy. For instance, suppose a quantity tax of t is added to good 1. The consumer has to pay tx_1 in tax. This gives us a new budget of $p_1x_1 + tx_1 + p_2x_2 = m$. This can also be written as $(p_1 + t)\, x_1 + p_2x_2 = m$ which demonstrates very clearly that a quantity tax simply increases the price of a good by t.

In an *ad valorem tax*, consumers are charged a percentage of the total amount they spend on a good. For instance, suppose a ad valorem tax is 1%. The consumer spends a total of x_1p_1 on good 1. So the tax changed is $0.01\,(x_1p_1)$. More generally if the size of the ad valorem tax is τ (pronounced "tau"), this gives us a new budget of $p_1x_1 + \tau\,(p_1x_1) + p_2x_2 = m$. This can also be written as $[(1 + \tau)\, p_1]\, x_1 + p_2x_2 = m$. This demonstrates, again, that an ad valorem tax is just another way of increasing the price of a good.

We will discuss other scenarios in class such as situations where the price of a good changes depending on how much you buy.

3 The Preference Relation $\succsim$

3.1 Definitions

A preference relation is a set of statements about pairs of bundles. The statement *bundle x is preferred to bundle y* is shortened to $x \succsim y$.

Preference Relation: $x \succsim y$. "Bundle x is preferred to bundle y"

Let's look back at our ice cream example. Again, let (x_1, x_2) be a bundle with x_1 scoops of vanilla and x_2 scoops of chocolate. Suppose someone likes a scoop of vanilla more than a scoop of chocolate. Then the following would be true for them: $(1, 0) \succsim (0, 1)$. They might also like *any* number of scoops of vanilla more than that same number of chocolate. Then the following would also be true of their preferences: $(2,0) \succsim (0,2)$ and $(3,0) \succsim (0,3)$ and $(100,0) \succsim (0,100)$.

The following might be true about a consumer who does not care about flavor at all: $(1,0) \succsim (0,1), (0,1) \succsim (1,0)$. Notice that we have both $(1,0) \succsim (0,1)$ and $(0,1) \succsim (1,0)$. That is, a scoop of vanilla is just as good as a scoop of chocolate and a scoop of chocolate is just as good as a scoop of vanilla. When this is the case, we say the consumer is indifferent.

Indifference Relation: $x \sim y$ "Bundle x is indifferent to bundle y"
$x \sim y$ if and only if $x \succsim y$ and $y \succsim x$.

When a consumer is not indifferent, we say they have strict preference for some bundles. Suppose a consumer would *much* rather have two scoops of vanilla to one scoop then we have $(2,0) \succsim (1,0)$ and $(1,0) \not\succsim (2,0)$. In this case we way that they strictly prefer $(2,0)$ to $(1,0)$ and use the symbol $\succ$. That is $(2,0) \succ (1,0)$.

Strict Preference Relation: $x \succ y$ "Bundle x is strictly preferred to bundle y"
$x \succ y$ if and only if $x \succsim y$ and not $y \succsim x$.

3.2 Assumptions on $\succsim$

Axiom 1. Reflexive: For all bundles, the bundle is at least as good as itself.

In set notation: $\forall x \in X : x \succsim x$

Axiom 2. Complete: For every pair of distinct bundles, either one is at least as good as the other or the consumer is indifferent.

In set notation: $\forall x, y \in X \& x \neq y : x \succsim y \, or \, y \succsim x \, or \, both$

This ensures a consumer can say "I'm indifferent." but not "I don't know" when comparing two bundles.

Axiom 3. Transitivity: If bundle A is preferred to B and bundle B is preferred to C, then bundle A is preferred to C.

In set notation: $x \succsim y, y \succsim z \,\text{implies}\, x \succsim z$

Transitivity (along with the other assumptions) implies we can always put a set of objects (or bundles) into a ranking.

3.3 Indifference Curves and Other Sets

For every bundle, we can use the preference relation to define a few sets. $\succsim (x)$ is the set of bundles at least as good as x. $\succ (x)$ is the set of bundles strictly better than x. $\sim (x)$ is the set of bundles indifferent to x.

Set of points weakly preferred to x: $\succsim (x) = \{y \mid y \in X, y \succsim x\}$
Set of points strictly preferred to x: $\succ (x) = \{y \mid y \in X, y \succ x\}$
Set of points indifferent to x: $\sim (x) = \{y \mid y \in X, y \sim x\}$

Sets of indifferent bundles are very important in studying preferences. We call such a set of bundles an "indifference curve". We use indifference curves to visualize preferences. Note: There are many indifference curves. We only sketch a few to get an idea of the "shape" of preferences.

3.4 Indifference Curves Cannot Cross

If preferences are rational, there's really only one thing we can say about indifference curves.

If $\succsim$ is rational, two distinct indifference curves cannot cross.

The proof of this is remarkably simple and we will discuss it in class, though you are not responsible for the proof.

3.5 Examples of Preferences

3.5.1 Perfect Substitutes

Perfect substitutes preferences have indifference curves that are straight lines.

These preferences are such that the willingness to trade-off between the goods is the same everywhere. The indifference curves are always downward sloping lines with the same slope. The slope measures the amount of x_2 you are willing to give up to get 1 more unit of x_1. A steep slope indicates a stronger relative preference for x_1. A shallow slope indicates a stronger relative preference for x_2.

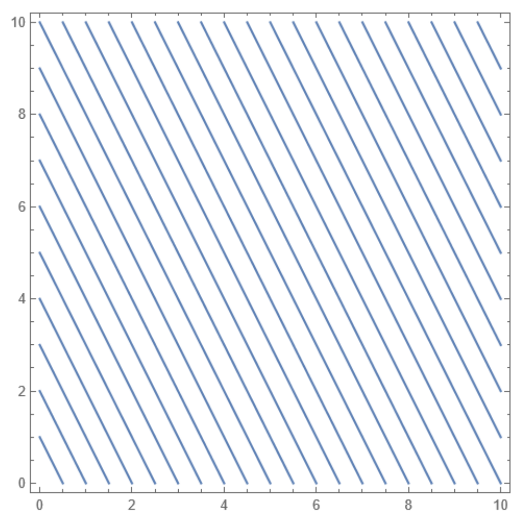

Figure 3.1: Indifference curves of perfect substitutes preferences where the consumer will always give up 2 units of x_2 to get 1 unit of x_1.

3.5.2 Perfect Complements

Perfect complements preferences have indifference curves that are L-shaped.

Perfect complements preferences represent situations where one good cannot substitute for the other. For example: left and right shoes. No matter how many left shoes you have, they cannot replace a right shoe. You must consume them in a 1-to-1 ratio.

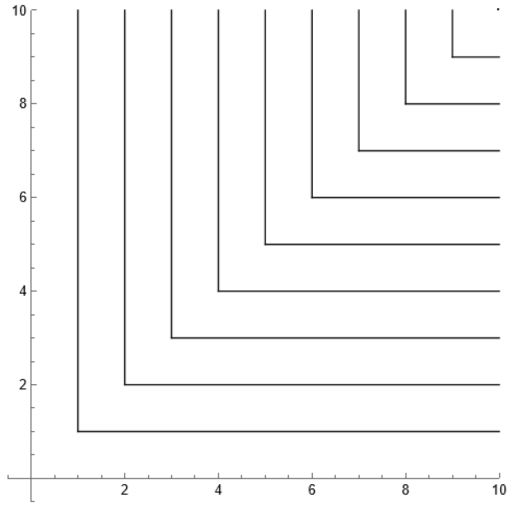

Figure 3.2: Indifference curves of perfect complements preferences where the consumer needs the goods in 1-to-1 combinations.

3.5.3 Bads

A "bad" is a product that a consumer actively dislikes. They prefer none of a bad to any positive amount of a bad. In a bundle with one good and one bad, indifference curves slope upwards. If both products are bads then indifference curves slope downwards, but preference increases towards the origin.

3.6 Further Assumptions: "Well Behaved Preferences"

3.6.1 Monotonicity

Monotonicity is the assumption that everything is a "good". That is, more of either good makes the consumer better off. There are two types of monotonicity:

> Strict Monotonicity: For two bundles (x_1, x_2) and (y_1, y_2), $(x_1, x_2) \succsim (y_1, y_2)$ if $x_1 \geq y_1$ and $x_2 \geq y_2$. $(x_1, x_2) \succ (y_1, y_2)$ if either $x_1 > y_1$ or $x_2 > y_2$.
> Monotonicity: For two bundles (x_1, x_2) and (y_1, y_2), $(x_1, x_2) \succsim (y_1, y_2)$ if $x_1 \geq y_1$ and $x_2 \geq y_2$. $(x_1, x_2) \succ (y_1, y_2)$ if both $x_1 > y_1$ and $x_2 > y_2$

For example, perfect substitutes are monotonic and strictly monotonic. Perfect complements are monotonic but not strictly monotonic. Monotonicity of either kind implies that indifference curves are downward sloping and that preference increases to the north east (away from the origin).

3.6.2 Convexity

Convexity is the assumption that mixtures are better than extremes. Again, there are two forms of this assumption. *Strict* and *Weak.*

> Strictly Convex: For two indifferent bundles $(x_1, x_2) \sim (y_1, y_2)$, for any $t \in (0, 1)$, the mixture given by $(tx_1 + (1-t)\, y_1, tx_2 + (1-t)\, y_2) \succ (x_1, x_2)$ and $(tx_1 + (1-t)\, y_1, tx_2 + (1-t)\, y_2) \succ (y_1, y_2)$.
> Weakly Convex: For two indifferent bundles $(x_1, x_2) \sim (y_1, y_2)$, for any $t \in [0, 1]$, the mixture given by $(tx_1 + (1-t)\, y_1, tx_2 + (1-t)\, y_2) \succsim (x_1, x_2)$ and $(tx_1 + (1-t)\, y_1, tx_2 + (1-t)\, y_2) \succsim (y_1, y_2)$.

As an example, perfect substitutes and perfect complements are both weakly convex because their indifference curves include "flat" portions. On the other hand, the preferences below are strictly convex.

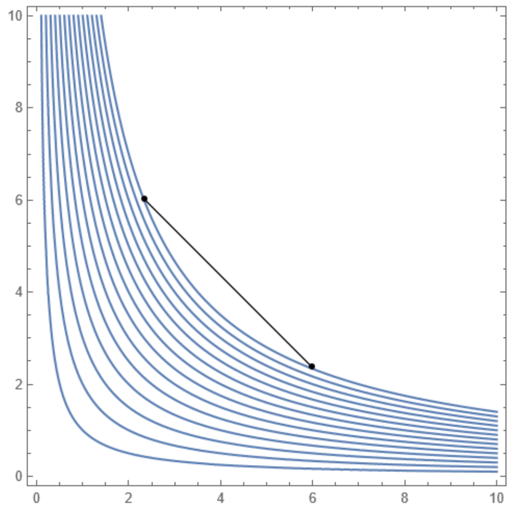

Figure 3.3: Indifference curves of a strictly convex preference relation. A line segment has been placed between two points on one of the indifference curves. Notice that it lies completely above that indifference curve.

It is useful to think about the shape of convex preferences in terms of their indifference curves. Assuming preferences are monotonic, if preferences are strictly convex, then each indifference curve lies strictly below a line between any two points on that indifference curve. If preferences are weakly convex, then the indifference curve always lies on or below a line between any two points on that indifference curve.

3.7 Marginal Rate of Substitution and Slope of the Indifference Curve

The marginal rate of substitution, or MRS, is defined as the slope of the indifference curve at a point. The MRS measures willingness to trade off between good 1 and good 2. Approximately, it's how much x_2 the consumer would give up to get one more unit of x_1.

4 Utility

4.1 Definition

A utility function is a way of assigning "scores" to bundles, such that better bundles according to $\succsim$ get a higher score.

A utility function $U(x)$ represents preferences $\succsim$ when, for every pair of bundles x and y, $U(x) \geq U(y)$ if and only if $x \succsim$

Example 1. Suppose $A \succ B \succ C \sim D$. Some utility functions that represent these preferences are $U(A) = 10$, $U(B) = 5$, $U(C) = U(D) = 0$ and also $U(A) = 12$, $U(B) = 1$, $U(C) = U(D) = -100$.

Definition. A utility function $U(x)$ represents preferences $\succsim$ when for every pair of bundles x and y, $U(x) \geq U(y)$ if and only if $x \succsim y$. That is, if x is better than y according to $\succsim$ it gets a higher utility according to $U()$.

Because the magnitude of exact utility scores are meaningless, and only the relationships between scores matter, we say that these utility functions are "ordinal" rather than "cardinal". There is no sense in which two times higher utility means that the preference is two times stronger. We can only infer the ranking of things, but not how strong the preferences are from $\succsim$ and a utility function that represents $\succsim$.

4.2 Transformations

Any strictly increasing function of a utility function represents the same preferences as the original utility function.

Because utility is ordinal, we are free to transform one utility function into another, as long as it maintains the same preferences. Any strictly increasing function of a utility function represents the same preferences as the original utility function. For example, suppose: $U(x_1, x_2) = x_1 + x_2$. This represents the preferences of someone who only cares about the total amount of stuff, but not the composition. Here are some other utility functions that represent the same preferences: $U'(x_1, x_2) = x_1 + x_2 + 100 = U(x_1, x_2) + 100$. Another one is $U'(x_1, x_2) = (x_1 + x_2)^2 = (U(x_1, x_2))^2$.

4.3 MRS from Utility Function

The MRS is the (negative of) the ratio of marginal utilities: $MRS = -\frac{mu_1}{mu_2} = -\frac{\frac{\partial u(x_1,x_2)}{\partial x_1}}{\frac{\partial u(x_1,x_2)}{\partial x_2}}$

Recall that the marginal rate of substitution (MRS) is the slope of the indifference curve.

Definition. The marginal utility of good i (mu_i) is $\frac{\partial u(x_1,x_2)}{\partial x_i}$.

We can get the MRS at any point by taking the ratio of marginal utilities. $MRS = -\frac{mu_1}{mu_2} = -\frac{\frac{\partial u(x_1,x_2)}{\partial x_1}}{\frac{\partial u(x_1,x_2)}{\partial x_2}}$

Note that, because two preferences that are the same will have the same indifference curves, they will also have the same MRS.

Two utility functions with the same MRS everywhere represent the same preferences.

Same MRS, same preferences.

4.4 Examples of Utility Functions

4.4.1 Perfect Substitutes

Perfect Substitutes: $u(x_1, x_2) = ax_1 + bx_2$

Perfect substitutes preferences can be represented with a utility function of the form $u(x_1, x_2) = ax_1 + bx_2$. The MRS is $-\frac{a}{b}$ everywhere. This constant MRS implies a constant willingness to trade off between the two goods.

4.4.2 Quasi-Linear

Quasi-Linear: $u(x_1, x_2) = x_1 + f(x_2)$

With quasi-linear preferences, the consumer eventually gets tired of one of the two goods. One common quasi-linear utility function is $u(x_1, x_2) = x_1 + ln(x_2)$. The marginal rate of substitution for these preferences are $MRS = -\frac{\frac{\partial(x_1+ln(x_2))}{\partial x_1}}{\frac{\partial(x_1+ln(x_2))}{\partial x_2}} = -x_2$. To interpret this, notice the that amount of x_2 the consumer is willing to give up increases as x_2 increases but does not depend on the amount of x_1 they have.

Another example of a quasi-linear utility function is $u(x_1, x_2) = 10x_1 + \sqrt{x_2}$. Practice taking the MRS of this function. Notice that it only depends on the amount of x_2.

4.4.3 Cobb-Douglas

Cobb-Douglas: $u(x_1, x_2) = x_1^\alpha x_2^\beta$

With Cobb-Douglas preferences, the consumer gets tired of both goods. A general form of these preferences is represented by the utility function $u(x_1, x_2) = x_1^\alpha x_2^\beta$. Let's look at the marginal utilities. $mu_1 = \frac{\partial\left(x_1^\alpha x_2^\beta\right)}{\partial x_1} = \alpha x_1^{\alpha-1} x_2^\beta$, $mu_2 = \frac{\partial\left(x_1^\alpha x_2^\beta\right)}{\partial x_2} = \beta x_1^\alpha x_2^{\beta-1}$. Thus, the MRS is $MRS = -\frac{MU_1}{MU_2} = -\frac{\alpha x_1^{\alpha-1} x_2^\beta}{\beta x_1^\alpha x_2^{\beta-1}} = -\frac{\alpha}{\beta}\frac{x_2}{x_1}$. Note that as the ratio $\frac{x_2}{x_1}$ increases (the consumer has proportionately more x_2) the MRS increases and they are willing to give up more x_2 to get x_1.

4.4.4 Perfect Complements

Perfect Complements: $u(x_1, x_2) = min\{ax_1, bx_2\}$

With perfect complements preferences, there is no substitution possible. The MRS is not defined for this function *because there are no trade-offs the consumer will make.*

4.5 Properties of Preferences and Utility

5 Choice

The goal of a consumer is to choose their most-preferred bundle from the budget set. Formally, they look for a bundle (x_1, x_2) such that for all other bundles $(y_1, y_2) \in B$, $(x_1, x_2) \succsim (y_1, y_2)$. In general, this is not that easy to work out. But, when we have a utility function that represents U along with a competitive budget set, the problem is not so hard. We can write it this way:

The consumer problem is to $max_{x_1,x_2} U(x_1, x_2)$ subject to $p_1x_1 + p_2x_2 \leq m$. This says they find the bundle (x_1, x_2) that gives the highest value of $U(x_1, x_2)$ such that it costs no more than m.

For any consumer with monotonic preferences, more is better, so we know that the consumer will not spend less than their entire income. Note that if we were writing a model where savings was important, they might. We will see that later. But for now, our models are simple. The consumer has only one time period to buy goods in and they have a fixed amount of money to do it. In this case, we can be sure they will spend all of their money. Because of this, for a consumer with monotonic preferences, we can write their maximization problem this way:

Consumer Maximization Problem: $max_{x_1,x_2} U(x_1, x_2)$ subject to $p_1x_1 + p_2x_2 = m$

We can use a little calculus to solve this problem.

5.1 Three Possibilities

Assuming $\succsim$ is monotonic, there are only three possibilities for where an optimal bundle can occur on a budget line. These come from one key observation.

At the optimal bundle, the indifference curve cannot pass through the budget line.

The argument for this is simple. Look at the drawing below. Suppose we think the bundle X is optimal. Here, the indifference curve passes through the budget line and into the interior of the budget set. That means there must a bundle on the indifference curve in the part of the budget set that does not cost all of income m. A point like X'. But, since X' costs less than M, the consumer could take the leftover money and buy more of everything. This would give them a point like X'' which, if preferences are monotonic, must be strictly better than X' which is indifferent to X. Thus X'' is affordable and strictly better than X.

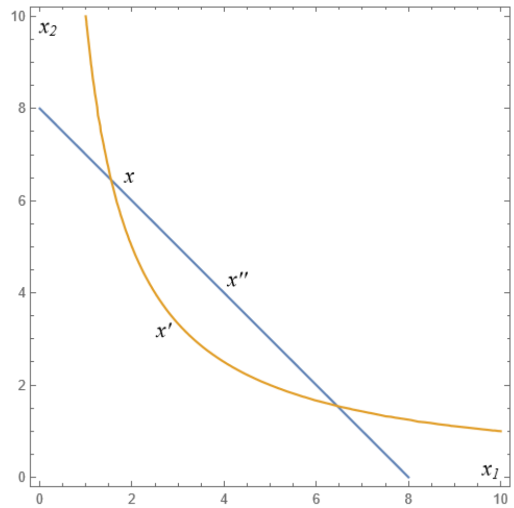

Figure 5.1: An optimal bundle cannot be on an indifference curve that passes "into" the budget set. The indifference curve containing x and x' is shown in orange. The budget line is shown in blue.

There are only three possibilities for an optimal budget on an indifference curve that does not cross into the budget set. These are enumerated below and demonstrated graphically.

1. (Tangent) It is at bundle where the indifference curve at that bundle had the same slope as the budget line.
2. (Touching but not smooth) The bundle is a "non-smooth" point on the indifference curve, but the that point just touches the budget line.
3. (Boundary) We are at one of the boundaries ($x_1 = 0$ or $x_2 = 0$) in this case the slope of the indifference curve and the slope of the budget need not be equal.

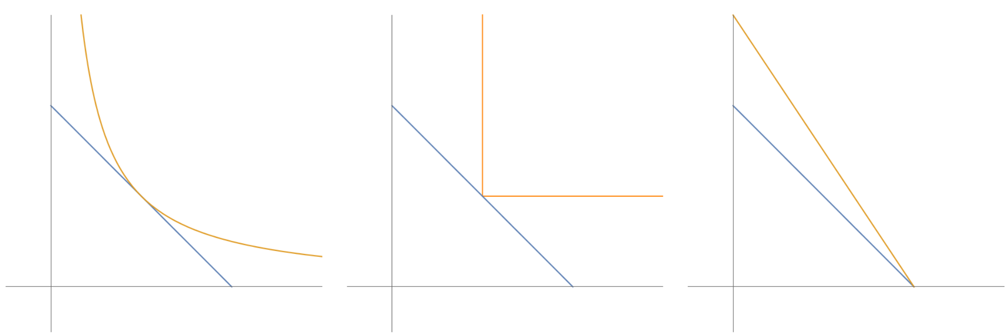

Figure 5.2: Graphical Examples of the three possibilities. Budget is shown in blue. The indifference curve through the optimal point is shown in orange.

Importantly, if we know the indifference curves are smooth everywhere (we can identify this case because we will be able to take the derivative of the utility function), then for a bundle containing some of each good (i.e. not on the boundary of the budget set) to be optimal it must be on a tangent point. The slope of the indifference curve at that optimal bundle is the same as the slope of the budget line. This condition is formalized by the familiar equation: $MRS = -\frac{p_1}{p_2}$.

If the utility function is smooth,

$MRS = -\frac{p_1}{p_2}$ is necessary for an optimal bundle that contains some of both goods.

Another way to interpret the tangency condition is as follows. Normally we have $-\frac{MU_1}{MU_2} = -\frac{p_1}{p_2}$ for the tangency condition. But we can rearrange this to $\frac{MU_1}{p_1} = \frac{MU_2}{p_2}$. Notice that if we divide marginal utility by the price of that good, it tells us the marginal utility of a dollar spent on the good. To see this, suppose the marginal utility of a good is 2, but the good costs \$2 per unit. Then spending \$1 gets half a unit which increases utility by about 1. If we take $\frac{MU}{p}$ we also get 1. Thus, we have the following alternative interpretation of the tangency condition.

The tangency condition ensures that the marginal unit per dollar spent on both goods is the same.

5.2 Examples

5.2.1 Cobb-Douglas:

Let's solve $u(x_1, x_2) = x_1 x_2$, $p_1 x_1 + p_2 x_2 = m$. The MRS is $MRS = -\frac{\frac{\partial(x_1 x_2)}{\partial x_1}}{\frac{\partial(x_1 x_2)}{\partial x_2}} = -\frac{p_1}{p_2}$. Thus the tangency condition is $-\frac{x_2}{x_1} = -\frac{p_1}{p_2}$ or $x_1 p_1 = x_2 p_2$. We also know the optimal bundle occurs on the budget line $x_1 p_1 + x_2 p_2 = m$. We have two equations and two unknowns. We can solve the tangency condition and budget condition together to get. $x_1^* = \frac{1}{2}\frac{m}{p_1}$ and $x_2^* = \frac{1}{2}\frac{m}{p_2}$. Thus, the optimal bundle, or the demand, is $\left(\frac{\frac{1}{2}m}{p_1}, \frac{\frac{1}{2}m}{p_2}\right)$.

5.2.2 Perfect Substitutes

Let's solve $u(x_1, x_2) = 2x_1 + x_2$ with budget $1x_1 + 1x_2 = 10$. That is $p_1 = 1$,$p_2 = 1$,$m = 10$. The MRS is $MRS = -\frac{\frac{\partial(x_1 x_2)}{\partial x_1}}{\frac{\partial(x_1 x_2)}{\partial x_2}} = -\frac{2}{1} = -2$. Thus the tangency condition is $-2 = -1$. Since this can never happen, the only possible solution is one on the boundary. That is, it does not contain some of both goods.

If the consumer only buys x_1, they can get the bundle $\left(\frac{m}{p_1}, 0\right) = (m, 0)$ which has utility 20. If they only buy x_2, they can get $\left(0, \frac{m}{p_2}\right) = (0, m)$ which has utility 10.

Thus, the optimal bundle is $(m, 0)$.

Note: If we had made the budget $p_1 = 2, p_2 = 1, m = 10$ or $2x_1 + 1x_2 = 10$ then we would have gotten $-\frac{2}{1} = -\frac{2}{1}$ for the tangency condition. In this case, the indifference curves have the same slope as the budget line. As long as they spend all of their money, any bundle is optimal. This is because they tradeoff they are willing to make is the same as the trade-off that the market asks them to make to stay affordable.

All of the bundles such that: $p_1 x_1 + p_2 x_2 = m$ are optimal.

5.2.3 Perfect Complements (Left and Right Shoes)

Suppose a consumer buys x_1 left shoes and x_2 right shoes. Their utility function is $u(x_1, x_2) = min\{x_1, x_2\}$ and suppose $p_1 = 2$,$p_2 = 1$,$m = 15$. The budget line is $2x_1 + x_2 = 15$. We still know the budget condition must be true at the optimum. What is the other condition? Notice that here we cannot take derivatives. However, the only possible place an indifference curve could just touch the budget line without crossing into it is at the kink points of the indifference curves. Thus, the equation for the kink points, or what I call the "no waste condition" serves the place of our tangency condition in the previous problems. The no waste condition in this problem specifies: $x_1 = x_2$. Solving the no waste condition along with the budget condition gets us: $x_1 = 5$ and $x_2 = 5$.

5.2.4 Perfect Complements (2 Apples, 1 Crust)

Suppose a consumer makes pies using 2 apples x_1 and 1 crust x_2 for every pie. Their utility function is $u(x_1, x_2) = min\left\{\frac{1}{2}x_1, x_2\right\}$ and suppose $p_1 = 2$,$p_2 = 1$,$m = 15$. The budget line is $2x_1 + x_2 = 15$.

Again, the budget condition must be true at the optimum: $2x_1+x_2 = 15$. And the no waste condition... $\frac{1}{2}x_1 = x_2$. Solving these together, we get $x_1 = 6$, $x_2 = 3$.

5.2.5 Max Preferences

Suppose a consumer's utility is the maximum of the amounts of x_1 and x_2 they have. $u(x_1, x_2) = max\{x_1, x_2\}$. Suppose the budget line is $2x_1 + x_2 = 15$. What is the optimal bundle?

6 Demand

In this section, we explore how demand for goods changes when we changes the "parameters" of the consumer's problem. Those parameters are the prices and income. We will look at how demand changes when one parameter changes but the other are held fixed. This exercise is known as **comparative statics.**

6.1 Marshallian Demand

In the examples in the previous chapter, we solved for demand given specific prices and income. However, we can also solve for demand while leaving these parameters unspecified. When we do that, we find the "Marshallian demand".

Marshallian demand is the optimal amount of a good as a functions of the prices and income. $x_1^*(p_1, p_2, m)$, $x_2^*(p_1, p_2$

The process of solving for Marshallian demand is the same as solving for the optimal bundle at some particular prices and income. We just leave the prices and income as variables rather than plug in values.

6.1.1 Example of Marshallian Demand for Cobb Douglass Preferences.

Let's find the Marshallian demand for this utility function: $u(x_1, x_2) = x_1x_2$.

First, we write down the tangency condition:

$$-\frac{x_2}{x_1} = -\frac{p_1}{p_2}$$

$$p_2x_2 = p_1x_1$$

Notice that this says "spend the same on both goods" and thus implies the consumer will **spend half of my money on good one and half on good two**. This will always be the case for any Cobb Douglass utility function with the same exponent on both goods like $u(x_1, x_2) = x_1x_2$ or $u(x_1, x_2) = x_1^2x_2^2$ or $u(x_1, x_2) = x_1^{\frac{1}{2}}x_2^{\frac{1}{2}}$.

And now the budget equation:

$$p_1x_1 + p_2x_2 = m$$

Solve this system of equations. Plug the tangency condition in to the budget equation and solve for x_1:

$$p_1x_1 + p_1x_1 = m$$

$$2p_1x_1 = m$$

$$p_1x_1 = \frac{m}{2}$$

$$x_1 = \frac{\frac{1}{2}m}{p_1}$$

Plug this back into the tangency condition to get demand for good 2:

$$p_1x_1 = p_2x_2$$

$$p_1\left(\frac{\frac{1}{2}m}{p_1}\right) = p_2x_2$$

$$\frac{1}{2}m = p_2x_2$$

$$x_2 = \frac{\frac{1}{2}m}{p_2}$$

We now have the Marshallian demands:

$$x_1^*\left(p_1, p_2, m\right) = \frac{\frac{1}{2}m}{p_1}, x_1^*\left(p_1, p_2, m\right) = \frac{\frac{1}{2}m}{p_2}$$

6.2 Changes in Income

Now that we know how to find a Marshallian demand, we can start looking at how demand changes when we change one parameter of the problem. We will start with the question of "how does demand change with income?" There are two possibilities.

Normal: Demand goes up when income goes up.
Inferior: Demand goes down when income goes up.

Example. Cobb-Douglas utility: $U\left(x_1, x_2\right) = x_1x_2$

As we saw above, for this utility function, demand for $x_1 = \frac{\frac{1}{2}m}{p_1}$. Since m is only in the numerator, demand must be increasing as income increases. Thus, the *good is normal.* Another way to see this is to take the derivative of x_1 with respect to m. If we do that we get $\frac{\frac{1}{2}}{p_1}$. Since this is a positive number regardless of what p_1 is, it tells us that x_1 increases as m increases.

6.2.1 Engel Curve

The Engel Curve is the relationship between income and demand for a good. Because of... tradition I guess, the Engel curve had income (m) on the vertical axis and the demand on the horizontal axis. This is not how I would personally chose to plot this, but it is what you will see in other courses and textbooks, so I will perpetuate this oddity. One way to interpret the Engle curve is telling us the amount of income a consumer would need to have to demand some amount of that good given some prices.

For instance, suppose demand for x_1 is $x_1 = \frac{\frac{1}{2}m}{p_1}$ and we want to look the the Engel curve for when $p_1 = 2$. In that case, we get $x_1 = \frac{1}{4}m$. Since we need to put m on the vertical axis, it is easier to isolate m in this equation and the plot the result. Doing that we get: $m = 4x_1$. If we wonder how much income the consumer would need such that they can buy 10 units of x_1 we just plut in $x_1 = 10$ and get $m = 4\,(10) = 40$. If the consumer has \$40, they would buy 10 units of the good. Here is the Engle curve plotted:

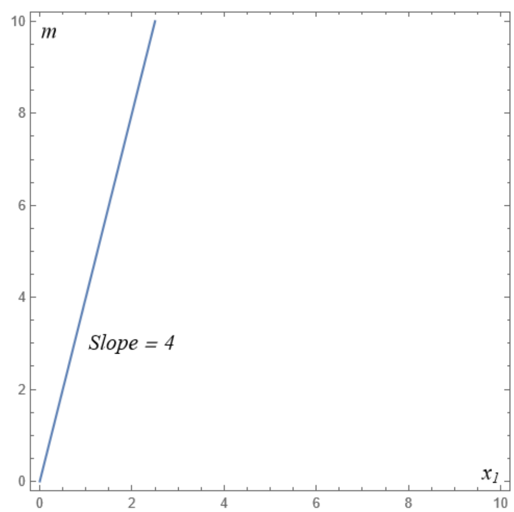

Figure 6.1: Engel curve for $x_1 = \frac{1}{4}m$.

Notice that the Engle curve is upward sloping. That is because demand for x_1 is normal. On the other hand, if the good is inferior, demand decreases as income increases and this is reflected in a "backwards bend" of the Engel curve. This is demonstrated in the plot below.

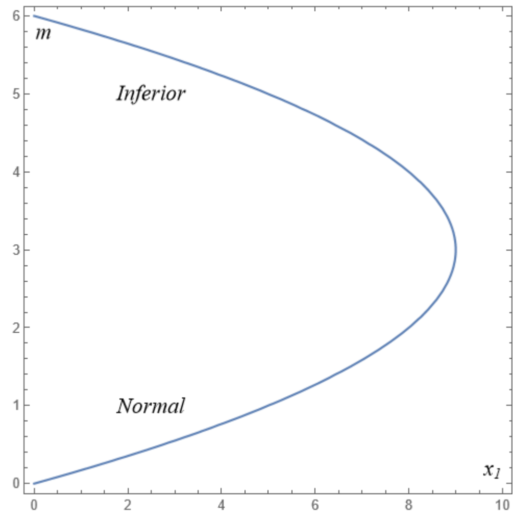

Figure 6.2: Engel curve for a good that is normal for low income and inferior for high income.

6.2.2 A good is never "always inferior".

An inferior good is one for which demand decreases as income increases. But to decrease, the demand for the good must be non-zero. That means, *at some point, it must have increased.* Because of this, a good can never be "always inferior". To be inferior, a good has to start out as normal and *become inferior.* Notice how in the plot of the inferior good above, x_1 starts as normal but becomes inferior as we move "up" in terms of income. This creates the backward bend.

6.2.3 Example: Perfect Complements

Suppose utility is $U(x_1, x_2) = min\{x_1, x_2\}$. Price are: $p_1 = 2$, $p_2 = 1$. Let's solve for demand and then plot the income offer and Engel curve for x_1. At the optimum, we know $x_1 = x_2$ (this is the no waste condition). Since the budget constraint is $2x_1 + 1x_2 = m$ we have two equations and two unknowns. Solve these together, we get demand $x_1 = \frac{m}{3}$, $x_2 = \frac{m}{3}$.

To get the Engel curve. Solve for m in the demand for x_1. We get $m = 3x_1$. This is the equation for the Engel curve for x_1. It is a line with slope of 3. This is plotted below.

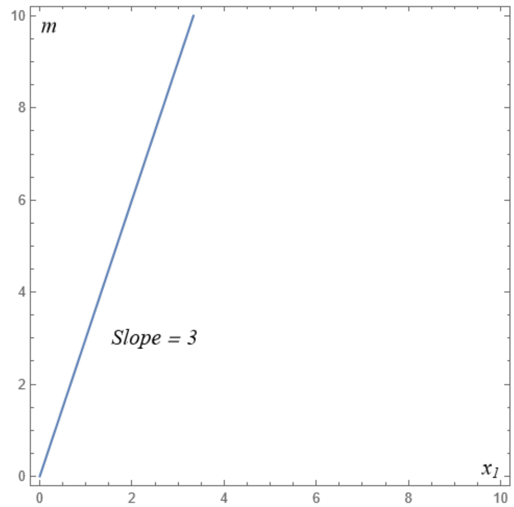

Figure 6.3: Engel Curve of x_1 for $min\{x_1, x_2\}$ with $p_1 = 2$ and $p_2 = 1$

6.3 Changes in "Own" Price

We now look at what happens to demand for a good when the price of that good changes. There are two possibilities.

Ordinary: Demand goes **down** when it's price goes up.
Giffen: Demand goes **up** when it's price goes up?!?

The possibility of Giffen goods might come as a surprise, but they are mathematically possible in our framework. Interestingly, **a Giffen good has to be inferior**. Here is how they arise. When the price of some good goes up, the consumer will naturally trade off to buy other things. But whatever amount of that good they continue to buy will now be more expensive. This makes their income effectively lower – it cannot buy as much. Thus, we can think of a price increase as also lowering effective income. If a good is inferior, this can lead to an increase in the demand for that good. If this effect overwhelms the decrease in demand due to the consumer trading-off to other goods, the net effect *can* be positive. However, in practice, such goods are hard to find and we will not study them extensively in this class.

6.3.1 Plotting the Inverse Demand Curve

Like we plot the relationship between demand and income using the Engle curve, we can also plot the relationship between income and price. If we plot this with demand on the vertical axis, we call is the **demand graph.** However, as with the Engle curve, it is common to put price on the vertical axis. Some people also refer to this plot as the demand graph, but I think that is confusing, so I like to use it's more appropriate name: the **inverse demand graph.** The inverse demand tells us the price that would be responsible for the consumer buying some amount of a good

For example, suppose demand for x_1 is:

$$x_1 = \frac{\frac{1}{2}m}{p_1}$$

Let's plug in an income $m = 10$. We get $x_1 = \frac{5}{p_1}$. This is the demand. To plot the inverse demand, we need to isolate p_1. When we do this we get $p_1 = \frac{5}{x_1}$. Suppose we want to know what price would result in a consumer buying $x_1 = 10$ units of the good. We get $p_1 = \frac{5}{10} = \frac{1}{2}$. So, if price is $\frac{1}{2}$, they will buy 10 units. Here is the inverse demand plotted:

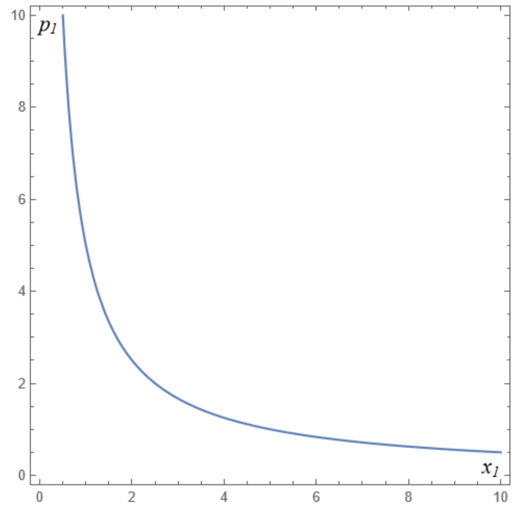

Figure 6.4: Plotting demand for $x_1 = \frac{5}{p_1}$.

6.4 Changes in "Other" Price

So far we have looked at what happens to a good when we change income and it's own price. We might also be interested in how demand changes for a good when the price of another good changes.

Complements: Demand for a good goes down when the price of the other good goes up.
Substitutes: Demand for a good goes up when the price of the other good goes up.
Neither: Demand for a good does not change when the price of the other good goes up.

6.4.1 Example of Perfect Complements

$u = min\{x_1, x_2\}\, has$ demand $x_1 = \frac{m}{p_1+p_2}$ and $x_2 = \frac{m}{p_1+p_2}$. For both goods, as you increase the price of the other good, the demand goes down. They are complements. Hopefully this is not a surprise since we call them perfect complements.

6.4.2 Example of Perfect Substitutes

$u = x_1 + x_2$ has demand $x_1 = \frac{m}{p_1}$ $x_2 = 0$ if $p_1 < p_2$ and $x_1 = 0$ $x_2 = \frac{m}{p_2}$ if $p_1 > p_2$. Let's look at the change in p_1. If $p_1 < p_2$ and p_1 increases, then if it increases enough such that $p_1 > p_2$ the demand for x_2 increases from 0 to $x_2 = \frac{m}{p_2}$. So, as long as the change in price p_1 has any effect on the demand for p_2 (it might not if it does not change which price is higher in this example) then the goods are substitutes.

6.4.3 Example of Cobb-Douglas

Suppose $u = x_1x_2$. Demand is $x_1 = \frac{\frac{1}{2}m}{p_1}$ and $x_2 = \frac{\frac{1}{2}m}{p_2}$. Neither good's demand depends on the price of the other good. They are neither complements nor substitutes.

7 Slutsky Decomposition

This process decomposes the change in demand for a good into two parts:

Substitution Effect: *Price went up, so I will demand less because I buy other things instead. This will always lead to a decrease in demand.*

Income Effect: *Price went up, so what I continue to buy is now more expensive. My effective income is now lower and my demand will change. This effect may be positive or negative.*

> Law of Demand: For a change in price of good i the substitution effect (on good i) will always lead to a decrease or no change in demand x_i.
> Thus, if price of a normal good increases, demand will decrease.

There are three combinations possible:

Ordinary/Normal: *Both effects decrease demand.*

Ordinary/Inferior: *Substitution decreases demand (it always does) and income effect increases demand, but not enough to overcome the decrease due to substitution.*

Giffen/Inferior: *Substitution decreases demand (it always does) and income effect increases demand so much that it overcomes the decrease due to substitution and increases demand overall.*

7.1 Slutsky Decomposition

This decomposition is a sort of thought experiment. Suppose price of a good increases, we go from the budget $p_1x_1 + p_2x_2 = m$ to $p_1^{'}x_1 + p_2x_2 = m$. When the budget changes, demand for x_1 will change as well. The total effect of this change on demand is: $x_1^*\left(p_1, p_2, m\right) - x_1^*\left(p_1^{'}, p_2, m\right)$.

How can we decompose this change into substitution and income effects? To study the substitution effect only, we need to know what the consumer would choose if the price had changed, but their demand could not change due to income. Thus, we think about how much income they would need at the new prices to afford the old bundle. If we were to give the consumer this extra income and ask what they buy at the new prices, then the change in their demand could not be due to the income effect! It is due only to the substitution effect.

To find this, we first calculate the compensating income: cost of the original bundle under the new prices. If we are analyzing a change in p_1 this would be $\tilde{m} = p_1^{'}x_1^*\left(p_1, p_2, m\right) + p_2x_2^*\left(p_1, p_2, m\right)$. Notice this is the cost of the old optimal bundle $x_1^*\left(p_1, p_2, m\right), x_2^*\left(p_1, p_2, m\right)$ but at the new price $p_1^{'}$.

Now we construct a new budget with this compensated income: $p_1^{'}x_1 + p_2x_2 = \tilde{m}$. We ask: what does the consumer choose given this budget? This is denoted $x_1^*\left(p_1^{'}, p_2, \tilde{m}\right)$. The substitution effect is the difference between the original optimal amount and the amount on this new budget line (new prices, compensating income). : $x_1^*\left(p_1, p_2, m\right) - x_1^*\left(p_1^{'}, p_2, \tilde{m}\right)$.

The income effect is the remainder: $x_1^*\left(p_1^{'}, p_2, \tilde{m}\right) - x_1^*\left(p_1^{'}, p_2, m\right)$. That is, the difference between what they choose on the thought experiment budget (new prices, extra income) and what they choose under the new prices with their actual income.

7.2 Example Problem

Suppose $u = x_1x_2$. Demand is $x_1^* = \frac{\frac{1}{2}m}{p_1}, x_2^* = \frac{\frac{1}{2}m}{p_2}$. Suppose $p_1 = 1$, $p_2 = 2$, $m = 10$. The optimal bundle (original prices): is $x_1^* = \frac{\frac{1}{2}10}{1} = 5, x_2^* = \frac{\frac{1}{2}10}{2} = 2.5$.

Now suppose the price of good 1 changes to $p_1^{'} = 2$. The new optimal bundle is $x_1^* = \frac{\frac{1}{2}10}{2} = 2.5, x_2^* = \frac{\frac{1}{2}10}{2} = 2.5$. The total effect is $(5 - 2.5) = 2.5$

Let's calculate the income needed to afford the old bundle at the new prices. Old bundle: $(5, 2.5)$. Cost of this under the new prices: $p_1 = 2$, $p_2 = 2$. Thus the *compensating income is* $5\,(2) + 2.5\,(2) = 15$. We need to construct a budget that has the new prices but enough income to afford the old bundle: $p_1 = 2, p_2 = 2, m = 15$. What does the consumer actually demand here? $x_1\,(2, 2, 15) = \frac{\frac{1}{2}15}{2} = 3.75$. With this we can calculate the substitution effect: $5 - 3.75 = 1.25$.

This leaves the income effect (total effect - substitution effect): 2.5 - 1.25 = 1.25

The price change decreases demand by 2.5. Demand is decreased by 1.25 due to substitution and decreased by 1.25 due to the income effect.

8 Buying and Selling

8.1 Income to Endowments

Until this point our consumers had income in terms of money. $m = \$10$ for instance. Now we will think of the consumers as having an endowment of goods to start with. An endowment is a bundle of goods denoted (w_1, w_2). For instance, if x_1 is apples and x_2 is crusts, an apple farmer might have the endowment $w_1 = 10, w_2 = 0$. This would be an endowment of 10 apples and zero crusts. A baker might have an endowment of 5 crusts and zero apples $w_1 = 0, w_2 = 5$.

When we move from income to endowments, the new budget condition requires that the cost of chosen bundles is less than or equal to the value of the endowment: $p_1x_1 + p_2x_2 \leq p_1w_1 + p_2w_2$. The new budget equation is $p_1x_1 + p_2x_2 = p_1w_1 + p_2w_2$.

Notice how the consumer's effective income, given by $p_1w_1 + p_2w_2$ reacts to changes in price. This is an important distinction from when income was given in terms of an amount of money.

As before we can get the slope of the budget equation, which is still $-\frac{p_1}{p_2}$ and the intercepts. The x_1 intercept (the amount of x_1 afford if I only buy x_1) is $w_1 + \frac{p_2w_2}{p_1}$ and the x_2 intercept is $\frac{p_1w_1}{p_2} + w_2$. Look carefully at these intercepts and see if you can give them an economic interpretation.

8.2 Gross Demand vs Net Demand

In this model we distinguish between what a consumer demands, the gross demand: x_i and the difference between their demand and what they started with, the net demand: $x_i - w_i$.

When net demand is positive, we say the consumer is a net demander or a buyer of that good. When net demand is negative, we say that they are a net supplier or seller of that good.

We can also write the budget equation in terms of net demand by rearranging things: $p_1\,(x_1 - w_1) - p_2\,(w_2 - x_2) = 0$. In this form of the budget equation, it makes clear that a consumer must have budget balance in terms of net demand. This also shows that if a consumer is a buyer of one good, they are a seller of the other.

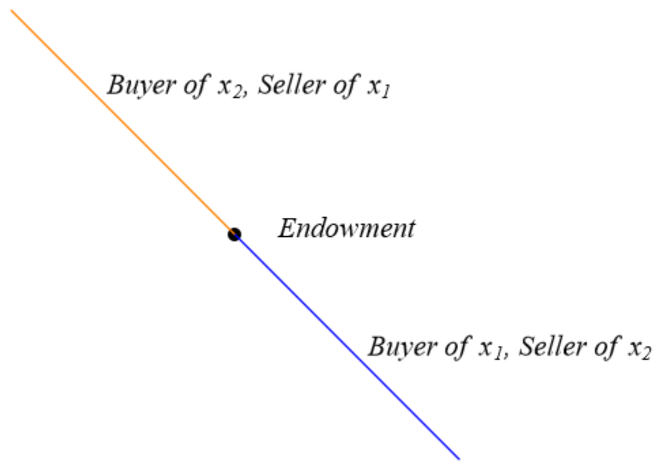

Figure 8.1: The budget line. The area of the budget line above the endowment is where the consumer is a buyer of x_2 and a seller of x_1. The area below is where the consumer is a buyer of x_1 and seller of x_2.

8.3 Drawing the Budget Line and Changes to Price

The budget line always passes through the endowment (w_1, w_2) . If prices change, the slope changes, and the budget pivots through this point.

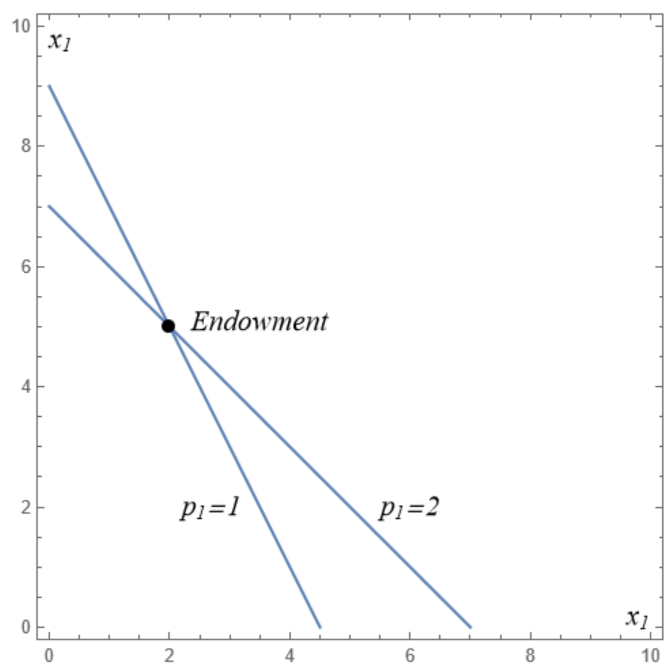

Figure 8.2: An example of how a change in p_1 affects the budget equation with endowment $(w_1, w_2) = (2, 5)$.

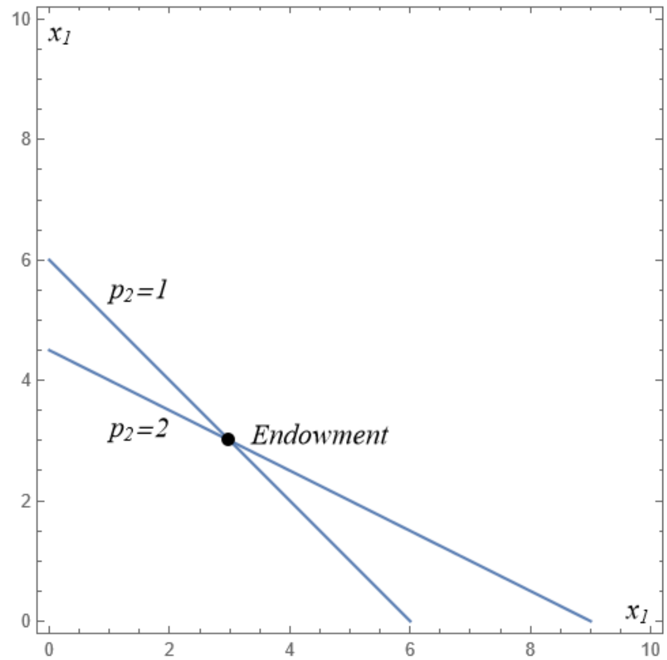

Figure 8.3: An example of how a change in p_2 affects the budget equation with endowment $(w_1, w_2) = (3, 3)$.

8.4 Price Changes and Net Buyers/Sellers

Unlike with a model where the consumer has a monetary income, when a consumer has an endowment, there are some situations where we can say for sure how a price change affects them. The proof of this will be given in class, but you are not responsible for it.

For a consumer who is a net buyer of a good, if the price of that good decreases they will remain a net buyer and will be strictly better off. For a consumer who is a net seller of a good, if the price of that good increases they will remain a net seller and be strictly better off.

8.5 Example Problem

Suppose we have an apple farmer with an endowment of $w_1 = 10$ apples and $w_2 = 0$ crusts. Their utility function is $u = min\left\{\frac{1}{2}x_1, x_2\right\}$. Initially the prices are $p_1 = 1$, $p_2 = 1$.

The consumer's budget equation if $1x_1 + 1x_2 = 1\,(10) + 1\,(0)$ or $x_1 + x_2 = 10$. With this we can solve for demand.

The *No Waste Condition* is $\frac{1}{2}x_1 = x_2$. The budget condition is $x_1 + x_2 = 10$. Solving these two equations gives us: $x_1 = \frac{20}{3}, x_2 = \frac{10}{3}$.

At these prices, the consumer is a seller of apples x_1 and a buyer of crusts x_2. If the price of apples increases or the price of crusts decreases, the consumer will remain a seller of apples x_1 and a buyer of crusts x_2 and will be strictly better off.

9 Intertemporal Choice

9.1 Bundles (Consumption Today, Consumption Tomorrow)

We can use the new model presented in the last chapter to study borrowing and saving behavior. Consider a two period model. The bundles are denoted (c_1, c_2) where c_1 is consumption in period 1, and c_2 is consumption in period 2. The endowments are denoted (m_1, m_2) where m_1 is income in period 1 and m_2 is income in period 2.

9.2 Budget Constraint

"Price" in this model is determined by the interest rate. This is the rate r that the consumer can borrow or save at.

For instance, if the consumer wants to borrow \$1000 in period 1, they have to pay back $1000\,(1+r)$ in period 2. If the consumer saves \$1000 in period 1, they get back $1000\,(1+r)$ in period 2.

Suppose the consumer saves some money in period 1 such that $(m_1 - c_1) > 0$. Consumption in period 2 is m_2 plus how much they saved in period 1 $(m_1 - c_1)$ multiplied my $1+r$. This gives us: $c_2 = m_2 + (1+r)(m_1 - c_1)$.

Suppose the consumer borrows some money in period 1 such that $(c_1 - m_1) > 0$. Consumption in period 2 is m_2 minus how much they have to pay back to cover the loan from period 1 $(1+r)(c_1 - m_1)$. This gives us: $c_2 = m_2 - (1+r)(c_1 - m_1)$ which we can rewrite as $c_2 = m_2 + (1+r)(m_1 - c_1)$.

Notice that these are exactly the same expression and represent the consumer's budget equation whether they choose to borrow or save. $c_2 = m_2 + (1+r)(m_1 - c_1)$. This can be rewritten as: $(1+r)\,c_1 + c_2 = (1+r)\,m_1 + m_2$. This looks like a standard budget equation now. Notice that the right hand side is $(1+r)\,m_1 + m_2$. This is the future value of income (how much c_2 can the can consume if they only consume c_2).

We can also divide both sides of this equation by $1+r$ and get the budget in terms of present value of income: $c_1 + \frac{c_2}{1+r} = m_1 + \frac{m_2}{1+r}$.

9.3 Plotting the Budget Equation

Starting with the budget equation $(1+r)\,c_1 + c_2 = (1+r)\,m_1 + m_2$, we can plot this as usual. The slope is $-\frac{1+r}{1}$ and the c_2 intercept is $(1+r)\,m_1 + m_2$.

As before, we can define situations when $(c_1 - m_1)$ is positive or negative. If $(c_1 - m_1)$ is positive then $c_1 > m_1$. This implies that the consumer is a borrower. If $c_1 - m_1$ is negative, the $m_1 > c_1$ and so the consumer is a saver or a lender.

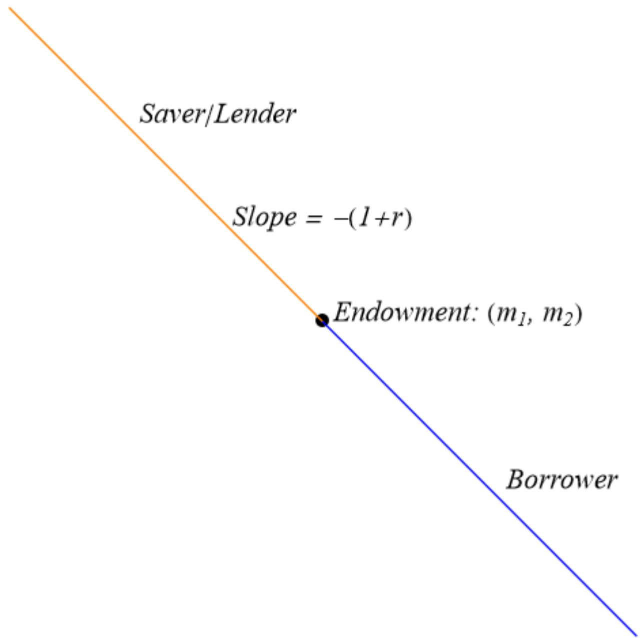

Figure 9.1: The budget line for a model of saving/borrowing.

9.4 Comparative Statics

We can map the results from the last chapter into the language of this chapter. We get:

A borrower, when the interest rate goes down, remains a borrower and must be strictly better off.
A lender (saver), when the interest rate goes up, remains a lender (saver) and must be strictly better off.

9.5 Example Problem

Suppose $m_1 = 200$, $m_2 = 600$, and $r = \frac{1}{2}$. Utility is: $u(c_1, c_2) = c_1 c_2$.

Let's write down the budget equation: $(1 + r)\, c_1 + c_2 = (1 + r)\, m_1 + m_2$. Plugging in the interest rate and incomes: $\frac{3}{2}c_1 + c_2 = \frac{3}{2}200 + 600$

Suppose the consumer chooses to only consume this month (set $c_2 = 0$): $c_1 = m_1 + \frac{m_2}{(1+r)} = 600$. If they choose to only consume next month (set $c_1 = 0$): $c_2 = (1 + r)\, m_1 + m_2 = 900$.

Let's find their actual demand. The tangency condition is $MRS = -(1 + r)$ or $-\frac{c_2}{c_1} = -\frac{3}{2}$. Plugging this into the budget equation and solving gives us: $c_1 = 300$ and $c_2 = 450$.

At this interest rate the consumer is a borrower since $c_1 = 300 > 200 = m_1$. If the interest rate were to decrease to $\frac{1}{4}$ we know that he will remain a borrower. See if you can figure out what the interest rate would have to change to to make the consumer a lender/saver.

10 Market Demand

In this section, we will study the process of adding individual consumer demands together to get a market demand curve.

10.1 Adding Demand Curves

Suppose we have n consumers, each with a demand for good 1 and a demand for good 2. We need a notation for both the consumer and the good. To do this, we use a subscript to refer to the consumer and superscript to refer to the good.

For instance, the demand of consumer 2 for good 1 is $x_2^1(p_1, p_2, m_2)$. Notice we use m_2 to indicate the income for consumer 2. As another example, the demand for consumer 3 for good 2 is $x_3^2(p_1, p_2, m_3)$.

The market demand for a good is the sum of individual consumer demands. We use a capital X to refer to market demand. Formally, the market demand for good 1 is: $X^1(p_1, p_2, m_1, ..., m_n) = \sum_{i=1}^n x_i^1(p_1, p_2, m_i)$. The market demand for good 2 is $X^2(p_1, p_2, m_1, ..., m_n) = \sum_{i=1}^n x_i^2(p_1, p_2, m_i)$.

10.2 Example of Cobb-Douglas Demand and the Representative Consumer Condition

Suppose we have n Cobb-Douglas consumers who all have the utility function: $u_i\left(x_i^1, x_i^2\right) = \left(x_i^1\right)^1\left(x_i^2\right)^1$ (the 1 and 2 superscripts are not exponents, but rather the label for the good). The consumer's demands are: $x_i^1 = \frac{\frac{1}{2}m_i}{p_1}$ and $x_i^2 = \frac{\frac{1}{2}m_i}{p_2}$.

Market demand for good 1 is the sum of the individual demands: $\sum_{i=1}^n\left(x_i^1\right) = \sum_{i=1}^n\left(\frac{\frac{1}{2}m_i}{p_1}\right)$. Suppose $p_1 = 1$ and $m_1 = 10$, $m_2 = 20$, $m_3 = 30$. Market demand is $\left(\frac{\frac{1}{2}10}{1}\right) + \left(\frac{\frac{1}{2}20}{1}\right) + \left(\frac{\frac{1}{2}30}{1}\right) = 30$. Notice, if we let M be the aggregate income $M = \sum_{i=1}^n m_i$. We can write the market demand using the aggregate income:

$$\sum_{i=1}^n\left(\frac{\frac{1}{2}m_i}{p_1}\right) = \frac{1}{2}\frac{1}{p_1}\sum_{i=1}^n m_i = \frac{\frac{1}{2}M}{p_1}$$

In fact, $\frac{\frac{1}{2}M}{p_1}$ is exactly the demand any one of our consumers would have if they happened to have all of the aggregate income M. This is very convenient. Instead of calculating each individual consumer demand and adding them, we can just imagine giving all of the income to one consumer and figuring out what they would choose. If we can, we say that the demand meets the representative consumer condition.

Not all situations meet this condition. First, all the consumers need to have the same utility function. For instance, suppose our consumers had different utility functions. One has perfect substitutes preferences $x_1 + x_2$ one has Cobb-Douglas x_1x_2 and one has perfect complements preferences $min\{x_1, x_2\}$. Who should we pick to be the representative consumer? Whoever we choose, there is no way their choices would represent the entire market.

Furthermore, *even if the consumers all have the same utility function,* the choices of a single consumer with all of the income might not represent the market.

For example, suppose all three consumers have utility function $u(x_1, x_2) = x_1 + ln(x_2)$ and prices are $p_1 = 30$ and $p_2 = 1$. As an exercise, try confirming that the demands for the three consumers are $(0, 10), (0, 20), (0, 30)$ respectively. They each buy no x_1. On the other hand, if we gave one consumer all of the income $M = 60$ then that representative consumer would choose $(1, 30)$. This demand is not the sum of the individual consumer's demands.

So what conditions need to be met for use to use a representative consumer?

1. Utility needs to be the same for all consumers.
2. Preferences of those consumers need to be homothetic.

10.3 Homothetic Preferences

Preferences are homothetic if $x \succsim y$ implies that $tx \succsim ty$ for all possible t. For instance, suppose $(1,2) \succsim (2,1)$, then $(2,4) \succsim (4,2)$ as well if preferences are homothetic.

Homothetic preferences imply that the "shape" of a consumer's preferences depend only on the proportion of goods, but not the scale of the bundle. In the above example, $(1,2)$ and $(2,4)$ are both bundles with a 2-to-1 ratio of x_2 to x_1 and $(2,1)$ and $(4,2)$ are both bundles with a 2-to-1 ratio of x_1 to x_2. The ratios are the same and so the preferences must be the same between the relevant pairs.

A quick way to see if preferences are homothetic (when the utility function is differentiable) is to check that MRS depends on the ratio of goods but not the amount. That is, check whether this condition holds: $MRS\left(x_1,x_2\right) = MRS\left(tx_1,tx_2\right)$. Let's check this for Cobb-Douglas utility function $x_1^{\alpha}x_2^{\beta}$:

$$MRS\left(x_1,x_2\right) = -\frac{\frac{\partial\left(x_1^{\alpha}x_2^{\beta}\right)}{\partial x_1}}{\frac{\partial\left(x_1^{\alpha}x_2^{\beta}\right)}{\partial x_2}} = -\frac{\alpha x_2}{\beta x_1}$$

$$MRS\left(tx_1,tx_2\right) = -\frac{\alpha\left(tx_2\right)}{\beta\left(tx_1\right)} = -\frac{\alpha x_2}{\beta x_1}$$

Let's try $x_1 + ln\left(x_2\right)$. In the previous section we saw we could not use a representative consumer for these preferences. Notice that they are not homothetic.

$$MRS\left(x_1,x_2\right) = -\frac{\frac{\partial\left(x_1+ln\left(x_2\right)\right)}{\partial x_1}}{\frac{\partial\left(x_1+ln\left(x_2\right)\right)}{\partial x_2}} = -x_2$$

$$MRS\left(tx_1,tx_2\right) = -tx_2$$

There are two useful facts to know about homothetic preferences. Some intuition for these will be given in class.

For homothetic preferences:

1. Indifference curves are parallel along a ray through the origin.
2. Engel curves are linear through the origin.

10.4 Elasticity

In economics, we like to use elasticity to measure how demand changes. Elasticity is a *unit-free* measure of change. It allows us to compare the behavior of demand across goods. To see why a unit-free measure of change is useful, consider the following scenario:

Suppose the price of a good changes from 1 to 2. Consumer 1's demand changes from 100 to 50 and consumer 2's changes from 10 to 5. Their behavior in terms of absolute changes in demand $\frac{\Delta x_i}{\Delta p_i}$ is wildly different, a 50 unit change for the first consumer and a 5 unit change for the second consumer. However, their behavior in terms of percentage terms $\frac{\frac{\Delta x_i}{x_i}}{\frac{\Delta p_i}{p_i}}$ is identical– a 50% decrease since $\frac{\frac{\Delta x_i}{x_i}}{\frac{\Delta p_i}{p_i}} = \frac{\frac{100-50}{100}}{\frac{1-2}{1}} = -\frac{1}{2}$.

Elasticity is simply a way of quantifying comparative statics in unit-free percentage terms.

10.5 Elasticities

We want to take a unit-free percentage change for some finite change in price $\frac{\frac{\Delta x_i}{x_i}}{\frac{\Delta p_i}{p_i}}$ and turn it into a way of measuring changes for *very small* changes in price. To do this, we take the infinitesimal analogy to $\frac{\frac{\Delta x_i}{x_i}}{\frac{\Delta p_i}{p_i}} \to \frac{\frac{\partial x_i}{x_i}}{\frac{\partial p_i}{p_i}}$. This is the price elasticity. We can rearrange it a bit as follows:

$$\epsilon_{i,i} = \frac{\frac{\partial x_i}{x_i}}{\frac{\partial p_i}{p_i}} = \frac{\partial x_i}{\partial p_i}\frac{p_i}{x_i}$$

Price elasticity measures the percent change in demand for a one percent increase in the good's own price.

We can also look at the cross-price elasticity:

$$\epsilon_{i,j} = \frac{\frac{\partial x_i}{x_i}}{\frac{\partial p_j}{p_j}} = \frac{\partial x_i}{\partial p_j}\frac{p_j}{x_i}$$

Cross price elasticity measures the percent change in demand for a one percent increase in the price of another good.

Lastly, we can ask how demand changes with income. This gives us the income elasticity:

$$\eta_i = \frac{\partial x_i}{\partial m}\frac{m}{x_i}$$

Income elasticity measures the percent change in demand for a one percent increase in income.

10.6 Cobb-Douglas Example

Suppose utility is $u = x_1 x_2$. Demand for good 1 is: $x_1 = \frac{\frac{1}{2}m}{p_1}$.

The price elasticity is: $\epsilon_{1,1} == -\left(\frac{\frac{1}{2}m}{p_1^2}\right)\frac{p_1}{x_1}$ we now plug in for x_1 to get $-\left(\frac{\frac{1}{2}m}{p_1^2}\right)\frac{p_1}{\frac{\frac{1}{2}m}{p_1}}$. We can simplify this to get $\epsilon_{1,1} = -1$. Elasticity is -1. Thus, a 1% increase in price leads to a 1% decrease in demand. We call this unit-elastic demand.

The cross price elasticity is: $\epsilon_{1,2} = (0)\frac{p_2}{x_1} = 0$. This should make sense. For Cobb-Douglas consumer, the demand for one good does not depend on the price of another.

The income elasticity is: $\eta = \frac{\frac{1}{2}}{p_1}\frac{m}{x_1}$. Plugging in for x_1, we get $\frac{\frac{1}{2}}{p_1}\frac{m}{\frac{\frac{1}{2}m}{p_1}}$. Simplifying this, we get $\eta = 1$. If income increases by 1% demand will increase by 1%. This should be intuitive. The Cobb-Douglas consumer budgets all of their their income. If income increase by 1%, their budget for all goods will also increase by 1%.

10.7 Classifications of Price Elasticity

In the Cobb-Douglas example, price elasticity was -1. A 1% increase in price leads to a 1% decrease in demand.

Unit-Elastic $|\epsilon| = 1$.
Elastic $|\epsilon| > 1$. (For instance -2). Demand responds sharply to changes in price. A 1% increase in price leads to more than 1% decrease in demand.
Inelastic $|\epsilon| < 1$. (For instance $-\frac{1}{2}$). Demand responds weakly to changes in price. A 1% increase in price leads to less than 1% decrease in demand.

Think of some goods that might have elastic and inelastic demand.

11 Equilibrium

11.1 Market Demand and Supply

In this chapter, we focus on where prices come from by looking at the market for *one good at a time*. This is called *partial equilibrium*. A market is made up of the demand side (buyers) and the supply side (sellers). For each side we need to know how their quantity (bought or sold) depends on price. Market demand $Q_d(p)$ is the total amount demanded at price p. Market supply $Q_s(p)$ is the total amount supplied at price p.

When we "plot" a market, we tend to put p on the vertical axis. For this reason, it is also useful to define the inverse market demand $p_d(Q)$ (at what price are Q units are demanded) and inverse market supply $p_s(Q)$ (at what price are Q units are supplied).

Example. Suppose all consumers have utility $x_1 x_2$. Each consumer demands $\frac{\frac{1}{2}m_i}{p_1}$ units of x_1. If we look at the market for x_1 then demand is $Q_d = \frac{\frac{1}{2}M}{p}$ and inverse market demand is $p = \frac{\frac{1}{2}M}{Q_d}$.

11.2 What is an equilibrium?

An equilibrium is defined as a price p^* such that supply is equal to demand. That is $Q_d(p^*) = Q_s(p^*)$. We focus on situations where supply equals demand for the following reason.

Suppose at some price p, supply exceeded demand $Q_s(p) > Q_d(p)$. In this case, price is too high. There are surplus units of the good, and any firm with a surplus unit would be willing to sell at a lower price since otherwise it will be wasted. This creates downward pressure on prices. Suppose now that demand exceeds supply $Q_d(p) > Q_s(p)$. In this case price is too low. There is a shortage and consumers willing to buy at a higher price. There is upward pressure on prices. Thus, the only time there is no pressure for the market price to change is when supply at the price is equal to demand at the price.

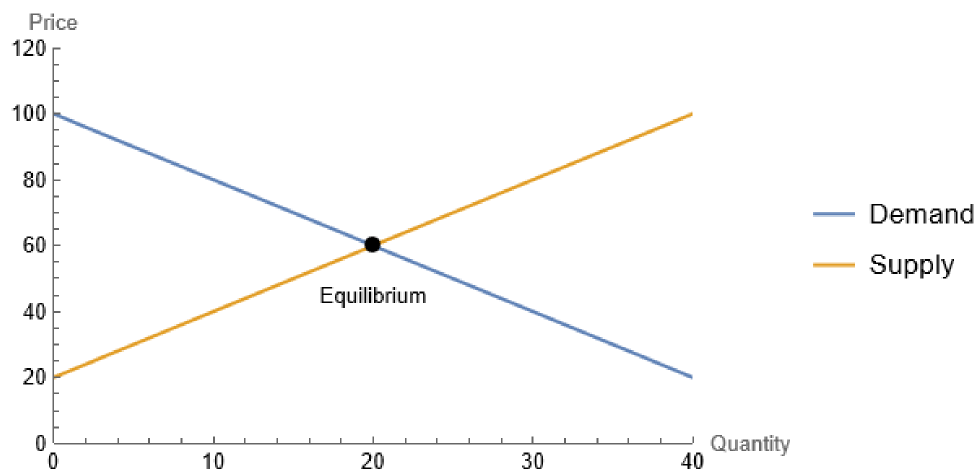

Figure 11.1: An Equilibrium Plot

11.3 Examples

Example. Suppose $Q_d = \frac{2500}{p}, Q_s = 100p$.

To find the equilibrium, we look for a price p^* such that $Q_d = Q_s$. This is done by solving $\frac{2500}{p} = 100p$ which has a solution $p = 5$. To get equilibrium quantity, plug this price into either supply or demand. We should get the same thing: $Q^* = Q_s = Q_d = 500$. Notice, we use Q^* to refer to the equilibrium quantity.

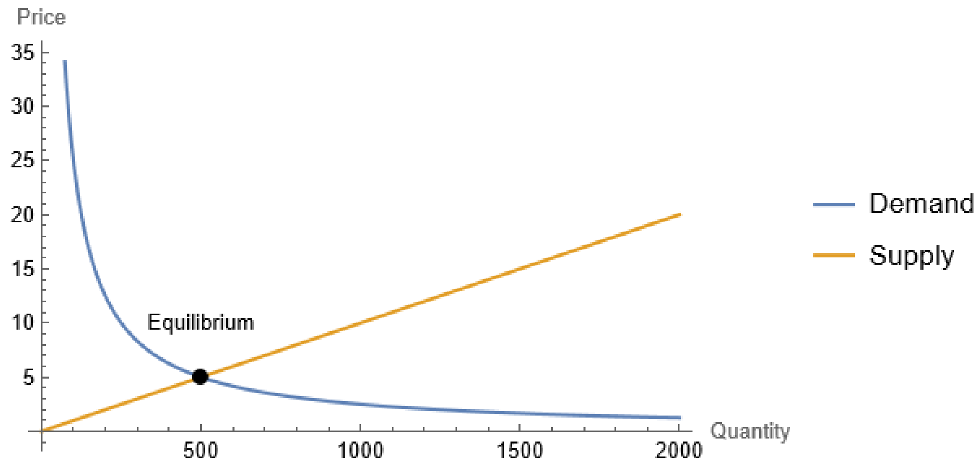

Figure 11.2: Equilibrium Example

Example. Fixed Supply

With fixed supply, the quantity supplied Q_s is constant for any price. The inverse supply curve is a vertical line. This would be the case, for instance, with concert tickets. The size of the venue is fixed regardless of the price of tickets. For example, suppose supply is fixed at $Q_s = 1000$ and demand is $Q_d = \frac{500}{p}$. To find the equilibrium price, solve $1000 = \frac{500}{p}$. This gives us the equilibrium price of $p^* = \frac{1}{2}$. This is the price at which consumers will demand the total supply of 1000. Equilibrium quantity, of course, is $Q^* = 1000$

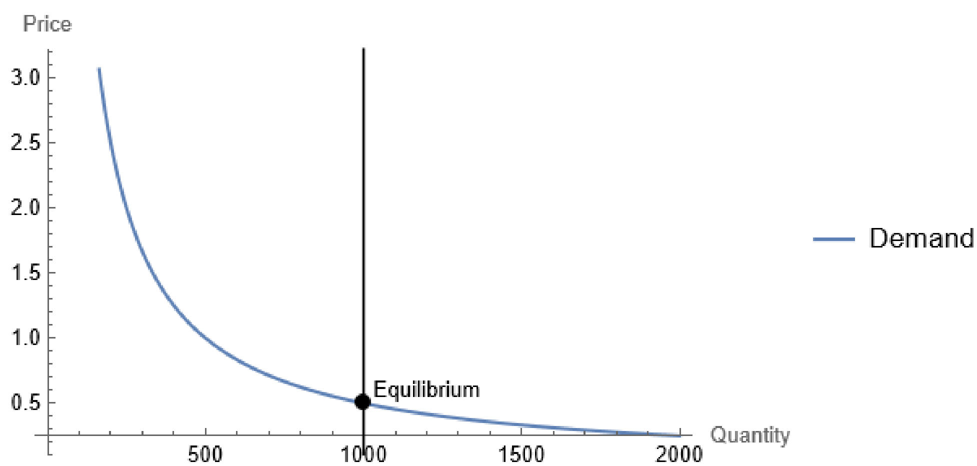

Figure 11.3: Equilibrium with Fixed Supply

11.4 Effect of a Tax

Suppose the government imposes a tax of t per unit of the good. If we think of p as being the price that firms charge for the good (the "sticker price") then firms will receive p for every good sold and consumer will have to pay $p + t$. This leaves us with the following equilibrium condition with a tax: $Q_s(p) = Q_d(p + t)$.

Example. Suppose $Q_s = 100p$ and $Q_d = 300 - 50p$. The government imposes a tax of $t = 3$.

The equilibrium price without a tax is the solution to $100p = 300 - 50p$. This gives us $p^* = 2$ and $Q^* = 200$. With the tax of $t = 3$, the new equilibrium condition is $300 - 50(p + 3) = 100p$. The new equilibrium price is $p^* = 1$ and new equilibrium quantity is $Q^* = 100$. Suppliers get $p = 1$ per unit and consumers pay $1 + 3 = 4$.

The effect of the tax is a lower quantity, consumers pay more then they used, and suppliers receive less than they used to.

Notice that both consumers and producers are worse off under this tax. To quantify how much "worse off" we use the concept of surplus.

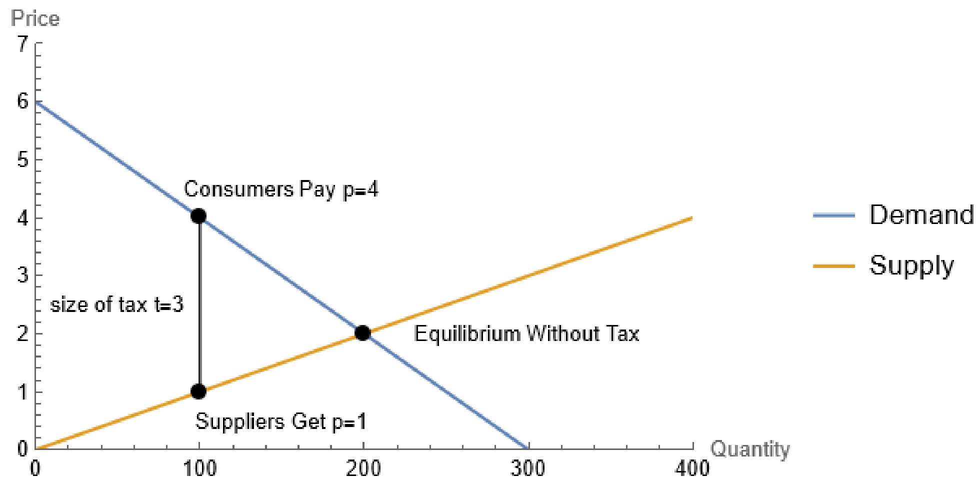

Figure 11.4: Effect of a Tax

11.5 Surplus and Deadweight Loss

Consumer surplus is a measure of welfare that tells us how much "better-off" the consumers are because the market sells them quantity q at price p compared to if the market did not exist at all. The consumer surplus is measured by the area under the inverse demand curve and above price. The producer surplus is the area above inverse supply and below price.

Using the area under inverse demand but above price to measure consumer surplus is motivated by thinking of the height of the inverse demand at some point as the price *some consumer* is willing to pay for a unit of that good. The difference between that height and the price the consumer actually has to pay is one measure of how happy they are to pay less than they were willing to. That is, one consumer's surplus from buying the good at price p. "Summing" over all the consumers gives that area below the inverse demand curve and above price. The same argument motivates the area above the inverse demand and below price as being the producer surplus.

In the tax example above, the consumer surplus (with no tax) is $\frac{1}{2}(4 * 200) = 400$. The producer surplus (with no tax) is $\frac{1}{2}(2 * 200) = 200$. To find these plot the inverse demand and supply, along with the price, and calculate the area of the resulting triangles. The total welfare is the sum of consumer and producer surplus. In this case, that is 600.

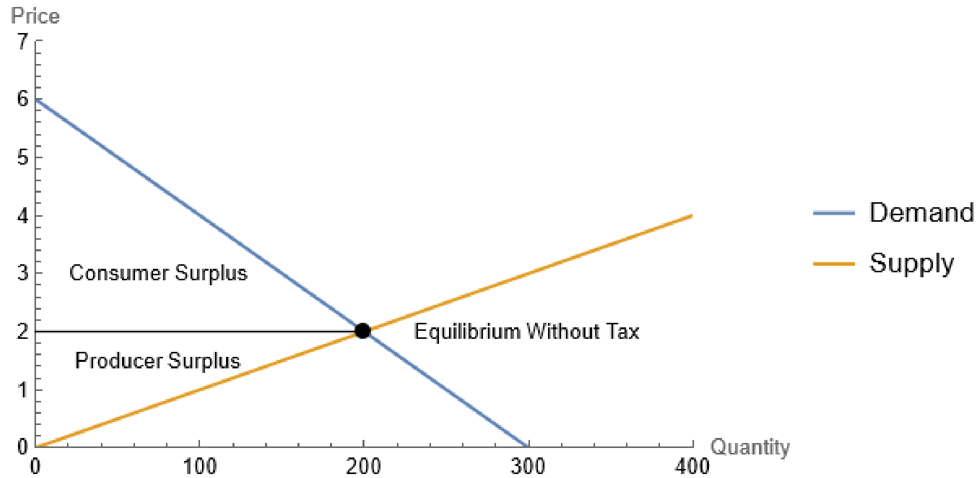

Figure 11.5: Surplus Without Tax

This total surplus of 600 is actually the most we could possibly get in this market. That is because to produce any more surplus we would need to sell more units. However, there is no consumer left who is willing to buy at a price that a firm is willing to sell. That is because to the left of equilibrium, the inverse demand curve is below the inverse supply curve. Since total surplus is maximized here, there is no way to make any consumer or firm better off without making some other consumer or firm worse off. When this is the case, we say that the market has reached Pareto efficiency. In the absence of taxes or other complicating factors, a market equilibrium will always be Pareto efficient.

A tax, however, will lower the total surplus. In the example above, after the tax is imposed, the consumer surplus is $\frac{(6-4)100}{2} = 100$ and the producer surplus is: $\frac{(1)100}{2} = 50$. When there is a tax, we include the government revenue in the calculation of total surplus. This is because tax revenue is not lost. It could be transferred back to consumer or producers in some way, so it contributes to total surplus. The government revenue under the tax in the example above is $3 * 100 = 300$. Total surplus is $100 + 50 + 300 = 450$. Compare this to the original surplus which was 600. The difference is 150. We call this amount the deadweight loss. It measures the amount of total surplus lost due to a tax. This deadweight loss occurs because the tax prevents some firms and consumers from trading even thought there is some price at which they would both be happy to trade at.

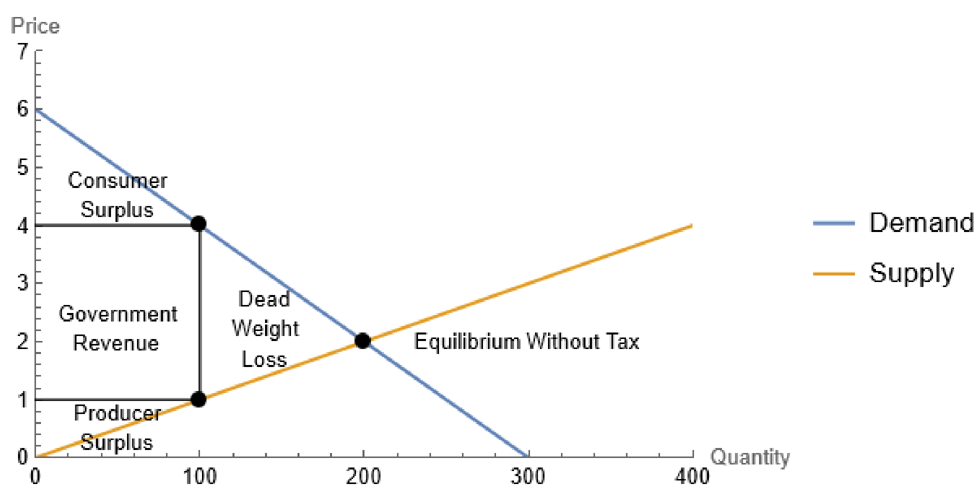

Figure 11.6: Surplus With Tax

11.6 Tax Burden

In the tax example above, after the tax is imposed, consumers pay 4. Before the tax, they only paid 2. Similarly, before the tax producers got 2 per unit sold but now only get 1. Consumers pay 2 more than before and producers get 1 less than before. These differences are called the tax burden or tax incidence. They allow us to determine who ends up "paying" for the tax.

Notice, these amounts sum to the size of the tax (3 in this case). This will always be the case. This also allows us to calculate the tax burden as a proportion of the tax. Just divide the burden on each "side" of the market by the size of the tax. Here, the proportion of the tax paid by consumers is $\frac{2}{3} \approx 66.67\%$ and the proportion paid by producers is $\frac{1}{3} \approx 33.34\%$.

The burden of the tax is determined by the relative elasticities of supply and demand. If demand is relatively elastic and supply is relatively inelastic, then most of burden will be on producers. This is because suppliers cannot not "pass on" much of the tax to consumers. If they did, because demand is relatively elastic, demand would decrease too much. On the other hand, when demand is relatively inelastic compared to supply, most of the burden of the tax will be on the consumers. The suppliers "pass on" most of the tax to consumers because demand is inelastic. We will discuss the graphical intuition for these claims more in class.

12 Technology

We now begin our study of the supply side of the market. Our first task is modeling firms in a mathematical and abstract way. The nature of a firm is that they use inputs to produce outputs and then they sell those outputs to consumers in order to maximize their profits. At lease, that's how we will think of them in this class. In this chapter, we focus just on defining the process by which firms turn inputs into outputs. We do this by defining a technology.

A technology is made up of inputs x_1, x_2 and an output y. For instance, x_1 and x_2 might be apples and crusts and y would be pies. The way that inputs become output is described by the firm's production function.

12.1 Production Functions

A production function maps an amount of each input into an amount of output. Generically, we will write it like this: $f(x_1, x_2)$.

Example. Baker.

A baker can always take 1 crust and 2 apples and produce a pie. If x_1 is crusts and x_2 is apples then we have $f(1,2) = 1, f(2,4) = 2, f(3,6) = 3$ and so on. Generally we can write: $f(x_1, x_2) = min\left\{x_1, \frac{1}{2}x_2\right\}$. This is the baker's production function for pie. Whereas we could with consumers, we cannot take transformations of this function. For instance, the production function $2min\left\{x_1, \frac{1}{2}x_2\right\}$ is one that turns 1 crust and 2 apples into 2 pies. That is a more productive technology, not the same technology. When we are given a production function, that's the one we are stuck with. No transformations allowed.

12.2 Isoquants

Isoquants are combinations of input that give you the same amount of output. They are analogous to indifference curves for consumers. Think of them as recipes for the same output. Let's look at the baker example again. One crust and two apples $(1,2)$ makes one pie, but so does $(1,3), (1,4), (2,2), (3,2)$. These are all input bundles on the isoquant for 1 unit of output. Isoquants for this production function are plotted below for $y = 1, 2, 3, 4, 5$.

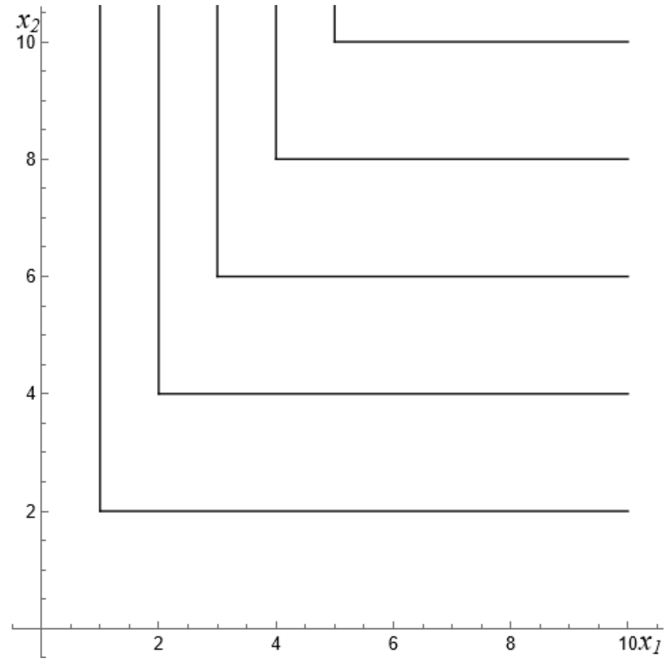

Figure 12.1: Isoquants for $f(x_1, x_2) = y$ for $y = 1, 2, 3, 4, 5$ with fixed-proportion technology $min\left\{x_1, \frac{1}{2}x_2\right\}$.

Working with isoquants is identical to working with indifference curves. The techniques we learned for indifference curves apply when finding isoquants, sketching them, analyzing them, etc., so we will not dwell on those techniques much here.

12.3 Marginal Products

When we talked about utility functions, the partial derivatives (the marginal utilities) were useful in finding the marginal rate of substitution, but were not that useful on their own. This is because "how much" utility increases when we increase a good is not meaningful since "how much" utility is not itself meaningful. The marginal utilities are only meaningful in comparison to each other. On the other hand, specific quantities of production are meaningful and tangible. 5 pies is 5 pies, while 5 points of utility is not tangible.

Because of the fact that the amount of production is a meaningful number, how that number changes when we change one of the inputs is also meaningful information. These are the marginal products $MP_i = \frac{\partial f(x_1, x_2)}{\partial x_i}$. This is the partial derivative of the production function with respect to input i. For example, suppose $f(x_1, x_2) = 2x_1 + x_2$. $MP_1 = 2, MP_2 = 1$. If we increase input 1 by one unit (holding x_2 fixed), we get 2 more units of output. If we increase x_2 by one unit (holding x_1 fixed) we get 1 more unit of output.

As another example, suppose $f(x_1, x_2) = (x_1 + x_2)^{\frac{1}{2}}$. The name of this function is the CES production function (constant elasticity of substitution). Don't worry about what constant elasticity of substitution means just yet, we will discuss it later if there is time. The marginal products are $MP_1 = \frac{\partial\left((x_1+x_2)^{\frac{1}{2}}\right)}{\partial x_1} = \frac{1}{2}\frac{1}{\sqrt{x_1+x_2}}$ and $MP_2 = \frac{1}{2}\frac{1}{\sqrt{x_1+x_2}}$. Notice the extra output for increasing either of the inputs only depends on the sum of the input amounts $x_1 + x_2$ and is decreasing in both.

As a final example, consider $f(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$ (Cobb-Douglas production). $MP_1 = \frac{\partial\left(x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}\right)}{\partial x_1} = \frac{1}{2}x_1^{\frac{1}{2}-1}x_2^{\frac{1}{2}} = \frac{1}{2}x_1^{-\frac{1}{2}}x_2^{\frac{1}{2}} = \frac{1}{2}\frac{x_2^{\frac{1}{2}}}{x_1^{\frac{1}{2}}} = \frac{\sqrt{x_2}}{2\sqrt{x_1}}$ and $MP_2 = \frac{\sqrt{x_1}}{2\sqrt{x_2}}$. The marginal product of 1 is decreasing in x_1 but increasing in x_2 and vise versa.

12.4 Diminishing Marginal Product

Diminishing marginal product is the property that as you increase one of the inputs while holding the other fixed, the additional output you get will decrease. That is, each input becomes less productive as you increase only that input. The definition is that, $\frac{\partial(MP_i)}{\partial x_i} = \frac{\partial f(x_1,x_2)}{\partial x_i \partial x_i} < 0$. This is the second derivative of f with respect to x_i. For marginal product to be diminishing in both inputs, the second partial derivative of f has to be negative for both inputs. Have a look at the Cobb-Douglas example above, $MP_1 = \frac{\sqrt{x_2}}{2\sqrt{x_1}}$. $\frac{\partial\left(\frac{\sqrt{x_2}}{2\sqrt{x_1}}\right)}{\partial x_1} = -\frac{\sqrt{x_2}}{4x_1^{3/2}} < 0$. Thus, there is diminishing marginal product for input one of the production function $f(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$. There will also be diminishing marginal product for input 2. Not all production functions have diminishing marginal product. For instance, the Cobb-Douglas production function $x_1^2 x_2^2$ has *increasing* marginal product for both x_1 and x_2.

12.5 Returns to Scale

While marginal product measures how production changes as we change one of the inputs, the returns to scale measures how production changes when all of the inputs are scaled up. Take our baker example. If we start with the input bundle $(1, 2)$ 1 crust, 2 apples, we get 1 pie. If we double either of the inputs, we still get one pie. For instance $f(2, 2)$ and $f(1, 4)$ both give 1 pie. However, if we double both inputs to $(2, 4)$, we get 2 pie. Doubling the inputs doubles the outputs. We call this linear returns to scale. However, for some production functions, when we double the inputs, we get less than double the outputs (decreasing returns to scale) or when we double inputs, we get more than double the output (increasing returns to scale).

Formally for any $t > 1$:

Linear (constant) returns to scale requires: $f(tx_1, tx_2) = tf(x_1, x_1)$.
Decreasing returns to scale requires: $f(tx_1, tx_2) < tf(x_1, x_1)$.
Increasing returns to scale requires: $f(tx_1, tx_2) > tf(x_1, x_1)$.

For example, consider $f(x_1, x_2) = (x_1 + x_2)^{\frac{1}{2}}$. $f(2, 2) = (2 + 2)^{\frac{1}{2}} = 2$. If we double the inputs we get: $f(4, 4) = (4 + 4)^{\frac{1}{2}} = 2.82843$. Doubling the inputs leads to less than double the output. We have decreasing returns to scale at $(2, 2)$. To prove this function has decreasing returns to scale everywhere note that $f(tx_1, tx_2) = (tx_1 + tx_2)^{\frac{1}{2}} = \sqrt{t}(x_1 + x_2)^{\frac{1}{2}}$. Since $\sqrt{t} < t$ for any $t > 1$, $f(tx_1, tx_2) < tf(tx_1, tx_2)$ which is the definition of decreasing returns to scale. In class we will look more closely at how we check the returns to scale of various production functions.

12.6 Technical Rate of Substitution

Along a particular isoquant, the slope of the isoquant measures how much x_2 you can give up if you add 1 unit of x_1 so that you continue producing the same amount of output. This slope and tradeoff are measured by the technical rate of substitution. It is analogous to the marginal rate of substitution. $TRS = -\frac{\frac{\partial f(x_1,x_2)}{\partial x_1}}{\frac{\partial f(x_1,x_2)}{\partial x_2}} = -\frac{MP_1}{MP_2}$. This is measuring the tradeoffs that a firm is *willing* to make to produce the same amount. Eventually it will play a key role in finding optimal input bundles.

13 Profit Maximization / Cost Minimization

13.1 Profit Function

In this class, we will assume a firm's goal is to maximize profits. The profit function of a firm is made up of revenue and costs.

Revenue is the price the firm gets multiplied by the output of the firm. Usually, the price a firm gets is a function of how much it produces. How this price is determined depends on details of the market. We will look at some possibilities in later chapters. But for now, we just leave this as some unknown price function $p(y)$. Output is simply y. Firm revenue is $p(y)\,y$. Notice, since $y = f(x_1, x_2)$, so we could also write revenue in terms of the amounts of inputs used instead of outputs. $p(f(x_1, x_2))\,f(x_1, x_2)$.

Costs are determined by the amount of inputs used and the price of those inputs. Here, we will assume the price of inputs is fixed at w_1 and w_2. (We use w because price of inputs are often called "wages"). Thus, costs are $w_1x_1 + w_2x_2$. This leads to the following profit function:

Firm profit function: $\pi(x_1, x_2) = p(f(x_1, x_2))\,f(x_1, x_2) - (w_1x_1 + w_2x_2)$

Sometimes a firm cannot change one of its inputs. When this is the case, we refer to the situation as "short-run". In the long-run, all of the firms inputs are variable. For instance, suppose x_2 is fixed at $\bar{x}_2$. Then the firm's short-run profit function is: $\pi(x_1, \bar{x_2}) = p(f(x_1, \bar{x_2}))\,f(x_1, \bar{x_2}) - (w_1x_1 + w_2\bar{x_2})$.

13.2 Profit Maximization Requires Cost Minimization

In the long-run, a firm's goal is to maximize this profit function by choosing x_1 and x_2. Doing this in one step by maximizing this function is possible. However, it is much easier to break the problem down into two parts using the following observation.

Profit maximization implies cost minimization.

To see this, notice that whatever values of x_1 and x_2 maximize profit, there is some amount of output produced y^*, this is the profit maximizing level of output. However, if the firm is using any x_1 and x_2 except the cheapest possible way of producing y^*: (x_1^*, x_2^*) then if could produce the same output and get the same revenue while reducing cost. This would lead to an increase in profit. Thus, if a firm was not minimizing the cost of producing what they thought was the profit maximizing level of output, there is a cheaper way to earn the same revenue, and thus get more profit.

This lets us break down the profit maximization problem into two steps:

1. Calculate the cheapest way to produce any level of output y.
2. Calculate the most profitable y.

Step 1 is what we will focus on in this chapter. It looks like this:

$$Min_{x_1,x_2} w_1x_1 + w_2x_2 \text{ subject to } f(x_1, x_2) = y.$$

Notice when we minimize cost to complete this step, we can ignore revenue, which depends on that pesky function $p(y)$ and allows us to put off talking about how price depends on output.

13.3 Cost Minimization

When we discussed utility maximization, we argued that a bundle that maximizes utility must be on an indifference curve that does not cross through the budget line. A very similar property will hold for cost minimization. First, we need to define the notion of "isocost" lines. These are sets of input bundles that cost the same to use. They are sets of (x_1, x_2) that meet the condition $w_1x_1 + w_2x_2 = c$ for some c. These look a lot like the budget line from the consumer problem. These are straight lines with slope $-\frac{w_1}{w_2}$. What a firm does in trying to find a cost minimizing bundle for producing output y is to look for a bundle of inputs that is on the isoquant for output y but is on the lowest isocost. This process is plotted below.

Notice that the bundle x is on the isoquant for y units of output. So it produces the right amount of output, but notice that the isocost through x' crosses the isoquant and includes the bundle x'' which costs the same as x' but produces strictly more that y. Thus, x could not possibly be cost minimizing, since it costs the same as another bundle that produces more than y. The firm could always find a bundle like x' that produces y and uses less of both inputs than x''. Thus a bundle like x could not possibly be cost minimizing. In this case, x' is the cost minimizing bundle.

A cost minimizing bundle must lie on an isocost that does not cross through the isoquant for y.

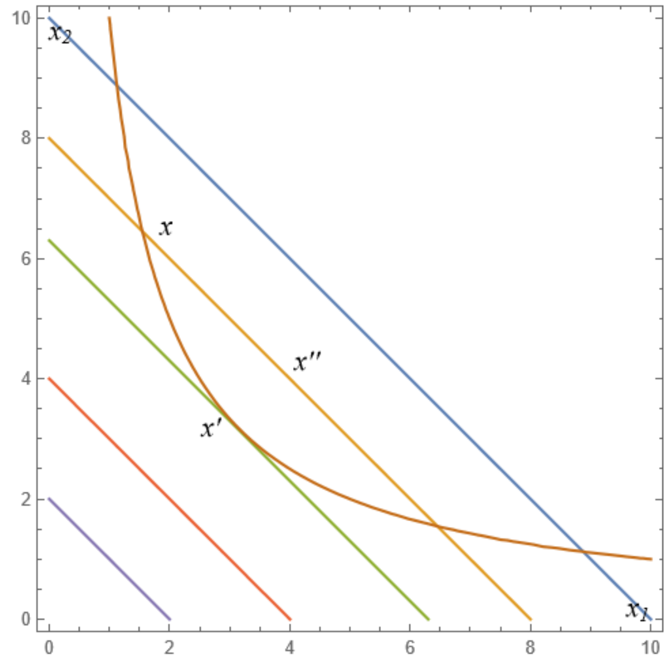

Figure 13.1: Demonstration of cost minimization. The bundle x' is cost minimizing.

The result of this is that, as long as the production function is smooth and we can take its derivatives, the cost minimizing bundle must occur where the slope of the isoquant is equal to the slope of the isocost. $TRS = -\frac{w_1}{w_2}$. This is identical to $-\frac{\frac{\partial f}{\partial x_1}}{\frac{\partial f}{\partial x_2}} = -\frac{w_1}{w_2}$ which can also be written $\frac{MP_1}{w_1} = \frac{MP_2}{w_2}$. This equation implies that the cost of increasing output by 1 unit using x_1 $\left(\frac{MP_1}{w_1}\right)$ is the same as the cost of increasing output by 1 unit using x_2 $\left(\frac{MP_2}{w_2}\right)$. I hope you will find this to be an intuitive condition for cost minimization. If it were not the case, the firm could reduce their use of the more expensive input (per unit of output), and increase their use of the less expensive (per unit of output) input and lower their cost while producing the same amount.

13.4 Minimizing Cost for a Cobb-Douglas Production Function

Since the mathematical conditions for cost minimization are so similar to the conditions for maximizing utility, you will find the examples to be very familiar. For instance, let's minimize the cost of producing y units of output with production function f$(x_1, x_2) = x_1^{\frac{1}{4}} x_2^{\frac{1}{4}}$.

The TRS is $-\frac{\frac{\partial\left(x_1^{\frac{1}{4}} x_2^{\frac{1}{4}}\right)}{\partial x_1}}{\frac{\partial\left(x_1^{\frac{1}{4}} x_2^{\frac{1}{4}}\right)}{\partial x_2}} = -\frac{x_2}{x_1}$. This gives us the equal-slope condition $-\frac{x_2}{x_1} = -\frac{w_1}{w_2}$. Solving this condition for x_1: $x_1 = \frac{x_2 w_2}{w_1}$.

Instead of plugging this into a budget equation like we would for the consumer utility maximization, we need to plug it into the producer's constraint, the production constraint, which is $x_1^{\frac{1}{4}}x_2^{\frac{1}{4}} = y$. Plug in the condition above for x_1 : $\left(\frac{x_2 w_2}{w_1}\right)^{\frac{1}{4}} x_2^{\frac{1}{4}} = y$. Solving this for x_2 gives us the so-called conditional factor demand for x_2: $x_2 = y^2\left(\frac{w_1}{w_2}\right)^{\frac{1}{2}}$. This is the amount of x_2 to use to produce y in the cheapest way possible. Plug this back into the tangency condition above to get the conditional factor demand for $x_1 : x_1 = y^2\left(\frac{w_2}{w_1}\right)^{\frac{1}{2}}$.

To calculate the cheapest cost for producing y, calculate the cost of the conditional factor demands using $w_1x_1 + w_2x_2$. Plugging these demands in gives us: $w_1\left(y^2\left(\frac{w_2}{w_1}\right)^{\frac{1}{2}}\right) + w_2\left(y^2\left(\frac{w_1}{w_2}\right)^{\frac{1}{2}}\right) = 2w_1^{\frac{1}{2}}w_2^{\frac{1}{2}}y^2$. This is the so-called cost function for the producer. $c(y) = 2w_1^{\frac{1}{2}}w_2^{\frac{1}{2}}y^2$. This is the amount it costs to produce y in the cheapest way possible. It is a very important function.

13.5 Profit Maximization Through Cost Minimization

Once we have the cost function for a firm, we can write the profit function as a function of y by replacing $w_1x_1 + w_2x_2$ with $c(y)$ to get $\pi(y) = p(y)y - c(y)$. Since $c(y)$ is the cheapest way of producing y, this will give the most profit a firm could possibly earn if it produces output y. The firm is just left to chose the optimal y. This is very easy to maximize since it is just one-dimensional. It only depends on y.

We still need to know what $p(y)$ is. But for now, let's use a simple assumption that price does not depend on output y, the firm just assumes the price they will get for each unit of output is fixed at p. This is called the price-taking assumption. This assumption is not valid in many cases. The idea that the price a firm can get for *any* amount of output it chooses is unreasonable. But if the firm is a *very small* part of a market (we call this perfect competition) it is probably an assumption we can get away with. We will discuss this more in class.

In any case, if we make the price-taking assumption, we can write profit as $\pi(y) = py - c(y)$. Let's look at the example from above. If we want to maximize profit with the production function f$(x_1, x_2) =$$x_1^{\frac{1}{4}}x_2^{\frac{1}{4}}$ and the price of output is assumed to be fixed at p, profit is $\pi(y) = py - 2w_1^{\frac{1}{2}}w_2^{\frac{1}{2}}y^2$. Notice we have plugged in the cost function we found above.

For an interior maximum (y is some number greater than 0), the slope of this will have to be zero at the optimum y^*. Otherwise, the firm could increase or decrease output and increase profit. The first order condition is: $\frac{\partial(\pi(y))}{\partial y} = 0$ which here is $p - 4\sqrt{w_1}\sqrt{w_2}y = 0$. Notice, we can rewrite this as: $p = 4\sqrt{w_1}\sqrt{w_2}y$. The left side of this is the extra revenue from increasing output by one unit (the marginal revenue (MR). The right side of this is the extra cost from increasing output by one unit (the marginal cost, MC). It will always be true that the firms optimal output solves where $MR = MC$.

Under the price taking assumption (that price p does not depen\d on y) the marginal revenue is just p and we have $p = MC$. Returning to the example, we can solve y to get the optimal y for any set of prices: $y^*(p, w_1, w_2) = \frac{p}{4\sqrt{w_1}\sqrt{w_2}}$.

This the optimal (profit maximizing) level of output for any price. We can also write the "profit function". Take this optimal level of output and plug it back into the "conditional profit function". We found previously that this conditional profit function is: $\pi(y) = py - 2w_1^{\frac{1}{2}}w_2^{\frac{1}{2}}y^2$.

Plugging in the optimal level of production yields the profit function: $\pi(y^*) = p\left(\frac{p}{4\sqrt{w_1}\sqrt{w_2}}\right) - 2w_1^{\frac{1}{2}}w_2^{\frac{1}{2}}\left(\frac{p}{4\sqrt{w_1}\sqrt{w_2}}\right)^2 = \frac{p^2}{8\sqrt{w_1}\sqrt{w_2}}$. Suppose $p = 10$ and $w_1 = w_2 = 1$ the maximum profit the firm can earn is (plug prices into the profit function above): $\pi^* = \frac{100}{8} = \frac{25}{2}$. Find the optimal level of output by plugging prices into the optimal output function $y^*(p, w_1, w_2) = \frac{p}{4\sqrt{w_1}\sqrt{w_2}}$. $y^* = \frac{10}{4} = \frac{5}{2}$.

13.6 More on Supply Under Price-Taking

When we make the price-taking assumption, firms do not consider how their quantity chosen affects the market price. Their profit function is $\pi = py - c(y)$. Maximizing this requires finding the point where the slope of the profit function is zero. This occurs where: $p = \frac{\partial(c(y))}{\partial y}$.

For example, suppose $c(y) = 5y^2$. Marginal cost (MC) is $10y$. Setting $p = mc$ we have $p = 10y$. This is the inverse supply. It tells us for any output y what the price p needs to be to get the firm to supply the output. We can invert it to get the supply function: $y = \frac{p}{10}$.

13.7 What can go wrong– Linear/Increasing Returns to Scale

If returns to scale are linear or increasing, then if we can find any output level y where the firm earns positive profit there is no profit maximizing level of y. The firm wants to produce as much as possible. This is because with linear or increasing returns to scale, doubling inputs will double cost and at least double output, so profit will at least double. Thus, if we can find a point were profit is positive, we can always use more of all inputs and increase profit.

Let's see this in an example. Suppose $f(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$. Price of output is $p = 100$ and $w_1 = 1$, $w_2 = 1$. In this case, the cost minimizing level of inputs are (try this yourself using cost minimization): $x_1 = x_2 = y$. The cost function is: $c(y) = 2y$. The profit function in terms of y is: $\pi(y) = 100y - 2y = 98y$. This profit function is increasing in y... there is not profit maximizing solution!

As another example, suppose you want to maximize profit using production function f$(x_1, x_2) = min\left\{\frac{1}{2}x_1, x_2\right\}$. To minimize costs, the firm should use: $\frac{1}{2}x_1 = x_2$. Plug this back into the production function to get the conditional factor demands: $x_1 = 2y$ and $x_2 = y$. The cost function is $c(y) = (2w_1 + w_2)y$.

The conditional profit function is: $\pi(y) = py - (2w_1 + w_2)y = (p - 2w_1 - w_2)y$. If $p > 2w_1 + w_2$ there is no profit maximizing level, and thus the firm wants to produce as much as possible. If $p < 2w_1 + w_2$ optimal level is $y = 0$ and profit is 0. If they are equal the profit is always zero and the firm can choose whatever they want.

14 Monopoly

14.1 Monopolies and the Price-Taking Assumption

A single firm serving a market cannot reasonably assume that price is fixed in their output. If a monopoly wants to sell 100 units of a good, they will try to sell it at the highest price they can. What will consumers pay for 100 units? They will pay $p(100)$ where $p()$ is the inverse demand function. Suppose demand is: $Q(p) = \frac{\frac{1}{2}(200)}{p}$. Then inverse demand is p$(Q) = \frac{\frac{1}{2}(200)}{Q} = \frac{100}{Q}$. The most this monopolist could charge to sell 100 units is \$1. If the monopolist wants to sell 200, the most they could charge is $p(200) = \frac{100}{200} = \frac{1}{2}$. As y increases, they amount they can charge will decrease. The point of this is that a monopolist cannot possibly take price as fixed. They have to take into account the fact that they can charge more if they produce lower output and less if they produce higher output.

14.2 The Monopolist's Profit Function

The price that the monopolist can sell y units of a good for is the inverse demand function $p(y)$. We write the profit function as: $\pi(y) = p(y)y - c(y)$.

$$\boxed{\text{Profit function: } \pi(y) = p(y)y - c(y)}$$

The revenue is $p(y)\,y$ and the cost is $c(y)$. The first-order condition for the firm is still setting marginal revenue to marginal cost. Suppose that was not the case. If marginal revenue is higher, then the firm can increase revenue more than cost by increasing output. This will increase profit. If marginal cost is more than marginal revenue, decreasing output will increase profit by lowering cost more than revenue. The key difference between this and perfect competition (where p is fixed) is the the marginal revenue now has an extra term and takes into account the indirect effect of increasing output on price.

14.3 Example of Maximizing Profit

Suppose demand is $q_d(p) = 100 - p$. Cost is $c(q) = 10q$. The inverse demand function is: $p = 100 - q$. The firm's profit function is: $\pi(q) = (100 - q)\,q - 10q$. The first order condition is: $\frac{\partial((100-q)q-10q)}{\partial q} = 0$ which is $90 - 2q = 0$.

Solving this for q gives us the optimal level of output: $q^* = 45$. Plugging this into the inverse demand function gives us the price the monopolist can sell for $p = 55$.

We can now calculate the firm's profit: $\pi(45) = (55)(45) - 10(45) = 2025$.

What if we consider this monopolist to be a price taker instead? Let price be some fixed amount p. A price take has the profit function $\pi(q) = pq - c(q)$. At any $q > 0$ which maximizes profit the derivative must be zero. Thus for a price taker we get: $p - \frac{\partial c(q)}{\partial q} = 0$ or $p = \frac{\partial c(q)}{\partial q}$ since $\frac{\partial c(q)}{\partial q}$ is the marginal cost, we can write this as $p = mc$. Thus, for price takers, price will always be equal to marginal cost if they are selling $q > 0$. In this case, that implies that for a price taker, $p = 10$ since marginal cost of $c(q) = 10q$ is 10. At a price of 10, consumers will buy 90 units and the firms profit is: $\pi = 10(90) - 10(90) = 0$. Compare again to the monopolist actual solution: $p = 55, q = 45, \pi = 2025$.

I will reiterate that price-taking is not a valid assumption for monopolists. The effect that a monopolist's output has on the price in the market cannot be ignored like it can if there are many firms and each firm is just a small part of the market. In this case, we can see how far off that assumption gets us. The predicted quantity for the monopolist is half of the price-taking quantity, and the predicted price is more than five times higher!

14.4 What does a monopoly do?

Inelastic demand implies that raising price by 1 percent lowers demand by less than one percent. On the other hand, lowering quantity by one percent allows them to raise price by more than one percent. This implies lowering quantity by one percent will increase revenue because price increases proportionally more than quantity decreases. Lowering quantity will also lower costs. This has to increase profit! This tells us that if a monopolist is acting optimally, it will always continue lowering quantity as long as demand remains inelastic, and can only be choosing an optimal quantity when demand is elastic. Thus, a monopolist will always operate in the elastic portion of demand (assuming one exists).

Looking again at the example in the previous section (Section 14.3)

$q = 100 - p$ and $c(q) = 10p$, we found the firm operated where $p = 55$. Let's check this is in the inelastic portion of the demand curve. The elasticity is:

$$\frac{\partial(100-p)}{\partial p}\frac{p}{100-p} = -\frac{p}{100-p}$$

This is inelastic when $-\frac{p}{100-p} < -1$ which occurs where $50 < p < 100$. Since $p = 55$ in that example, we confirm the firm is operating in the *elastic* portion of the demand curve.

We can calculate something called a "markup" using the first order condition. I will go over the derivation in class but you are not responsible for knowing exactly how it is derived. The markup calculation is as follows: $p = \frac{\epsilon}{\epsilon+1}mc$. This says the firm marks up price $\frac{\epsilon}{1+\epsilon}$ over their cost. For example, suppose $\epsilon = -2$. This is slightly elastic demand and $\frac{\epsilon}{1+\epsilon} = \frac{-2}{-2+1} = 2$. So the firm will charge 2 times more than their cost.

14.5 In Action

Suppose a firm has a cost function $c(y) = y$. Thus, marginal cost is constant at 1. Suppose demand is $y = \frac{100}{p^2}$. Let's construct the monopolist's profit function. First we need the inverse demand: $p = \left(\frac{100}{y}\right)^{\frac{1}{2}}$. Profit is $\pi = \left(\frac{100}{y}\right)^{\frac{1}{2}} y - y$. The first order condition is $\frac{\partial\left(\left(\frac{100}{y}\right)^{\frac{1}{2}} y - y\right)}{\partial y} = 5\sqrt{\frac{1}{y}} - 1 = 0$. This is solved where $y = 25$. At this price, the consumers will pay $p = \left(\frac{100}{25}\right)^{\frac{1}{2}} = 2$ and this is what the firm will charge. Notice, the firm marks up price 2 times over marginal cost (which is 1).

Let's check the markup. First we need the consumer's elasticity of demand: $\epsilon = \frac{\partial y}{\partial p}\frac{p}{y} = \frac{\partial\left(\frac{100}{p^2}\right)}{\partial p}\frac{p}{\frac{100}{p^2}} = -2$. This gives a markup of $\frac{-2}{-2+1} = 2$ and confirms the price the firm is charging. But notice we could have simplified our work by using elasticity in the first place.

Suppose you knew consumers have a constant elasticity demand of -1.5 (demand is $y = \frac{100}{p^{1.5}}$) and the firm has a constant marginal cost of 1. What would the monopolist charge? We calculate the markup: $\frac{-1.5}{-1.5+1} = 3$. Thus, they would charge 3. At a price of 3 consumers demand $y = \frac{100}{3^{1.5}} \approx 19.245$. Thus, we get the price and quantity without even needing to write down the firm's profit function.

14.6 Consumer Surplus Under Monopoly

As we saw above, suppose a firm had constant marginal cost of $c(y) = 10y$ and demand is $q = 100 - p$, then, if the firm acts as a monopolist, it will sell 45 units at a price of 55.

If we calculate the producer surplus it is the area below the inverse demand of $p = 100 - q$ but above the price of 55. This is the triangle labeled in the graph below. It has an area of $\frac{45*(100-55)}{2} = 1012.5$. Thus, despite trying their best to capture as much of consumer surplus as possible by finding the q that will maximize profit and setting p as high as possible, the firm leaves some surplus on the table. How can they get even more? They will need to charge different prices. We will look at that in the next chapter.

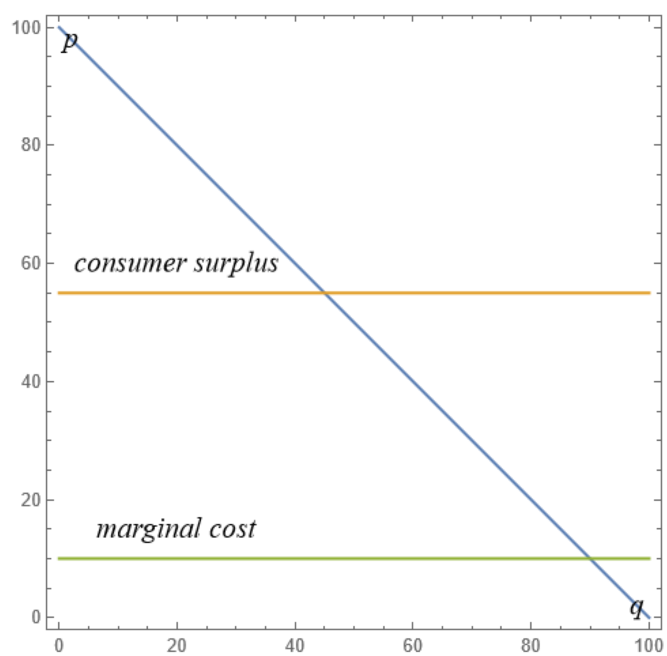

Figure 14.1: Monopoly Profit, Consumer Surplus, and Deadweight Loss.

15 Monopoly Behavior

In the pervious chapter, we looked at what a monopolist can achieve by leveraging their market power to create scarcity and drive up the price in order to increase profit. However, as seen in figure 14.1, the monopolist ends up leaving a lot of surplus on the table both in terms of remaining consumer surplus (charging some consumers less than they are willing to pay) and deadweight loss (lost surplus due to excluding some consumers from the market through the artificial scarcity). In this chapter, we look at more complex ways a monoplist can sell their products to try and capture some of this surplus.

15.1 Types of Price Discrimination

There are several methods monopoilists use when pricing their goods to get more profit. Here is a list:

- First Degree Price Discrimination: The firm can identify every consumer, learn their willingness to pay, and charge different prices.
  - This is an extreme form of price discrimination. It is best understood as a thought experiment about what an extreme monopolist could do rather than something actually achievable.
  - Examples: Airlines sometimes come close to this when they use complex pricing schemes to try and extract more and more surplus from consumers. Everyone on the airplane probably paid a different price for their tickets, but it is unlikely they all paid their highest willingness.
- Second Degree Price Discrimination: The firm cannot identify individual consumers, but can offer different packages or qualities of goods at different prices.
  - Examples: Quantity discounts, quality differences (first-class/coach tickets, "reserve" wines, "flagship" high-end products that differ little from cheaper counterparts.
- Third Degree Price Discrimination: Can identify groups and charge those groups different amounts.
  - Examples: Student tickets, senior discounts.
- Bundling: Combine different goods and force consumers to buy them in bundles.
  - Examples: Cable TV packages, Microsoft Office Software Bundle.
- Two-Part Tariff: Charge the consumer an entry fee for membership that gives the consumer the right to buy the good at the lowest efficient price (the marginal cost of the firm).
  - Examples: Netflix (compare this to streaming rental services), theme park tickets (rides are free), free-coffee for the month when you pay $19.99 to buy a special mug.

15.2 First Degree Price Discrimination in Action

Suppose there are three people willing to pay $3, $2, $1 for a good respectively. Suppose the monopolist has zero marginal cost. Here is the monopoly profit at different prices if it charges everyone the same price:

Price	# Buyers	Profit
$3	1	$3
$2	2	$4
$1	3	$3

The most the firm can make is \$4. But if it knew everyone's willingness to pay and could charge them different prices, the firm could earn \$6!

When the monopolist charges one price, the consumers get some consumer surplus (refer to Figure 14.1) This is because there are some consumers who get the good at a price lower than they are willing to pay. On other hand, there is also some deadweight loss because the monopolist restricts quantity from the efficient level where price is equal to marginal cost. When a monopolist uses first degree price discrimination there is no deadweight loss and they capture all of the consumer surplus! This is because the firm can sell to everyone who it is efficient to sell to (they are willing to pay more than marginal cost) at exactly the price they are willing to pay.

15.3 Third Degree Price Discrimination in Action

Suppose there are two groups of people: students and non-students. A movie theater sells tickets to both groups. Assume the firm has zero marginal cost so that $c(y) = 0$ (cost is zero regardless of output). Students have demand function: $y_s = 100 - 2p$ and non-students have demand function: $y_n = 100 - p$.

If we add up both types of consumer, entire market demand is: $Y = 100 - 2p + 100 - p = 200 - 3p$ (as long as $p \leq 50$). The inverse demands for both groups, and the market as a whole are $p_s = \frac{100-y_s}{2}, p_n = 100 - y_n, p = \frac{200-Y}{3}$.

Suppose the monopolist was going to set one price for the entire market. Their profit function would be: $\pi = \frac{200-Y}{3}Y$. By taking the first order condition and solving we find that the optimal $Y = 100$ and the optimal price is $p = \frac{100}{3}$. At this price $y_s = \frac{1}{3}(100)$ (about 33) is the student demand and $y_n = \frac{2}{3}(100)$ about 66 is the non-student demand. The firm's profit is: $\pi \approx 3333.33$.

What if the firm wanted to set prices differently for students and non-students?

The profit earned from students is: $\pi_s = \frac{100-y_s}{2}y_s$. The profit earned from non-students is: $\pi_n = (100 - y_n)\,y_n$. Solving the first-order conditions, we get that the optimal $y_s = 50$ and the optimal $y_n = 50$. The prices the firm can charge are $p_s = 25$ and $p_n = 50$. The profits are: $\pi_s = 1250$ and $\pi_n = 2500$.

The total profit is: $\pi = \pi_s + \pi_n = 3750$. Notice the firm can earn about 416.67 more by setting different prices!

15.4 Bundling

Bundling can occur when a firm sells multiple products. The goal of bundling is to take advantage of differences in types of demand by forcing consumers to buy bundles of goods at a single price rather than selling each good at a separate price.

For example, suppose a firm sells pants and shirts. There are two consumers who each demand up to one shirt and one pair of pants. They are willing to pay the following:

	Shirt	Pants	Both
Consumer 1	50	30	80
Consumer 2	10	80	90

Pricing Shirts.

If they price shirts at \$50, they sell one shirt and earn \$50. If they price at \$10, they sell two shirts and earn \$20.

Pricing Pants.

If they price pants at \$80, they sell one pair of pants and earn \$80. If they price at \$30, they sell two pairs of pants and earn \$60.

Thus, the best they can do is sell one shirt at \$50 and one pair of pants at \$80 to earn \$130.

Pricing Bundles.

If the firm forces consumers to buy a bundle of a shirt and a pair of pants they can price that bundle at \$80, sell two bundles and earn \$160.

15.5 Two-Part Tariff

Two-part tariffs can be used when consumers demand multiple units of a good. An example of this is theme park tickets. The theme park could charge a price per ride. In fact, this happens at some fairs. However, instead, rides are free once you have purchased the ticket. The goal of a two-part tariff is to create as much consumer surplus as possible by selling the consumer as much as is efficient (this occurs where price is marginal cost). This will create the most consumer surplus possible. Instead of leaving that consumer with the surplus, charge them an "entry fee" (this is the other part of the tariff) equal to their consumer surplus.

For example, suppose a consumer's demand for coffee is $q = 10 - p$ and the firm has zero marginal cost for coffee. If the firm sells to that consumer at a single price it's profit of selling the consumer q cups of coffee at the most they will pay for those q cups is: $\pi = (10 - q)\, q$. The best thing to do is sell them 5 cups of coffee at 5 dollars and earn \$25.

If the firm prices at marginal cost (\$0) the consumer will demand 10 cups of coffee. Their surplus is the area below the inverse demand but above price of zero. That surplus is \$50, so they would be willing to pay up to \$50 for the right to buy cups of coffee at \$0 (assuming you don't give them the option to buy at \$5 per cup). So the firm can earn \$50 by forcing the consumer to pay an "entry fee" of \$50 and then give them coffee for free.

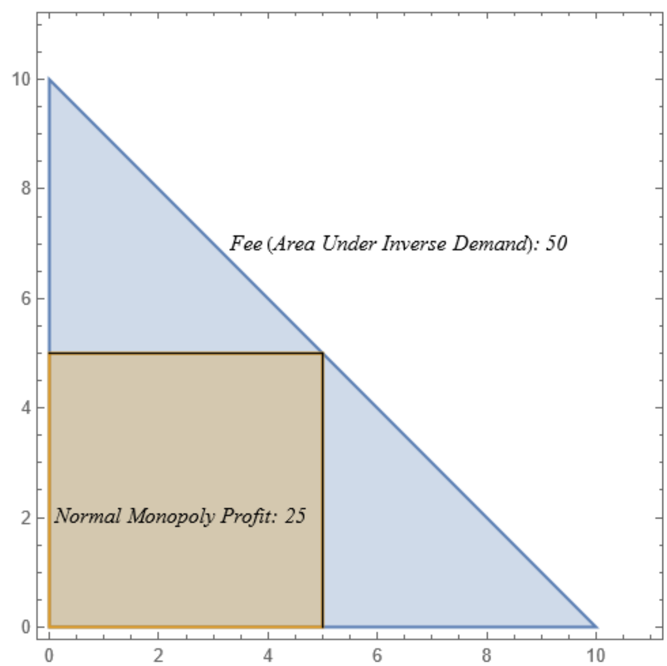

Figure 15.1: Earning more with a two-part tariff. Orange is the profit under optimal unit pricing. Blue plus orange area is the consumer surplus under marginal cost pricing that is then captured with an entry fee.

16 The Cournot Model of Competition

16.1 Extending the Monopoly Model

To relax the price-taking assumption, we need to know the relationship between a firm's quantity and the market price. For a monopoly, that relationship is easy to figure out, it is the inverse demand. The inverse demand is the *most* that consumers will pay to buy q units of the good. So, if a monopolist produces q units of the good, they can charge the inverse demand $p(q)$ at that quantity.

But what about when there are multiple firms? The inverse demand represents the relationship between the price and the market quantity. But when there are many firms, each firm's quantity is only a part of the market quantity. The first set to extending the model is to make sure we can keep track of each firm's quantity, and the total market quantity. We use the following notation:

q_i: firm i's quantity

Q: Total (market) quantity $Q = \sum_{i=1}^{n} q_i$

Q_{-i}: Total (market) quantity of all firms except i. $Q_{-i} = Q - q_i$.

We know that the amount consumers will pay for the total **market quantity** is the inverse demand $p(Q)$. Thus, the profit of a firm is dependent both on the quantity they produce and the total market quantity. This lets us write each firm's profit function like this: $\pi(q_i, Q) = p(Q)\, q_i - c(q_i)$.

Notice here, the market price is determined by the total market quantity. Revenue is that price multiplied by the firms own quantity. Of course, costs are specific to a firm and dependent only on their own quantity. One issue with writing the profit this way is that Q itself depends on q_i. It is more convenient to write the profit function in terms of q_i and Q_{-i}. Noting that $Q = Q_{-i} + q_i$, we can write: $\pi(q_i, Q_{-i}) = p(Q_{-i} + q_i)\, q_i - c(q_i)$.

If the firm knew (or even had some assumption about) what Q_{-i} is, they could calculate their profit for any quantity q_i they produce and even maximize it. Let's look at an example.

16.2 Example of Maximizing Profit with Two Firms

Suppose inverse demand is $p(Q) = 100 - Q$, there are two firms, and the cost function of each firm is $c(q_i) = 10q_i$.

Firm 1's profit function is $\pi_1(q_1, q_2) = (100 - q_1 - q_2)\, q_1 - 10q_1$ which simplifies to $\pi_1(q_1, q_2) = 90q_1 - q_i^2 - q_1 q_2$.

Similarly 2's profit function is $\pi_2(q_1, q_2) = 90q_2 - q_2^2 - q_1 q_2$.

16.2.1 Game Theory

This model is a **Game**. A game is a formal mathematical object studied in game theory. For our purposes, we can think of a game as a set of players, a set of actions for each player, and a set of payoffs for each player that depend on the actions chosen by everyone.

There are many ways to "solve" a game. That is, to make predictions about what might happen in that game given the strategic sophistication of players in the game. The most common way to solve a game in game theory is to use the **Nash equilibrium.** To define a Nash Equilibrium, let's first look at **best responses.**

16.2.2 Nash Equilibrium for 2 Firms

Suppose firm 1 believes firm 2 will produce $q_2 = 50$. Then they believe their profit function is $\pi_1(q_1, 50) = 40q_1 - q_1^2$. To maximize this, find where the derivative is zero. This occurs at $q_2 = 20$.

In game theory, $q_1 = 20$ is what we call a **best response** to $q_2 = 50$.

When a set of quantities are *simultaneously* best responses to each other, we say that are a **Nash Equilibrium** of the game. Since Nash Equilibrium requires all players to be simultaneously best responding to each other, it will also imply that no player has incentive to change their action.

While $q_1 = 20$ is a best response to $q_2 = 50$. Firm 2's profit if firm 1 produces $q_1 = 20$ is: $\pi_2(20, q_2) = 90q_2 - q_2^2 - 20q_2$ which is maximized at 35. Thus, $q_2 = 35$ is the best response to $q_1 = 50$ and so the pair $(20, 50)$ is **not** a Nash equilibrium since firm 2 is not best responding in choosing $q_2 = 50$.

To find the Nash equilibrium, let's look again at each firm's profit function:

$\pi_1(q_1, q_2) = (100 - q_1 - q_2)\, q_1 - 10q_1 = 90q_1 - q_1^2 - q_1q_2$.

$\pi_2(q_2, q_1) = (100 - q_1 - q_2)\, q_2 - 10q_2 = 90q_2 - q_2^2 - q_2q_1$.

To find firm 1's best response to any q_2, maximize their profit by finding where their marginal profit is zero (the first-order condition). The marginal profit is: $\frac{\partial\left(90q_1 - q_1^2 - q_1q_2\right)}{\partial q_1} = 90 - 2q_1 - q_2$. This is equal to zero where $q_1 = \frac{90-q_2}{2}$. This is firm 1's **best response function**. It tells firm 1 what quantity to choose given what firm 2 chooses. Firm 2's best response is likewise $q_2 = \frac{90-q_1}{2}$.

A **Nash Equilibrium** is a pair (q_1, q_2) that solves both of these at the same time. That is, firm 1 is best responding to firm 2 and firm 2 is best responding to firm 1.

Solving this system of equations, $\left\{q_1 = \frac{90-q_2}{2}, q_2 = \frac{90-q_1}{2}\right\}$, the only pair of quantities that solve this are $q_1 = 30, q_2 = 30$.

16.3 Equilibrium with N firms.

Now suppose we have N firms. Each firm's profit is $\pi_i(q_i, Q_{-i}) = (100 - q_i - Q_{-i})\, q_i - 10q_i$ which simplifies to $\pi_i(q_i, Q_{-i}) = 90q_i - q_i^2 - q_iQ_{-i}$. If we maximize the profit for firm i conditional on Q_{-i}, we need to look for where the marginal profit of i is zero. This occurs where $\frac{\partial\left(90q_i - q_i^2 - q_iQ_{-i}\right)}{\partial q_i} = 90 - 2q_i - Q_{-i} = 0$. Solving this for q_i gives firm i's best response function to Q_{-i} by the other firms $q_i = \frac{90-Q_{-i}}{2}$.

To find the Nash equilibrium, we would need to solve N equations of the form $q_i = \frac{90-Q_{-i}}{2}$ for the n unknowns. However, notice each firm's best response function looks identical because they all have **the same cost function**. When this is the case, instead of solving N equations to find the N firm's quantities in equilibrium, we can **impose symmetry** on the set of equations. That is, assume in equilibrium all firms will produce the same output. *There will always be an equilibrium like this in our simple models when the costs are symmetric.* Thus, we can assume that $q_i = q_j = q$ for all i and j and solve just one equation for one unknown... q. Noting that $Q_{-i} = (N-1)\, q$ since there are $(N-1)$ firms that aren't firm i, imposing symmetry, we get the equation: $q = \frac{90-((N-1)q)}{2}$.

Solving this for q gives us $q^* = \frac{90}{N+1}$. Thus, in equilibrium, with N firms, all will produce $q = \frac{90}{N+1}$. The market quantity in equilibrium will be $Q = (N)\, q^* = \frac{N}{N+1}90$ and the market price will be. $p^* = 100 - \frac{N}{N+1}90$. Let's look at this market price and market quantity as the number of firms changes:

N	p	Q
1	55	45
2	40	60
5	25	75
100	10.9	89.1
1000	10.1	89.9

Figure 16.1: Price and quantity in Nash equilibrium of the Cournot model with N firms when each has cost $c(q) = 10q$ and inverse demand is $p = 100 - Q$.

For $N = 1$ we get the monopoly solution $p = 55$, $q = 45$. As N increases, the price approaches 10 and the market quantity approaches 90. This is what would occur if firms were price takers. That is, under perfect competition. This is because as N increases, each firm becomes a less less important part of the market, their market power shrinks to zero and they become, effectively, price takers.

17 Collusion and Cooperation

17.1 A Basic Example

Let's start with a two firm Cournot model. Inverse demand is $p(Q) = 25 - Q$, and the cost function of each firm is $c(q_i) = q_i^2$. We can write the profit functions of each firm as follows:

$$\pi_1(q_1, q_2) = (25 - (q_1 + q_2))\, q_1 - q_1^2$$

$$\pi_2(q_1, q_2) = (25 - (q_1 + q_2))\, q_2 - q_2^2$$

If each firm maximizes its profit by choosing their quantity, we get the best response functions (see the section on the Cournot model for details on this): $q_1 = \frac{25-q_2}{4}$, $q_2 = \frac{25-q_1}{4}$. Since the firms have the same cost function, we can impose symmetry and solve to find the symmetric Nash equilibrium. This is where $q = \frac{25-q}{4}$. The solution to this is $q = 5$, which gives the Nash equilibrium of the game. When both firms play the Nash equilibrium, both earn $\pi = 50$.

Now, suppose the firms get together and collude to maximize their joint profit. They look for a quantity they can both choose that maximizes the sum of their profits $\pi_{joint} = 2\left((25 - (q + q))\, q - q^2\right)$. Taking the derivative of this with respect to q we get $2(25 - 6q)$. Solving for where this is equal to zero gives us the quantity that jointly maximizes their profits. This is $q_{collusion} = \frac{25}{6}$. When both firms play this collusive quantity, they both earn $\pi_{collusion} = \frac{625}{12} \approx 52.1$.

Note that collusion is *not* an equilibrium. Playing $\frac{25}{6}$ when the other plays $\frac{25}{6}$ is sub-optimal. The optimal strategy is to best respond to $\frac{25}{6}$. Plugging $\frac{25}{6}$ in for the quantity of the other firm in the above the best response functions gives us: that given the other will stick to the agreement of playing $\frac{25}{6}$, it is actually best to play $q_{deviating} = \frac{125}{24} \approx 5.2$. By doing this the deviating firm gets $\pi \approx 54.3$ and the other firm who sticks to $\frac{25}{6}$ gets $\pi = 47.7431$. If both happen t deviate by playing $\frac{125}{24}$, they both get $\pi \approx 48.8281$.

If we boil down this scenario for the firms to the choice of whether to cooperate by playing the collusive quantity, or to back-stab the other firm and deviate by playing $\frac{125}{24}$, we get what game theory called a "2x2 game". These are the simplest "games" in game theory. There are many of these 2x2 games. The one that results from this particular strategic scenario is probably the most famous. It is called the **prisoner's dilemma**. In a prisoner's dilemma game, each player has incentive to deviate from some cooperative outcome regardless of what they think the other person will do. We can represent this 2x2 game in the following table. Firm 1's strategy determines the row and firm 2's determines the column. Notice that regardless of what the other firm chooses, a firm always does better by deviating.

	cooperate $q_2 = \frac{25}{6}$	deviate $q_2 = \frac{125}{24}$
cooperate $q_1 = \frac{25}{6}$	$52.1, 52.1$	$47.7, \mathbf{54.3}$
deviate $q_1 = \frac{125}{24}$	$\mathbf{54.3}, 47.7$	$\mathbf{48.8}, \mathbf{48.8}$

Figure 17.1: The "prisoner's dilemma" game resulting from two firms agreeing to collude in a market where both have cost function $c(q) = q^2$ and inverse demand is $25 - Q$. Firm 1's decision determines the row, firm 2's decision determines the column. Payoffs are written with firm 1's payoff first.

17.2 Prisoner's Dilemma

In the previous section, we saw how two firms agreeing to collude in a market results in a strategic scenario known in game theory as the "prisoner's dilemma". In this section, we will simplify the payments from that game slightly and look at how, despite having incentive to deviate from the agreement when played one time, firms that repeat this game over and over can agree to cooperate in a way that is self-enforcing.

Let's suppose we worked through the analysis for two firms colluding as in the previous section and got this simplified game instead of the one in Figure 17.1.

	cooperate	defect
cooperate	4,4	0,10
defect	10,0	2,2

Figure 17.2: A simplified prisoner's dilemma game.

Now, let's suppose these firms play this game repeatedly, forever. Each firm is a little impatient and likes money in the next period β (with $0 < \beta < 1$) as much as they like money today. And they like money in two periods β times as much than money one periods from now so that they like money in two periods β^2 as much as money right now... and so on. The payoff of a firm in this game can be represented as a discounted stream of the stage-game payoffs like this. If time $t = 0$ is now the payoffs from getting payoffs π_0 now, π_1 next period, π_2 in the period after that , and so on:

$$\sum_{t=0}^{\infty} \beta^t \pi_t = \pi_0 + \beta\pi_1 + \beta^2\pi_2 + \beta^3\pi_3 + ...$$

As an example, suppose the firms cooperate in every period. Using the fact that $\sum_{t=0}^{\infty} \beta^t = \frac{1}{1-\beta}$, the payoff is:

$$4 + 4\beta + 4\beta^2 + 4\beta^3 + ... = \sum_{t=0}^{\infty} 4\beta^t = 4\sum_{t=0}^{\infty} \beta^t = \frac{4}{1-\beta}$$

17.3 Sustaining Cooperation

We have seen that if firms play this game once, both have incentive to deviate from cooperation. However, what if they repeat the game? If they are patient enough, the following agreement will be self-enforcing: "Cooperate as long we have always (both) cooperated in the past. If either of us have ever deviated, we will both deviate forever."

To see that this is self-enforcing, we need to check whether any firm ever has incentive to deviate from the agreement as long as the other is following the agreement. If someone has deviated in the past, then the other firm will deviate forever, no matter what. The best a firm can do is to deviate for ever. What if no one has deviated? Does a firm have incentive to deviate? If they go along with the agreement, they will cooperate forever and get 4 in every period. If they deviate, they will get 10 today, but both of the firms will then deviate forever after. We can summarize the payments of both options:

Follow the agreement and cooperate: $4 + \beta\,(4) + \beta^2\,(4) + \beta^3\,(4) + ... = 4\sum_{t=0}^{\infty} \beta^t = \frac{4}{1-\beta}$

Deviate: $10 + 2\beta 2 + 2\beta^2 + 2\beta^3 + ... = 10 + \sum_{t=1}^{\infty} 2\beta^t = 10 + \beta\sum_{t=0}^{\infty} 2\beta^t = 10 + \beta\frac{2}{1-\beta}$

Cooperating is better than deviating, and thus the agreement is self-enforcing as long as $\frac{4}{1-\beta} > 10 + \beta\frac{2}{1-\beta}$. Solving this for β, we get $\beta > \frac{3}{4}$. As long as these firms are patient enough (they care about tomorrow at least $\frac{3}{4}$ what they care about today) then they can sustain cooperation.

18 Externalities

Externalities occur when the choices of one person affect the outcomes of other people. For instance, if someone smokes, it can affect the health of others. If one person drives to work, it creates a little more traffic congestion that slows down everyone else (even if it is just by a little bit). These are both externalities. Let's see how externalities can create inefficient outcomes in simple models and how policy can improve efficiency.

18.1 Tragedy of the Commons Example

A simple model that demonstrates negative externalities without a lot of complexity is the *tragedy of the commons* model. In this model, a group of people share a common resource. When use of this resource is unregulated the outcomes will be inefficient because everyone ignores how their use of the resource diminishes the effectiveness for everyone else.

For example, suppose $100\sqrt{B}$ fish will be caught on a lake when B boats are fishing. Fish can be sold for \$1 but it costs \$10 to buy fuel and supplies to fish. If there are B boats, each catches $\frac{100\sqrt{B}}{B}$. The profit of each boat is: $\pi = \frac{100\sqrt{B}}{B} - 10 = \frac{100}{\sqrt{B}} - 10$. Notice that as more boats come to the lake, it reduces the profit of all other boats, a negative externality.

How many boats will be on the lake if fishing is unregulated? Each boat will decide to fish as long as they can earn positive profit. That is, as long as $\frac{100}{\sqrt{B}} - 10 \geq 0$. Solving this for B we get $B \leq 100$. Thus, if fishing is not regulated, there will be 100 boats on the lake. 1000 fish will be caught. Each boat catches 10 fish and earns \$0 profit.

The total profit earned by all B fishing boats is $\pi(B) = B\left(\frac{100\sqrt{B}}{B} - 10\right) = 100\sqrt{B} - 10B$. This is maximized where $\frac{\partial\left(100\sqrt{B}-10B\right)}{\partial B} = 0$. Solving this, we get $B = 25$. Thus, the optimal number of boats to have on the lake is 25. They would catch 500 fish and each boat would earn \$10.

Suppose the government wants to charge for a permit to fish $\$p$ to get the optimal number of fish on the lake. What permit fee will bring the number of boats to the efficient level? We need to find what price p will make profit of each boat zero when there are 25 boats $\pi = \frac{100\sqrt{B}}{B} - 10 - p$. Plugging in $B = 25$ and solving for where the profit is zero: $\frac{100\sqrt{25}}{25} - 10 - p = 0$. We get $p = 10$. If the government charges a \$10 fishing fee, the efficient number of boats will fish.

18.2 Positive Externalities - Public Goods

As we have seen above, when there are negative externalities, people tend to use "too much" of something that makes others worse off. The opposite happens when there are positive externalities. When using or doing something makes others better, off individuals will tend to not do it enough.

For example, suppose 100 people share a park and are asked to donate money to it. Each individual has income m_i and their contribution is g_i. The total contributions are $G = \sum_{i=1}^{100} g_i$. Each individuals utility function is quasilinear in money and the total contributed to the public good and is given by $u(g_i) = m_i - g_i + 100\sqrt{G_{-i} + g_i}$ where $G_{-i} = G - g_i$. For convenience, let's assume $m_i = 1000000$ for all consumers.

What does each individual contribute? The individual marginal utility of contribution is $\frac{\partial\left(10000 - g_i + 100\sqrt{G_{-i}+g_i}\right)}{\partial g_i} = \frac{50}{\sqrt{g_i+G_{-i}}} - 1$. This is maximized where $\frac{50}{\sqrt{g_i+G_{-i}}} = 1$ or where $g_i = 2500 - G_{-i}$. This is each individual's **best response.** Notice that each individual has the same best response function. We can look for a **symmetric Nash equilibrium** here just as we did in the Cournot model. If we impose symmetry by having all contribute g then we get the best response function: $g = 2500 - (100 - 1)\,g$ which is solved by $g = 25$. Thus, in the symmetric Nash equilibrium, all contribute \$25.

Each individual's utility is $1000000 - 25 + 100\sqrt{2500} = 1004975$. Could we make these individuals better off?

18.2.1 Social Optimum

What if the government could tax individuals and use those taxes to contribute to the park? What would the optimum tax be? Suppose we set the tax to t for everyone. Each individual's utility function is: $u(t) = 10000 - t + 100\sqrt{100t}$. Let's maximize this with respect to t. $u(t) = 1000000 - t + 100\sqrt{100t}$. The marginal utility with respect to t is $\frac{\partial\left(1000000 - t + 100\sqrt{100t}\right)}{\partial t} = \frac{500}{\sqrt{t}} - 1$. Utility for every individual is maximized where this derivative is equal to zero. Thus, utility is maximized where $500 = \sqrt{t}$ or where $t = 250000$. At this tax level the utility of every individual is $1000000 - 250000 + 100\sqrt{100\,(250000)} = 1250000$.

19 Exchange

19.1 Pies - An Equilibrium

Suppose there is a baker who has 30 crusts and a farmer who has 60 apples. Both eat only pies (perfect complements) that use 1 crust and 2 apples. Prices are $p_1 = 2, p_2 = 1$. Since it requires two apples and a crust to make a pie, a pie costs \$4 at these prices. Notice there are ingredients for 10 pies available between the endowments of the two consumers. We call such a model an **exchange economy**.

Since the baker has 30 crusts and each is worth \$2, this endowment is worth \$60. The baker can afford 15 pies. Since the farmer has 60 apples and each is worth \$1, this endowment is worth \$260. The farmer can afford 15 pies. At these prices, the baker demands 15 crusts, 30 apples and the farmer demands: 15 crusts, 30 apples.

At these prices, the pair would agree on the following trade: the farmer gives 30 apples to the barker in exchange for 15 crusts. Whenever the sum of the demands at some set of prices is exactly equal to the total endowments of the goods in the economy, a mutually agreeable trade like this is possible. We call this an equilibrium. An **equilibrium** in an exchange economy is a set of prices such that the demands of each good at those prices sum to the total endowment of those goods. That is, there is no over supply or over demand of any of the goods at these prices.

19.2 Pies - Not an Equilibrium

Suppose there is a baker who has 30 crusts and a farmer who has 60 apples. The baker eat pies that use 1 crust, 2 apples. The farmer eats anything (perfect substitutes utility $u = x_1 + x_2$). Prices are $p_1 = 2, p_2 = 1$. Notice the only difference between this exchange economy and the previous example is the farmer's preferences.

Again, a pie costs \$4. Since the baker has 30 crusts and each is worth \$2, this endowment is worth \$60. The baker can afford 15 pies. Since the farmer has 60 apples and each is worth \$1, this endowment is worth \$60. The demands are... **baker**: 15 crusts, 30 apples. **farmer**: 0 crusts, 60 apples.

This is **not a equilibrium** because at these prices, the demands of the consumers do not sum to the endowments. There is an over demand for apples (90 demanded with 60 available) and an over supply of crusts (15 demanded with 30 available).

Since there is an over demand for apple and an oversupply of crusts, to find an equilibrium, the price of crusts will need to increase relative to the price of apples. Let's try: $p_1 = p_2 = 1$. At these prices, the baker demands: $(10, 20)$. The farmer demands anything that cots \$60. For instance, the farmer might demand: $(20, 40)$. At these prices and with these demands, total demand for crusts is 30, total demand for apples is 60. There is no over-supply or over demand. This is an equilibrium.

19.3 General Environment

To see how we can find an equilibrium more generally, let's set up the problem. We need the demands and endowments for two people. Let's call them A and B.

Person A: Demands: x_1^A, x_2^A. Endowments: ω_1^A, ω_2^A. Utility: $u_A\left(x_1^A, x_2^A\right)$.

Person B: Demands: x_1^B, x_2^B. Endowments: ω_1^B, ω_2^B. Utility: $u_B\left(x_1^B, x_2^B\right)$.

The prices apply to both people, so we will not need to designate person A or person B. The prices are just as we are used to: p_1, p_2.

19.3.1 Equilibrium Conditions

There are two conditions for equilibrium.

1. Consumers maximize utility $u_i\left(x_1^i, x_2^i\right)$ subject to their budget equations: $p_1 x_1^i + p_2 x_2^i = p_1 \omega_1^i + p_2 \omega_2^i$.
2. Those demands "clear" the market: $x_1^A + x_1^B = \omega_1^A + \omega_1^B$ and $x_2^A + x_2^B = \omega_2^A + \omega_2^B$.

19.3.2 Walras' Law

One useful result is **Walras' Law.** This law tells us that if all but one markets clears, the last one will clear as well. If we just have two markets as in our examples, this allows us to find the prices that will clear one market and be assured they will clear the other market.

19.3.3 Normalizing Prices

Suppose some set of prices is an equilibrium. If we multiply all of the prices by some number, that will also to be an equilibrium. This is because, if we scale all of the prices by, let's say t, each consumer's budget equation goes from $p_1 x_1^i + p_2 x_2^i = p_1 \omega_1^i + p_2 \omega_2^i$ to $tp_1 x_1^i + tp_2 x_2^i = tp_1 \omega_1^i + tp_2 \omega_2^i$. Since we can factor out t from this second equation and get the first equation, the budget set for each consumer must be the same when we scale the prices. Thus, the markets will also clear at prices tp_1, tp_2.

Because an equilibrium only depends on the relative prices of the goods, we are always free to pick one of the goods and set it's price to 1 and then find the price of the other good that brings the markets into equilibrium. This is called normalizing the prices. I usually like to choose good 1 to normalize the prices to by setting $p_1 = 1$. I will demonstrate this in the example below.

19.4 Example

Suppose $u_A\left(x_1^A, x_2^A\right) = x_1^A x_2^A$, $\omega_1^A = 10, \omega_2^A = 0$, $u_B\left(x_1^B, x_2^B\right) = x_1^B x_2^B$, $\omega_1^B = 0, \omega_2^B = 20$.

Find the demands for each consumer by maximizing utility. Consumer A's demands are:

$$x_1^A = \frac{\frac{1}{2}\left(10p_1\right)}{p_1}, x_2^A = \frac{\frac{1}{2}\left(10p_1\right)}{p_2}$$

Consumer B's demands are:

$$x_1^B = \frac{\frac{1}{2}\left(20p_2\right)}{p_1}, x_2^B = \frac{\frac{1}{2}\left(20p_2\right)}{p_2}$$

Now we need to apply the market clearing condition. By Walras' law, if one market clears, the other will clear as well. We only need to check one market. Let's pick market for good 1. For the market to clear we need $x_1^A + x_1^B = \omega_1^A + \omega_1^B$. Plugging in demands and endowments, we get: $\frac{\frac{1}{2}(10p_1)}{p_1} + \frac{\frac{1}{2}(20p_2)}{p_1} = 10$. Let's normalizes prices by setting $p_1 = 1$ and we have $\frac{\frac{1}{2}(10)}{1} + \frac{\frac{1}{2}(20p_2)}{1} = 10$. Solving this for p_2, we get $p_2 = \frac{1}{2}$.

Thus, $p_1 = 1$ and $p_2 = \frac{1}{2}$ are equilibrium prices. More generally, any set of prices such that $p_2 = \frac{1}{2}p_1$ are an equilibrium. Let's figure out what each consumer demands at these prices by plugging in the prices into demands above. $x_1^A = 5$, $x_2^A = 10$, $x_1^B = 5$, $x_2^B = 10$.

19.5 One More General Equilibrium Example.

Suppose we have $u_A = ln\,(x_1) + x_2$, $w_1^A = 10$, $w_2^A = 10$. $u_B = x_1 x_2$, $w_1^B = 10$, $w_2^B = 10$.

Let's find A's demand. Maximizing utility, we get: $x_1^A = \frac{p_2}{p_1}$, $x_2^A = \frac{(10p_1+10p_2)-p_1\left(\frac{p_2}{p_1}\right)}{p_2}$. B's demand is $x_1^B = \frac{\frac{1}{2}(10p_1+10p_2)}{p_1}$ and $x_2^B = \frac{\frac{1}{2}(10p_1+10p_2)}{p_2}$.

Now let's normalize prices by setting $p_1 = 1$ to simplify the expressions.

$$x_1^a = p_2, x_2^a = \frac{10 + 9p_2}{p_2}$$

$$x_1^b = \frac{1}{2}\left(10 + 10p_2\right), x_2^b = \frac{\frac{1}{2}\left(10 + 10p_2\right)}{p_2}$$

Finally, we find the p_2 that makes this an equilibrium by finding the p_2 that makes either market clear. By walras' law we only need to check this for one of the two markets. Let's find the p_2 that makes market one clear by solving $p_2 + \frac{1}{2}\left(10 + 10p_2\right) = 20$. We get $p_2 = \frac{5}{2}$.

19.6 Pareto Efficiency and Equilibrium

Pareto efficient. You can't make someone better off without making someone else worse off.

First welfare theorem. *An equilibrium in an exchange economy has to be pareto efficient.*

Second welfare theorem. Every possible Pareto efficient outcome can be arrived at in equilibrium from some starting set of endowments.

Having the equilibrium outcome of a market be Pareto efficient requires some assumptions. One of the key reasons that markets in a simple exchange economy are Pareto efficient is that there are no externalities. What one consumer chooses does not affect the utility of the other.

In markets that have externalities, we should not expect equilibria to be efficient. In these instances, there is good reason for regulators to get involved to change the market outcomes and improve efficiency. Let's have a look at some simple models with externalities.