Econ 121b: Intermediate Microeconomics

Dirk Bergemann, Spring 2012

1 Introduction

1.1 What's Economics?

This is an exciting time to study economics, even though may not be so exciting to be part of this economy. We have faced the largest financial crisis since the Great Depression. $787 billion has been pumped into the economy in the form of stimulus package by the US Government. $700 billion has been spent on the Troubled Asset Relief Programs for the Banks. The unemployment rate has been high for a long time. The August unemployment rate is 9.7%. Also there has been big debates going on at the same time on health care reform, government deficits, climate change etc. We need answers to all of these big questions and many others. And all of these come under the purview of the discipline of economics (along with other fields of study). But then how do we define this field of study? In terms of subject matter it can be defined as the study of allocation of scarce resources. A more pragmatic definition might be, economics is what economists do! In terms of methodology Optimiz! ation Theory, Statistical Analysis, Game Theory etc characterize the study of Economics. One of the primary goals of economics is to explain human behavior in various contexts, humans as consumer of commodities or decision maker in firms or head of families or politician holding a political office etc. The areas of research extends from international trade, taxes, economic growth, antitrust to crime, marriage, war, laws, media, corruption etc. There are a lot of opportunity for us to bring our way of thinking to these issues. Indeed, one of most active areas of the subject is to push this frontier.

Economists like to think that the discipline follows Popperian methods, moving from Stylized facts to Hypothesis formation to Testing hypothesis. Popperian tradition tells you that hypotheses can only be proven false empirically, not proven true. Hence an integral part of economics is to gather information about the real world in the form of data and test whatever hypothesis that the economists are proposing to be true. What this course builds up, however, is how to come up with sensible hypotheses that can be tested. Thus economic theory is the exercise in hypothesis formation using the language of mathematics to formalize assumptions (about certain fundamentals of human behavior, or market organization, or distribution of information among individuals etc). Some critics of economics say our models are too simplistic. We leave too many things out. Of course this is true - we do leave many many things out, but for a useful purpose. It is better to be clear about an argument! and focusing on specific things in one model helps us achieve that. Failing to formalize a theory does not necessarily imply that the argument is generic and holistic, it just means that the requirement of specificity in the argument is not as high.

Historically most economists rely on maximization as a core tool in economics, and it is a matter of good practice. Most of what we will discuss in this course follows this tradition: maximization is much easier to work with than alternatives. But philosophically I don't think that maximization is necessary for any work to be considered as part of economics. You will have to decide on your own. My own view is that there are 3 core tools:

- The principle that people respond to incentives
- An equilibrium concept that assumes that absence of free lunches
- A welfare criteria saying that more choices are better

Last methodological point: Milton Friedman made distinction of the field into positive and normative economics:

- Positive economics - why the world is the way it is and looks the way it does
- Normative economics - how the world can be improved

Both areas are necessary and sometimes merge perfectly. But there are often tensions. We will return to this throughout the rest of the class. What I hope you will get out of the course are the following:

- Ability to understand basic microeconomic mechanisms
- Ability to evaluate and challenge economic arguments
- Appreciation for economic way of looking at the world

We now try to describe a very simple form of human interaction in an economic context, namely trade or the voluntary exchange of goods or objects between two people, one is called the seller, the current owner of the object and the other the buyer, someone who has a demand or want for that object. It is referred to as bilateral trading.

1.2 Gains from Trade

1.2.1 Bilateral Trading

Suppose that a seller values a single, homogeneous object at c (opportunity cost), and a potential buyer values the same object at v (willingness to pay). Trade could occur at a price p, in which case the payoff to the seller is $p - c$ and to the buyer is $v - p$. We assume for now that there is only one buyer and one seller, and only one object that can potentially be traded. If no trade occurs, both agents receive a payoff of 0.

Whenever $v > c$ there is the possibility for a mutually beneficial trade at some price $c \leq p \leq v$. Any such allocation results in both players receiving non-negative returns from trading and so both are willing to participate ($p - c$ and $v - p$ are non-negative).

There are many prices at which trade is possible. And each of these allocations, consisting of whether the buyer gets the object and the price paid, is efficient in the following sense:

Definition 1. An allocation is Pareto efficient if there is no other allocation that makes at least one agent strictly better off, without making any other agent worse off.

1.2.2 Experimental Evidence

This framework can be extended to consider many buyers and sellers, and to allow for production. One of the most striking examples comes from international trade. We are interested, not only in how specific markets function, but also in how markets should be organized or designed.

There are many examples of markets, such as the NYSE, NASDAQ, E-Bay and Google. The last two consist of markets that were recently created where they did not exist before. So we want to consider not just existing markets, but also the creation of new markets.

Before elaborating on the theory, we will consider three experiments that illustrate how these markets function. We can then interpret the results in relation to the theory. Two types of cards (red and black) with numbers between 2 and 10 are handed out to the students. If the student receives a red card they are a seller, and the number reflects their cost. If the student receives a black card they are a buyer, and this reflects their valuation. The number on the card is private information. Trade then takes place according to the following three protocols.

1. Bilateral Trading: One seller and one buyer are matched before receiving their cards. The buyer and seller can only trade with the individual they are matched with. They have 5 minutes to make offers and counter offers and then agree (or not) on the price.

2. Pit Market: Buyer and seller cards are handed out to all students at the beginning. Buyers and sellers then have 5 minutes to find someone to trade with and agree on the price to trade.

3. Double Auction: Buyer and seller cards are handed out to all students at the beginning. The initial price is set at 6 (the middle valuation). All buyers

and sellers who are willing to trade at this price can trade. If there is a surplus of sellers the price is decreased, and if there is a surplus of buyers then the price is increased. This continues for 5 minutes until there are no more trades taking place.

2 Choice

In the decision problem in the previous section, the agents had a binary decision: whether to buy (sell) the object. However, there are usually more than two alternatives. The price at which trade could occur, for example, could take on a continuum of values. In this section we will look more closely at preferences, and determine when it is possible to represent preferences by "something handy," which is a utility function.

Suppose there is a set of alternatives $X = \{x_1, x_2, \ldots, x_n\}$ for some individual decision maker. We are going to assume, in a manner made precise below, that two features of preferences are true.

- There is a complete ranking of alternatives.
- "Framing" does not affect decisions.

We refer to X as a choice set consisting of n alternatives, and each alternative $x \in X$ is a consumption bundle of k different items. For example, the first element of the bundle could be food, the second element could be shelter and so on. We will denote preferences by $\succsim$, where $x \succsim y$ means that "x is weakly preferred to y." All this means is that when a decision maker is asked to choose between x and y they will choose x. Similarly, $x \succ y$, means that "x is strictly preferred to y" and $x \sim y$ indicates that the decision maker is "indifferent between x and y." The preference relationship $\succsim$ defines an ordering on $X \times X$. We make the following three assumptions about preferences.

Axiom 1. *Completeness. For all $x, y \in X$ either $x \succsim y$, $y \succsim x$, or both.*

This first axiom simply says that, given two alternatives the decision maker can compare the alternatives, and will weakly prefer one of the alternatives to the other, or will be indifferent, in case both are weakly preferred to each other.

Axiom 2. *Transitivity. For all triples $x, y, z \in X$ if $x \succsim y$ and $y \succsim z$ then $x \succsim z$.*

Very simply, this axiom imposes some level of consistency on choices. For example, suppose there were three potential travel locations, Tokyo (T), Beijing (B), and Seoul (S). If a decision maker, when offered the choice between Tokyo and

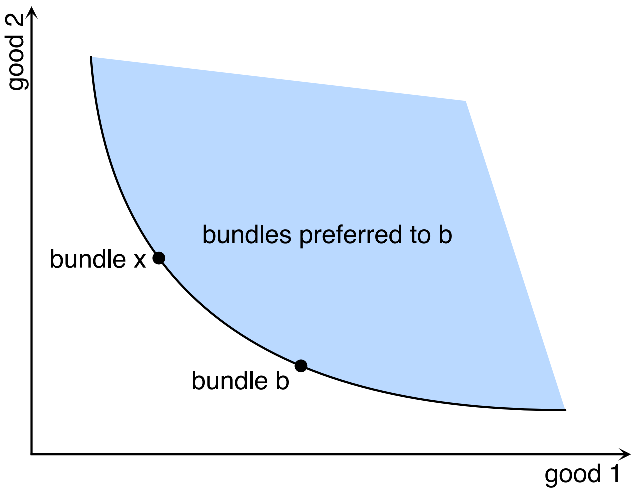

Figure 1: Indifference curve

Beijing, weakly prefers to go to Tokyo, and when given the choice between Beijing and Seoul weakly prefers to go to Beijing, then this axiom simply says that if she was offered a choice between a trip to Tokyo or a trip to Seoul, she would weakly prefer to go to Tokyo. This is because she has already demonstrated that she weakly prefers Tokyo to Beijing, and Beijing to Seoul, so weakly preferring Seoul to Tokyo would mean that their preferences are inconsistent.

But it is conceivable that people might violate transitivity in certain circumstances. One of them is "framing effect". It is the idea that the way the choice alternatives are framed may affect decision and hence in turn may violate transitivity eventually. The idea was made explicit by an experiment due to Danny Kahneman and Amos Tversky (1984). In the experiment students visiting the MIT-Coop to purchase a stereo for \$125 and a calculator for \$5 were informed that the calculator is on sale for 5 dollars less at Harvard Coop. The question is would the students make the trip?
Suppose instead the students were informed that the stereo is 5 dollars less at Harvard Coop.
Kahneman and Tversky found that the fraction of respondents who would travel for cheaper calculator is much higher than for cheaper stereo. But they were also told that there is a stockout and the students have to go to Harvard Coop, and will get 5 dollars off either item as compensation, and were asked which item do you care to get money off? Most of them said that they were indifferent. If $x =$ go to Harvard and get 5 dollars off calculator, $y=$ go to Harvard and get 5 dollars off stereo, $z =$ get both items at MIT. We have $x \succ z$ and $z \succ y$, but last question implies $x \sim y$. Transitivity would imply that $x \succ y$, which is the contradiction.

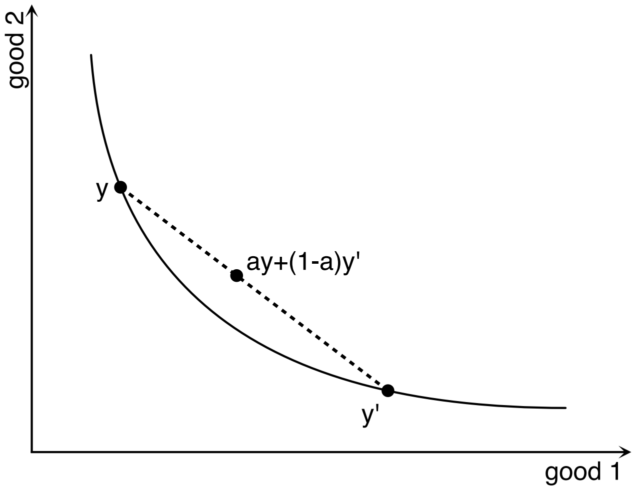

Figure 2: Convex preferences

We for the purposes of this course would assume away any such framing effects in the mind of the decision maker.

Axiom 3. *Reflexivity. For all $x \in X$, $x \succsim x$ (equivalently, $x \sim x$).*

The final axiom is made for technical reasons, and simply says that a bundle cannot be strictly preferred to itself. Such preferences would not make sense.

These three axioms allow for bundles to be ordered in terms of preference. In fact, these three conditions are sufficient to allow preferences to be represented by a utility function.

Before elaborating on this, we consider an example. Suppose there are two goods, Wine and Cheese. Suppose there are four consumption bundles $z = (2,2), y = (1,1), a = (2,1), b = (1,2)$ where the two elements of the vector represent the amount of wine or cheese. Most likely, $z \succ y$ since it provides more of everything (i.e., wine and cheese are "goods"). It is not clear how to compare a and b. What we can do is consider which bundles are indifferent with b. This is an indifference curve (see Figure 1). We can define it as

$$I_b = \{x \in X | b \sim x)$$

We can then (if we assume that more is better) compare a and b by considering which side of the indifference curve a lies on: bundles above and to the right are more preferred, bundles below and to the left are less preferred. This reduces the dimensionality of the problem. We can speak of the "better than b" set as the set of points weakly preferred to b. These preferences are "ordinal:" we can ask whether x is in the better than set, but this does not tell us how much x is

Figure 3: Perfect substitutes (left) and perfect complements (right)

preferred to b. It is common to assume that preferences are monotone: more of a good is better.

Definition 2. The preferences $\succsim$ are said to be (strictly) monotone if $x \geq y \Rightarrow x \succsim y$ ($x \geq y, x \neq y \Rightarrow x \succsim y$ for strict monotonicity).[1]

Suppose I want to increase my consumption of good 1 without changing my level of well-being. The amount I must change x_2 to keep utility constant, $\frac{dx_2}{dx_1}$ is the marginal rate of substitution. Most of the time we believe that individuals like moderation. This desire for moderation is reflected in convex preferences. A mixture between two bundles, between which the agent is indifferent, is strictly preferred to either of the initial bundle (see Figure 2).

Definition 3. A preference relation is convex if for all y and y' with $y \sim y'$ and all $\alpha \in [0,1]$ we have that $\alpha y + (1-\alpha)y' \succsim y \sim y'$.

While convex preferences are usually assumed, there could be instances where preferences are not convex. For example, there could be returns to scale for some good.

Examples: perfect substitutes, perfect complements (see Figure 3). Both of these preferences are convex.

Notice that indifference curves cannot intersect. If they did we could take two points x and y, both to the right of the indifference curve the other lies on. We would then have $x \succ y \succ x$, but then by transitivity $x \succ x$ which contradicts reflexivity. So every bundle is associated with one, and only one, welfare level.

Another important property of preference relation is continuity.

Definition 4. Let $\{x_n\}, \{y_n\}$ be two sequences of choices. If $x_n \succsim y_n$, $\forall n$ and $x_n \to x$, and $y_n \to y$, then $x \succsim y$.

[1] If $x = (x_1, \ldots, x_N)$ and $y = (y_1, \ldots, y_N)$ are vectors of the same dimension, then $x \geq y$ if and only if, for all i, $x_i \geq y_i$. $x \neq y$ means that $x_i \neq y_i$ for at least one i.

This property guarantees that there is no jump in preferences. When X is no longer finite, we need continuity to ensure a utility representation.

2.1 Utility Functions

What we want to consider now is whether we can take preferences and map them to some sort of utility index. If we can somehow represent preferences by such a function we can apply mathematical techniques to make the consumer's problem more tractable. Working with preferences directly requires comparing each of a possibly infinite number of choices to determine which one is most preferred. Maximizing an associated utility function is often just a simple application of calculus. If we take a consumption bundle $x \in \mathbb{R}^N_+$ we can take a utility function as a mapping from $\mathbb{R}^N_+$ into $\mathbb{R}$.

Definition 5. A utility function (index) $u : X \to \mathbb{R}$ represents a preference profile $\succsim$ if and only if, for all $x, y \in X$: $x \succsim y \Leftrightarrow u(x) \geq u(y)$.

We can think about a utility function as an "as if"-concept: the agent acts "as if" she has a utility function in mind when making decisions.

Is it always possible to find such a function? The following result shows that such a function exists under the three assumptions about preferences we made above.

Proposition 1. *Suppose that X is finite. Then the assumptions of completeness, transitivity, and reflexivity imply that there is a utility function u such that $u(x) \geq u(y)$ if and only if $x \succsim y$.*

Proof. We define an explicit utility function. Let's introduce some notation:

$$B(x) = \{z \in X | x \succsim z\}$$

Therefore $B(x)$ is the set of "all items below x". Let the utility function be defined as,

$$u(x) = |B(x)|$$

where $|B(x)|$ is the cardinality of the set $B(x)$, i.e. the number of elements in the set $B(x)$. There are two steps to the argument:
First part:

$$u(x) \geq u(y) \quad \Rightarrow \quad x \succsim y$$

Second part:

$$x \succsim y \quad \Rightarrow \quad u(x) \geq u(y)$$

First part of proof:

By definition, $u(x) \geq u(y) \quad \Rightarrow \quad |B(x)| \geq |B(y)|$. If $y \in B(x)$, then $x \succsim y$ by definition of $B(x)$ and we are done. Otherwise, $y \notin B(x)$. We will work towards a contradiction.

Since $y \notin B(x)$, we have

$$|B(x) - \{y\}| = |B(x)|$$

Since $y \in B(y)$ (by reflexivity), we have

$$|B(y)| - 1 = |B(y) - \{y\}|$$

Since $|B(x)| \geq |B(y)|$, $|B(x)| > |B(y)| - 1$ and hence,

$$|B(x) - \{y\}| > |B(y) - \{y\}|$$

Therefore, there must be some $z \in X - \{y\}$ such that $x \succsim z$ and $y \not\succsim z$. By completeness: $z \succsim y$. By transitivity: $x \succsim y$. But this implies that $y \in B(x)$, a contradiction. Second part of proof

Want to show: $x \succsim y \quad \Rightarrow \quad u(x) \geq u(y)$.
Suppose $x \succsim y$ and $z \in B(y)$.
Then $x \succsim y$ and $y \succsim z$, so by transitivity $x \succsim z$.
Hence, $z \in B(x)$.
This shows that when $x \succsim y$, anything in $B(y)$ must also be in $B(x)$.

$$B(y) \subset B(x) \quad \Rightarrow \quad |B(x)| \geq |B(y)| \quad \Rightarrow \quad u(x) \geq u(y)$$

This completes the proof. □

In general the following proposition holds:

Proposition 2. *Every (continuous) preference ranking can be represented by a (continuous) utility function.*

This result can be extended to environments with uncertainty, as was shown by Leonard Savage. Consequently, we can say that individuals behave as if they are maximizing utility functions, which allows for marginal and calculus arguments. There is, however, one qualification. The utility function that represents the preferences is not unique.

Remark 1. If u represents preferences, then for any increasing function $f : \mathbb{R} \to \mathbb{R}$, $f(u(x))$ also represents the same preference ranking

In the previous section, we claimed that preferences usually reflect the idea that "more is better," or that preferences are monotone.

Definition 6. The utility function (preferences) are monotone increasing if $x \geq y$ implies that $u(x) \geq u(y)$ and $x > y$ implies that $u(x) > u(y)$.

One feature that monotone preferences rule out is (local) satiation, where one point is preferred to all other points nearby. For economics the relevant decision is maximizing utility subject to limited resources. This leads us to consider constrained optimization.

3 Maximization

Now we take a look at the mathematical tool that will be used with the greatest intensity in this course. Let $x = (x_1, x_2, \ldots, x_n)$ be a n-dimensional vector where each component of the vector $x_i, i = 1, 2, \ldots, n$ is a non-negative real number. In mathematical notations we write $x \in \mathbb{R}^n_+$. We can think of x as description of different characteristics of a choice that the decision maker faces. For example, while choosing which college to go (among the ones that have offered admission) a decision maker, who is a student in this case, looks into different aspects of a university, namely the quality of instruction, diversity of courses, location of the campus etc. The components of the vector x can be thought of as each of these characteristics when the choice problem faced by the decision maker (i.e. the student) is to choose which university to attend. Usually when people go to groceries they are faced with the problem of buying not just! one commodity, but a bundle of commodities and therefore it is the combination of quantities of different commodities which needs to be decided and again the components of x can be thought of as quantities of each commodity purchased. Whatever be the specific context, utility is defined over the set of such bundles. Since $x \in \mathbb{R}^n_+$, we take $X = \mathbb{R}^n_+$. So the utility function is a mapping $u \colon \mathbb{R}^n_+ \to \mathbb{R}$.

Now for the time being let x be one dimensional, i.e. $x \in R$. Let $f \colon \mathbb{R} \to \mathbb{R}$ be a continuous and differentiable function that takes real numbers and maps it to another real number. Continuity is assumed to avoid any jump in the function and differentiability is assumed to avoid kinks. The slope of the function f is defined as the first derivative of the function and the curvature of the function is defined as the second derivative of the function. So, the slope of f at x is formally defined as:

$$\frac{df(x)}{dx} \triangleq f^{'}(x)$$

and the curvature of f at x is formally defined as:

$$\frac{d^2 f(x)}{dx^2} \triangleq f^{''}(x)$$

In order to find out the maximum of f we must first look into the slope of f. If the slope is positive then raising the value of x increases the value of f. So to find out the maximum we must keep increasing x. Similarly if slope is negative then reducing the value of x increases the value of f and therefore to find the maximum we should reduce the value of x. Therefore the maximum is reached when the slope is exactly equal to 0. This condition is referred to as the First Order Condition (F.O.C.) or the necessary condition:

$$\frac{df(x)}{dx} = 0$$

But this in itself doesn't guarantee that maximum is reached, as a perfectly flat slope may also imply that we're at the trough, i.e. at the minimum. The F.O.C. therefore finds the extremum points in the function. We need to look at the curvature to make sure whether the extremum is actually a maximum or not. If the second derivative is negative then it means that from the extremum point if we move x a little bit on either side $f(x)$ would fall, and therefore the extremum is a maximum. But if the second derivative is positive then by similar argument we know that its the minimum. This condition is referred to as the Second Order Condition (S.O.C) or the sufficient condition:

$$\frac{d^2f(x)}{dx^2} \leq 0$$

Now we look at the definitions of two important kind of functions:

Definition 7. (i)A continuous and differentiable function $f\colon \mathbb{R} \to \mathbb{R}$ is (strictly) concave if

$$\frac{d^2f(x)}{dx^2} \leq (<)0.$$

(ii) f is convex if

$$\frac{d^2f(x)}{dx^2} \geq (>)0.$$

Therefore a concave function the F.O.C. is both necessary and sufficient condition for maximization. We can also define concavity or convexity of functions with the help of convex combinations.

Definition 8. A convex combination of two any two points $x^{'}, x^{''} \in \mathbb{R}^n$ is defined as $x_\lambda = \lambda x^{'} + (1-\lambda)x^{''}$ for any $\lambda \in (0,1)$.

Convex combination of two points represent a point on the straight line joining those two points. We now define concavity and convexity of functions using this concept.

Definition 9. f is concave if for any two points $x', x'' \in \mathbb{R}$, $f(x_\lambda) \geq \lambda f(x') + (1-\lambda)f(x'')$ where x_λ is a convex combination of x' and x'' for $\lambda \in (0,1)$. f is strictly concave if the inequality is strict.

Definition 10. Similarly f is convex if $f(x_\lambda) \leq \lambda f(x') + (1-\lambda)f(x'')$. f is strictly convex if the inequality is strict.

If the utility function is concave for any individual then, given this definition, we can understand that, she would prefer to have a certain consumption of x_λ than face an uncertain prospect of consuming either x' or x''. Such individuals are called risk averse. We shall explore these concepts in full detail later in the course and then we would require these definitions of concavity and convexity.

4 Utility Maximization

4.1 Multivariate Function Maximization

Let $x = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n_+$ be a consumption bundle and $f\colon \mathbb{R}^n_+ \to \mathbb{R}$ be a multivariate function. The multivariate function that we are interested in here is the utility function $u\colon \mathbb{R}^n_+ \to \mathbb{R}$ where $u(x)$ is the utility of the consumption bundle x.
The F.O.C. for maximization of f is given by:

$$\frac{\partial f(x_1, x_2, \ldots, x_n)}{\partial x_i} = 0 \quad \forall\, i = 1, 2, \ldots, n$$

This is a direct extension of the F.O.C. for univariate functions as explained in Lecture 3. The S.O.C. however is a little different from the single variable case. Let's look at a bivariate function $f\colon \mathbb{R}^2 \to \mathbb{R}$. Let's first define the following notations:

$$f_i(x) \triangleq \frac{df(x)}{dx_i},\ f_{ii}(x) \triangleq \frac{d^2 f(x)}{dx_2^2},\ i = 1, 2,$$
$$f_{ij}(x) \triangleq \frac{d^2 f(x)}{dx_i dx_j},\ i \neq j$$

The S.O.C. for the maximization of f is then given by,

$$(i) \qquad f_{11} < 0$$
$$(ii) \ \begin{vmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{vmatrix} > 0$$

The first of the S.O.C.s is analogous to the S.O.C. for the univariate case. If we write out the second one we get,

$$f_{11}f_{22} - f_{12}f_{21} > 0$$

But we know that $f_{12} = f_{21}$. So,

$$f_{11}f_{22} > f_{12}^2 > 0$$

$$\Rightarrow f_{22} < 0 \ \ (\text{since } f_{11} < 0)$$

Therefore the S.O.C. for the bivariate case is stronger than the analogous conditions from the univariate case. This is because for the bivariate case to make sure that we are at the peak of a function it is not enough to check if the function is concave in the directions of x_1 and x_2, as it could not be concave along the diagonal and therefore the need to introduce cross derivatives in to the condition. For the purposes of this class we'd assume that the S.O.C. is satisfied for the utility function being given, unless it is asked specifically to check for it.

4.2 Budget Constraint

A budget constraint is a constraint on how much money (income, wealth) an agent can spend on goods. We denote the amount of available income by $I \geq 0$. $x_1, \ldots, x_N$ are the quantities of the goods purchased and $p_1, \ldots, p_N$ are the according prices. Then the budget constraint is

$$\sum_{i=1}^{N} p_i x_i \leq I.$$

As an example, we consider the case with two goods. In that case we get that $p_1x_1 + p_2x_2 \leq I$, i.e., the agent spends her entire income on the two goods. The points where the budget line intersects with the axes are $x_1 = I/p_1$ and $x_2 = I/p_2$ since these are the points where the agent spends her income on only one good. Solving for x_2, we can express the budget line as a function of x_1:

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

where the slope of the budget line is given by,

$$\frac{dx_2}{dx_1} = -\frac{p_1}{p_2}$$

The **budget line** here is defined as the equation involving x_1 and x_2 such that the decision maker exhausts all her income. The set of consumption bundles (x_1, x_2) which are feasible given the income, i.e. (x_1, x_2) for which $p_1x_1 + p_2x_2 \leq I$ holds is defined as the **budget set**.

4.3 Indifference Curve

Indifference Curve (IC) is defined as the locus of consumption bundles (x_1, x_2) such that the utility is held fixed at some level. Therefore the equation of the IC is given by,

$$u(x_1, x_2) = \bar{u}$$

To get the slope of the IC we differentiate the equation w.r.t. x_1:

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

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

where MU_i refers to the marginal utility of good i. So the slope of IC is the (negative of) ratio of marginal utilities of good 1 and 2. This ratio is referred to as the Marginal Rate of Substitution or MRS. This tells us the rate at which the consumer is ready to substitute between good 1 and 2 to remain at the same utility level.

4.4 Constrained Optimization

Consumers are typically endowed with money I, which determines which consumption bundles are affordable. The budget set consists of all consumption bundles such that $\sum_{i=1}^{N} p_i x_i \leq I$. The consumer's problem is then to find the point on the highest indifference curve that is in the budget set. At this point the indifference curve must be tangent to the budget line. The slope of the budget line is given by,

$$\frac{dx_2}{dx_1} = -\frac{p_1}{p_2}$$

which defines how much x_2 must decrease if the amount of consumption of good 1 is increased by dx_1 for the bundle to still be affordable. It reflects the opportunity cost, as money spent on good 1 cannot be used to purchase good 2 (see Figure 4).

The marginal rate of substitution, on the other hand, reflects the relative benefit from consuming different goods. The slope of the indifference curve is $-MRS$. So the relevant optimality condition, where the slope of the indifference curve equals the slope of the budget line, is

$$\frac{p_1}{p_2} = \frac{\frac{\partial u(x)}{\partial x_1}}{\frac{\partial u(x)}{\partial x_2}}.$$

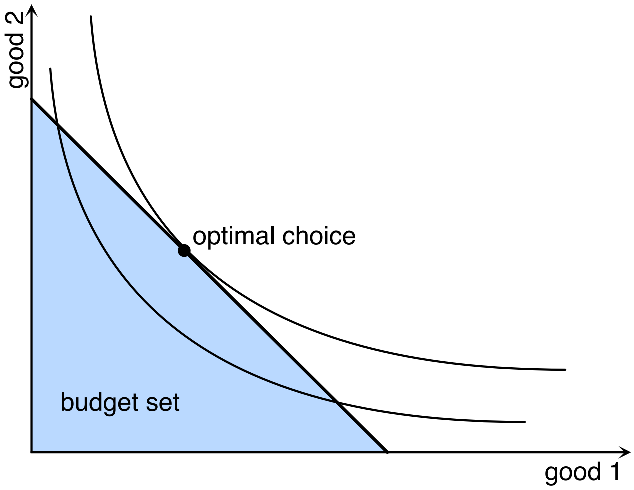

Figure 4: Indifference curve and budget set

We could equivalently talk about equating marginal utility per dollar. If

$$\frac{\frac{\partial u(x)}{\partial x_2}}{p_2} > \frac{\frac{\partial u(x)}{\partial x_1}}{p_1}$$

then one dollar spent on good 2 generates more utility then one dollar spent on good 1. So shifting consumption from good 1 to good 2 would result in higher utility. So, to be at an optimum we must have the marginal utility per dollar equated across goods.

Does this mean then that we must have $\frac{\partial u(x)}{\partial x_i} = p_i$ at the optimum? No. Such a condition wouldn't make sense since we could rescale the utility function. We could instead rescale the equation by a factor $\lambda \geq 0$ that converts "money" into "utility." We could then write $\frac{\partial u(x)}{\partial x_i} = \lambda p_i$. Here, λ reflects the marginal utility of money. More on this in the subsection on Optimization using Lagrange approach.

4.4.1 Optimization by Substitution

The consumer's problem is to maximize utility subject to a budget constraint. There are two ways to approach this problem. The first approach involves writing the last good as a function of the previous goods, and then proceeding with an unconstrained maximization. Consider the two good case. The budget set consists of the constraint that $p_1x_1 + p_2x_2 \leq I$. So the problem is

$$\max_{x_1, x_2} u(x_1, x_2) \text{ subject to } p_1x_1 + p_2x_2 \leq I$$

But notice that whenever u is (locally) non-satiated then the budget constraint holds with equality since there in no reason to hold money that could have been used for additional valued consumption. So, $p_1x_1 + p_2x_2 = I$, and so we can write x_2 as a function of x_1 from the budget equation in the following way

$$x_2 = \frac{I - p_1x_1}{p_2}$$

Now we can treat the maximization of $u\left(x_1, \frac{I-p_1x_1}{p_2}\right)$ as the standard single variable maximization problem. Therefore now the maximization problem becomes,

$$\max_{x_1} u\left(x_1, \frac{I - p_1x_1}{p_2}\right)$$

The F.O.C. is then given by,

$$\frac{du}{dx_1} + \frac{du}{dx_2}\frac{dx_2(x_1)}{dx_1} = 0$$
$$\Rightarrow \frac{du}{dx_1} - \frac{p_1}{p_2}\frac{du}{dx_2} = 0$$

The second equation substitutes $\frac{dx_2(x_1)}{dx_1}$ by $-\frac{p_1}{p_2}$ from the budget line equation. We can further rearrange terms to get,

$$\frac{\frac{du}{dx_1}}{p_1} = \frac{\frac{du}{dx_2}}{p_2}$$
$$\Rightarrow \frac{\frac{du}{dx_1}}{\frac{du}{dx_2}} = \frac{p_1}{p_2}$$

This exactly the condition we got by arguing in terms of budget line and indifference curves. In the following lecture we shall look at a specific example where we would maximize a particular utility function using this substitution method and then move over to the Lagrange approach.

5 Utility Maximization Continued

5.1 Application of Substitution Method

Example 1. We consider a consumer with Cobb-Douglas preferences. Cobb-Douglas preferences are easy to use and therefore commonly used. The utility function is defined as (with two goods)

$$u(x_1, x_2) = x_1^{\alpha}x_2^{1-\alpha},\ \alpha > 0$$

The goods' prices are p_1,p_2 and the consumer if endowed with income I. Hence, the constraint optimization problem is

$$\max_{x_1,x_2} \; x_1^{\alpha} x_2^{1-\alpha} \;\; \text{subject to} \;\; p_1x_1 + p_2x_2 = I.$$

We solve this maximization by substituting the budget constraint into the utility function so that the problem becomes an unconstrained optimization with one choice variable:

$$u(x_1) = x_1^{\alpha} \left(\frac{I - p_1x_1}{p_2} \right)^{1-\alpha}. \tag{1}$$

In general, we take the total derivative of the utility function

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

which gives us the condition for optimal demand

$$\frac{dx_2}{dx_1} = -\frac{\frac{\partial u}{\partial x_1}}{\frac{\partial u}{\partial x_2}}.$$

The right-hand side is the marginal rate of substitution (MRS).

In order to calculate the demand for both goods, we go back to our example. Taking the derivative of the utility function (1)

$$\begin{aligned} u'(x_1) &= \alpha x_1^{\alpha-1} \left(\frac{I - p_1x_1}{p_2} \right)^{1-\alpha} + (1-\alpha)x_1^{\alpha} \left(\frac{I - p_1x_1}{p_2} \right)^{-\alpha} \left(-\frac{p_1}{p_2} \right) \\ &= x_1^{\alpha-1} \left(\frac{I - p_1x_1}{p_2} \right)^{-\alpha} \left[\alpha \frac{I - p_1x_1}{p_2} - (1-\alpha)x_1 \frac{p_1}{p_2} \right] \end{aligned}$$

so the FOC is satisfied when

$$\alpha(I - p_1x_1) - (1-\alpha)x_1p_1 = 0$$

which holds when

$$x_1^*(p_1, p_2, I) = \frac{\alpha I}{p_1}. \tag{2}$$

Hence, we see that the budget spent on good 1, p_1x_1, equals the budget share αI, where α is the preference parameter associated with good 1.

Plugging (2) into the budget constraint yields

$$x_2^*(p_1, p_2, I) = \frac{I - p_1x_1}{p_2} = \frac{(1-\alpha)I}{p_2}.$$

These are referred to as the Marshallian demand or uncompensated demand.

Several important features of this example are worth noting. First of all, x_1 does not depend on p_2 and vice versa. Also, the share of income spent on each good $\frac{p_i x_i}{M}$ does not depend on price or wealth. What is going on here? When the price of one good, p_2, increases there are two effects. First, the price increase makes good 1 relatively cheaper ($\frac{p_1}{p_2}$ decreases). This will cause consumers to "substitute" toward the relatively cheaper good. There is also another effect. When the price increases the individual becomes poorer in real terms, as the set of affordable consumption bundles becomes strictly smaller. The Cobb-Douglas utility function is a special case where this "income effect" exactly cancels out the substitution effect, so the consumption of one good is independent of the price of the other goods.

Cobb - Douglass utility function $u(x_1, x_2) = {x_1}^{\alpha} {x_2}^{(1-\alpha)}$ sub. to budget constraint $p_1 x_1 + p_2 x_2 = I$

Therefore we get,

$$\max_{x_1} {x_1}^{\alpha} \left(\frac{I - p_1 x_1}{p_2} \right)^{(1-\alpha)}$$

The F.O.C. is then given by,

$$\alpha {x_1}^{\alpha - 1} \left(\frac{I - p_1 x_1}{p_2} \right)^{(1-\alpha)} + (1-\alpha) {x_1}^{\alpha} \left(\frac{I - p_1 x_1}{p_2} \right)^{\alpha} \left(-\frac{p_1}{p_2} \right) = 0$$

$$\Rightarrow \alpha \left(\frac{I - p_1 x_1}{p_2} \right) = (1-\alpha) x_1 \left(\frac{p_1}{p_2} \right)$$

$$\Rightarrow {x_1}^{*}(p_1, p_2, I) = \frac{\alpha I}{p_1}$$

$$\Rightarrow {x_2}^{*}(p_1, p_2, I) = \frac{(1-\alpha) I}{p_2}$$

This is referred to as the Marshallian Demand or uncompensated demand.

5.2 Elasticity

When calculating price or income effects, the result depends on the units used. For example, when considering the own-price effect for gasoline, we might express quantity demanded in gallons or liters and the price in dollars or euros. The own-price effects would differ even if consumers in the U.S. and Europe had the same underlying preferences. In order to make price or income effects comparable across different units, we need to normalize them. This is the reason why we use the concept of elasticity. The *own-price elasticity of demand* is defined as the percentage change in demand for each percentage change in its own price and is denoted by ϵ_i:

$$\epsilon_i = -\frac{\frac{\partial x_i}{\partial p_i}}{\frac{x_i}{p_i}} = -\frac{\partial x_i}{\partial p_i}\frac{p_i}{x_i}.$$

It is common to multiply the price effect by -1 so that ϵ is a positive number since the price effect is usually negative. Of course, the *cross-price elasticity of demand* is defined similarly

$$\epsilon_{ij} = -\frac{\frac{\partial x_i}{\partial p_j}}{\frac{x_i}{p_j}} = -\frac{\partial x_i}{\partial p_j}\frac{p_j}{x_i}.$$

Similarly the *income elasticity of demand* is defined as:

$$\epsilon_I = \frac{\frac{\partial x_i}{\partial I}}{\frac{x_i}{I}} = \frac{\partial x_i}{\partial I}\frac{I}{x_i}$$

5.2.1 Constant Elasticity of Substitution

Elasticity of substitution for a utility function is defined as the elasticity of the ratio of consumption of two goods to the MRS. Therefore it is a measure of how easily the two goods are substitutable along an indifference curve. In terms of mathematics, it is defined as,

$$\epsilon_S = \frac{d(x_2/x_1)}{dMRS}\frac{MRS}{x_2/x_1}$$

For a class of utility functions this value is constant for all (x_1, x_2). These utility functions are called Constant Elasticity of Substitution (CES) utility functions. The general form looks like the following:

$$u(x_1, x_2) = \left(\alpha_1 x_1{}^{-\rho} + \alpha_2 x_2{}^{-\rho}\right)^{-\frac{1}{\rho}}$$

It is easy to show that for CES utility functions,

$$\epsilon_S = \frac{1}{\rho + 1}$$

The following utility functions are special cases of the general CES utility function:
Linear Utility: Linear Utility is of the form

$$U(x_1, x_2) = ax_1 + bx_2, \quad a, b \text{ constants}$$

which is a CES utility with $\rho = -1$.
Leontief Utility: Leontief utility is of the form

$$U(x_1, x_2) = \max\{\frac{x_1}{a}, \frac{x_2}{b}\}, \quad a, b > 0$$

and this is also a CES utility function with $\rho = \infty$.

5.3 Optimization Using the Lagrange Approach

While the approach using substitution is simple enough, there are situations where it will be difficult to apply. The procedure requires that, as we know, before the calculation, the budget constraint actually binds. In many situations there may be other constraints (such as a non-negativity constraint on the consumption of each good) and we may not know whether they bind before demands are calculated. Consequently, we will consider a more general approach of Lagrange multipliers. Again, we consider the (two good) problem of

$$\max_{x_1,x_2} u(x_1,x_2) \text{ s.t. } p_1x_1 + p_2x_2 \leq I$$

Let's think about this problem as a game. The first player, let's call him the kid, wants to maximize his utility, $u(x_1, x_2)$, whereas the other player (the parent) is concerned that the kid violates the budget constraint, $p_1x_1 + p_2x_2 \leq I$, by spending too much on goods 1 and 2. In order to induce the kid to stay within the budget constraint, the parent can punish him by an amount λ for every dollar the kid exceeds his income. Hence, the total punishment is

$$\lambda(I - p_1x_1 - p_2x_2).$$

Adding the kid's utility from consumption and the punishment, we get

$$\mathcal{L}(x_1, x_2, \lambda) = u(x_1, x_2) + \lambda(I - p_1x_1 - p_2x_2). \tag{3}$$

Since, for any function, we have max $f = -$ min f, this game is a zero-sum game: the payoff for the kid is $\mathcal{L}$ and the parent's payoff is $-\mathcal{L}$ so that the total payoff will always be 0. Now, the kid maximizes expression (3) by choosing optimal levels of x_1 and x_2, whereas the parent minimizes (3) by choosing an optimal level of λ:

$$\min_{\lambda} \max_{x_1,x_2} \mathcal{L}(x_1, x_2, \lambda) = u(x_1, x_2) + \lambda(I - p_1x_1 - p_2x_2).$$

In equilibrium, the optimally chosen level of consumption, x^*, has to be the best response to the optimal level of λ^* and vice versa. In other words, when we fix a level of x^*, the parent chooses an optimal $\lambda^*$and when we fix a level of λ^*, the kid chooses an optimal x^*. In equilibrium, no one wants to deviate from their optimal choice. Could it be an equilibrium for the parent to choose a very large λ? No, because then the kid would not spend any money on consumption, but rather have the maximized expression (3) to equal λI.

Since the first-order conditions for minima and maxima are the same, we have the following first-order conditions for problem (3):

$$\frac{\partial \mathcal{L}}{\partial x_1} = \frac{\partial u}{\partial x_1} - \lambda p_1 = 0 \tag{4}$$

$$\frac{\partial \mathcal{L}}{\partial x_2} = \frac{\partial u}{\partial x_2} - \lambda p_2 = 0 \tag{5}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = I - p_1 x_1 - p_2 x_2 = 0.$$

Here, we have three equations in three unknowns that we can solve for the optimal choice x^*, λ^*.

Before solving this problem for an example, we can think about it in more formal terms. The basic idea is as follows: Just as a necessary condition for a maximum in a one variable maximization problem is that the derivative equals 0 ($f'(x) = 0$), a necessary condition for a maximum in multiple variables is that all partial derivatives are equal to 0 ($\frac{\partial f(x)}{\partial x_i} = 0$). To see why, recall that the partial derivative reflects the change as x_i increases and the other variables are all held constant. If any partial derivative was positive, then holding all other variables constant while increasing x_i will increase the objective function (similarly, if the partial derivative is negative we could decrease x_i). We also need to ensure that the solution is in the budget set, which typically won't happen if we just try to maximize u. Basically, we impose a "cost" on consumption (the punishment in the game above), proceed with unconstrained maximization for the induced problem, and set this cost so that the maximum lies in the budget set.

Notice that the first-order conditions (4) and (5) imply that

$$\frac{\frac{\partial u}{\partial x_1}}{p_1} = \lambda = \frac{\frac{\partial u}{\partial x_2}}{p_2}$$

or

$$\frac{\frac{\partial u}{\partial x_1}}{\frac{\partial u}{\partial x_2}} = \frac{p_1}{p_2}$$

which is precisely the "MRS = price ratio" condition for optimality that we saw before.

Finally, it should be noted that the FOCs are necessary for optimality, but they are not, in general, sufficient for the solution to be a maximum. However, whenever $u(x)$ is a concave function the FOCs are also sufficient to ensure that the solution is a maximum. In most situations, the utility function will be concave.

Example 2. We can consider the problem of deriving demands for a Cobb-Douglas utility function using the Lagrange approach. The associated Lagrangian is

$$\mathcal{L}(x_1, x_2, \lambda) = x_1^{\alpha} x_2^{1-\alpha} + \lambda(I - p_1 x_1 - p_2 x_2),$$

which yields the associated FOCs

$$\frac{\partial \mathcal{L}}{\partial x_1} = \alpha x_1^{\alpha-1} x_2^{1-\alpha} - \lambda p_1 = \alpha \left(\frac{x_2}{x_1}\right)^{1-\alpha} - \lambda p_1 = 0 \tag{6}$$

$$\frac{\partial \mathcal{L}}{\partial x_2} = (1-\alpha) x_1^{\alpha} x_2^{-\alpha} - \lambda p_2 = (1-\alpha) \left(\frac{x_1}{x_2}\right)^{\alpha} - \lambda p_2 = 0 \tag{7}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = (I - p_1 x_1 - p_2 x_2) = 0. \tag{8}$$

We have three equations with three unknowns (x_1, x_2, λ) so that this system should be solvable. Notice that since it is not possible that $\frac{x_2}{x_1}$ and $\frac{x_1}{x_2}$ are both 0 we cannot have a solution to equations (6) and (7) with $\lambda = 0$. Consequently we must have that $p_1 x_1 + p_2 x_2 = I$ in order to satisfy equation (8). Solving for λ in the above equations tells us that

$$\lambda = \frac{\alpha}{p_1} \left(\frac{x_2}{x_1}\right)^{1-\alpha} = \frac{(1-\alpha)}{p_2} \left(\frac{x_1}{x_2}\right)^{\alpha}$$

and so

$$p_2 x_2 = \frac{1-\alpha}{\alpha} p_1 x_1.$$

Combining with the budget constraint this gives

$$p_1 x_1 + \frac{1-\alpha}{\alpha} p_1 x_1 = \frac{1}{\alpha} p_1 x_1 = I.$$

So the Marshallian[2] demand functions are

$$x_1^* = \frac{\alpha I}{p_1}$$

and

$$x_2^* = \frac{(1-\alpha) I}{p_2}.$$

So we see that the result of the Lagrangian approach is the same as from approach that uses substitution. Using equation (6) or (7) again along with the optimal demand x_1^* or x_2^* gives us the following expression for λ:

$$\lambda^* = \frac{1}{I}.$$

Hence, λ^* equals the derivative of the Lagrangian $\mathcal{L}$ with respect to income I. We call this derivative, $\frac{\partial \mathcal{L}}{\partial I}$, the marginal utility of money.

[2] After the British economist Alfred Marshall.

6 Value Function and Comparative Statics

6.1 Indirect Utility Function

The indirect utility function

$$V(p_1, p_2, I) \triangleq u\left(x_1^*(p_1, p_2, I), x_2^*(p_1, p_2, I)\right)$$

Therefore V is the maximum utility that can be achieved given the prices and the income level. We shall show later that λ is same as

$$\frac{\partial V(p_1, p_2, I)}{\partial I}$$

6.2 Interpretation of λ

From FOC of maximization we get,

$$\begin{aligned}
\frac{dL}{dx_1} &= \frac{\partial u}{dx_1} - \lambda p_1 = 0 \\
\frac{dL}{dx_2} &= \frac{\partial u}{dx_2} - \lambda p_2 = 0 \\
\frac{dL}{d\lambda} &= I - p_1 x_1 - p_2 x_2 = 0
\end{aligned}$$

From the first two equations we get,

$$\lambda = \frac{\frac{\partial u}{dx_i}}{p_i}$$

This means that λ can be interpreted as the par dollar marginal utility from any good. It also implies, as we have argued before, that the benefit to cost ratio is equalized across goods. We can also interpret λ as shadow value of money. But we explain this concept later. Before that let's solve an example and find out the value of λ for that problem.
Let's work with the utility function:

$$u(x_1, x_2) = \alpha \ln x_1 + (1 - \alpha) \ln x_2$$

The F.O.C.s are then given by,

$$\frac{\partial L}{\partial x_1} = \frac{\alpha}{x_1} - \lambda p_1 = 0 \tag{9}$$

$$\frac{\partial L}{\partial x_1} = \frac{1 - \alpha}{x_2} - \lambda p_2 = 0 \tag{10}$$

$$\frac{\partial L}{\partial \lambda} = I - p_1 x_1 - p_2 x_2 = 0 \tag{11}$$

From the first two equations (3) and (4) we get,

$$x_1 = \frac{\alpha}{\lambda p_1} \quad \text{and} \quad x_2 = \frac{1-\alpha}{\lambda p_2}$$

Plugging it in the F.O.C. equation (5) we get,

$$I = \frac{\alpha}{\lambda} + \frac{1-\alpha}{\lambda}$$

$$\Rightarrow \lambda^* = \frac{1}{I}$$

6.3 Comparative Statics

Let $f\colon \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ be a function which is dependent on an endogenous variable, say x, and an exogenous variable a. Therefore we have,

$$f(x, a)$$

Let's define value function as the maximized value of f w.r.t. x, i.e.

$$v(a) \triangleq \max_x f(x, a)$$

Let $x^*(a)$ be the value of x that maximizes f given the value of a. Therefore,

$$v(a) = f(x^*(a), a)$$

To find out the effect of changing the value of the exogenous variable a on the maximized value of f we differentiate v w.r.t. a. Hence we get,

$$v'(a) = \frac{df}{dx}\frac{dx^*}{da} + \frac{df}{da}$$

But from the F.O.C. of maximization of f we know that,

$$\frac{df}{dx}(x^*(a), a) = 0$$

Therefore we get that,

$$v'(a) = \frac{df}{da}(x^*(a), a)$$

Thus the effect of change in the exogenous variable on the value function is only it's direct effect on the objective function. This is referred to as the Envelope Theorem.

In case of utility maximization the value function is the indirect utility function. We can also define the indirect utility function as,

$$V(p_1, p_2, I) \triangleq u(x_1^*, x_2^*, \lambda^*) + \lambda^*[I - p_1 x_1^* - p_2 x_2^*] = L(x_1^*, x_2^*, \lambda^*)$$

Therefore,

$$\begin{aligned}\frac{\partial V}{\partial I} = & \frac{\partial u}{\partial x_1}\frac{\partial x_1^*}{\partial I} + \frac{\partial u}{\partial x_2}\frac{\partial x_2^*}{\partial I} \\ & - \lambda^* p_1 \frac{\partial x_1^*}{\partial I} - \lambda^* p_2 \frac{\partial x_2^*}{\partial I} \\ & + [I - p_1 x_1^* - p_2 x_2^*]\frac{\partial \lambda^*}{\partial I} + \lambda^* \\ & = \lambda^* \quad \text{(by Envelope Theorem)}\end{aligned}$$

Therefore we see that λ^* is the marginal value of money in the optimum. So if the income constraint is relaxed by a dollar, it increases the maximum utility of the consumer by λ^* and hence λ^* is interpreted as the shadow value of money.

7 Expenditure Minimization

Instead of maximizing utility subject to a given income we can also minimize expenditure subject to achieving a given level of utility $\bar{u}$. In this case, the consumer wants to spend as little money as possible to enjoy a certain utility. Formally, we write

$$\min_x \; p_1 x_1 + p_2 x_2 \quad \text{s.t. } u(x) \geq \bar{u}. \tag{12}$$

We can set up the Lagrange expression for this problem as the following:

$$\mathcal{L}(x_1, x_2, \lambda) = p_1 x_1 + p_2 x_2 + \lambda[\bar{u} - u(x_1, x_2)]$$

The F.O.C.s are now:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x_1} &= p_1 - \lambda \frac{\partial u}{\partial x_1} = 0 \\ \frac{\partial \mathcal{L}}{\partial x_2} &= p_2 - \lambda \frac{\partial u}{\partial x_2} = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= \bar{u} - u(x_1, x_2) = 0\end{aligned}$$

Comparing the first two equations we get,

$$\frac{\frac{\partial u}{\partial x_1}}{\frac{\partial u}{\partial x_2}} = \frac{p_1}{p_2}$$

This is the exact relation we got in the utility maximization program. Therefore these two programs are equivalent exercises. In the language of mathematics it is called the duality. But the values of x_1 and x_2 that minimizes the expenditure is a function of the utility level $\bar{u}$ instead of income as in the case of utility maximization. The result of this optimization problem is a demand function again, but in general it is different from $x^*(p_1, p_2, I)$. We call the demand function derived from problem (1) *compensated demand* or Hicksian demand.[3] We denote it by,

$$h_1(p_1, p_2, \bar{u}) \text{ and } h_2(p_1, p_2, \bar{u})$$

Note that compensated demand is a function of prices and the utility level whereas uncompensated demand is a function of prices and income. Plugging compensated demand into the objective function $(p_1x_1 + p_2x_2)$ yields the *expenditure function* as function of prices and $\bar{u}$

$$E(p_1, p_2, \bar{u}) = p_1h_1(p_1, p_2, \bar{u}) + p_2h_2(p_1, p_2, \bar{u}).$$

Hence, the expenditure function measures the minimal amount of money required to buy a bundle that yields a utility of $\bar{u}$.

Uncompensated and compensated demand functions usually differ from each other, which is immediately clear from the fact that they have different arguments. There is a special case where they are identical. First, note that indirect utility and expenditure function are related by the following relationships

$$\begin{aligned} V(p_1, p_2, E(p_1, p_2, \bar{u})) &= \bar{u} \\ E(p_1, p_2, V(p_1, p_2, I)) &= I. \end{aligned}$$

That is, if income is exactly equal to the expenditure necessary to achieve utility level $\bar{u}$, then the resulting indirect utility is equal to $\bar{u}$. Similarly, if the required utility level is set equal to the indirect function when income is I, then minimized expenditure will be equal to I. Using these relationships, we have that uncompensated and compensated demand are equal in the following two cases:

$$\begin{aligned} x_i^*(p_1, p_2, I) &= h_i^*(p_1, p_2, V(p_1, p_2, I)) \\ x_i^*(p_1, p_2, E(p_1, p_2, \bar{u})) &= h_i^*(p_1, p_2, \bar{u}) \text{ for } i = 1, 2. \end{aligned} \tag{13}$$

Now we can express income and substitution effects analytically. Start with one component of equation (13):

$$h_i^*(p_1, p_2, \bar{u}) = x_i^*(p_1, p_2, E(p_1, p_2, \bar{u}))$$

[3] After the British economist Sir John Hicks, co-recipient of the 1972 Nobel Prize in Economic Sciences.

and take the derivative with respect to p_j using the chain rule

$$\frac{\partial h_i^*}{\partial p_j} = \frac{\partial x_i^*}{\partial p_j} + \frac{\partial x_i^*}{\partial I}\frac{\partial E}{\partial p_j}. \tag{14}$$

Now we have to find an expression for $\frac{\partial E}{\partial p_j}$. Start with the Lagrangian associated with problem (12) evaluated at the optimal solution $(h^*(p_1, p_2, \bar{u}),\ \lambda^*(p_1, p_2, \bar{u}))$:

$$\mathcal{L}(h^*(p_1, p_2, \bar{u}), \lambda^*(p_1, p_2, \bar{u})) = p_1 h_1^*(p_1, p_2, \bar{u}) + p_2 h_2^*(p_1, p_2, \bar{u}) + \lambda^*(p_1, p_2, \bar{u})[\bar{u} - u(x(p_1, p_2, \bar{u}))].$$

Taking the derivative with respect to any price p_j and noting that $\bar{u} = u(x(p, \bar{u}))$ at the optimum we get

$$\begin{aligned} \frac{\partial \mathcal{L}(h^*(p, \bar{u}), \lambda^*(p, \bar{u}))}{\partial p_j} &= h_j^* + \sum_{i=1}^{I} p_i \frac{\partial h_i^*}{\partial p_j} - \lambda^* \sum_{i=1}^{I} \frac{\partial u}{\partial x_i}\frac{\partial x_i}{\partial p_j} \\ &= h_j^* + \sum_{i=1}^{I} \left(p_i - \lambda^* \frac{\partial u}{\partial x_i} \right) \frac{\partial x_i}{\partial p_j}. \end{aligned}$$

But the first -order conditions for this Lagrangian are

$$p_i - \lambda \frac{\partial u}{\partial x_i} = 0 \text{ for all } i.$$

Hence

$$\frac{\partial E}{\partial p_j} = \frac{\partial \mathcal{L}}{\partial p_j} = h_j^*(p_1, p_2, \bar{u}).$$

This result also follows form the Envelope Theorem. Moreover, from equation (13) it follows that $h_j^* = x_j^*$. Hence, using these two facts and bringing the second term on the RHS to the LHS we can rewrite equation (14) as

$$\frac{\partial x_i^*}{\partial p_j} = \underbrace{\frac{\partial h_i^*}{\partial p_j}}_{SE} - x_j^* \underbrace{\frac{\partial x_i^*}{\partial I}}_{IE}.$$

This equation is known as the *Slutsky Equation*[4] and shows formally that the price effect can be separated into a substitution (SE) and an income effect (IE).

[4] After the Russian statistician and economist Eugen Slutsky.

8 Categories of goods and Ealsticities

Definition 11. A **normal good** is a commodity whose Marshallian demand is positively related to income, i.e. as income goes up the uncompensated demand of that good goes up as well. Therefore good i is normal if

$$\frac{\partial x_i^*}{\partial I} > 0$$

Definition 12. A **inferior good** is a commodity whose Marshallian demand is negatively related to income, i.e. as income goes up the uncompensated demand of that good goes down. Therefore good i is inferior if

$$\frac{\partial x_i^*}{\partial I} < 0$$

Definition 13. Two goods are **gross substitutes** if rise in the price of one good raises the uncompensated demand of the other good. Therefore goods i and j are gross substitutes if

$$\frac{\partial x_i^*}{\partial p_j} > 0$$

Definition 14. Two goods are **net substitutes** if rise in the price of one good raises the compensated or Hicksian demand of the other good. Therefore goods i and j are net substitutes if

$$\frac{\partial h_i^*}{\partial p_j} > 0$$

Definition 15. Two goods are **net complements** if rise in the price of one good reduces the compensated or Hicksian demand of the other good. Therefore goods i and j are net complements if

$$\frac{\partial h_i^*}{\partial p_j} < 0$$

8.1 Shape of Expenditure Function

The expression for expenditure function in a n commodity case is given by,

$$E(p_1, p_2, \ldots, p_n, \bar{u}) \triangleq \sum_{i=1}^{n} p_i h_i^*(p_1, p_2, \ldots, p_n, \bar{u})$$

Now let's look at the effect of changing price p_i on the expenditure. By envelope theorem we get that,

$$\frac{\partial E}{\partial p_i} = h_i^*(p_1, p_2, \ldots, p_n, \bar{u}) > 0$$

Therefore the expenditure function is positively sloped, i.e. when prices go up the minimum expenditure required to meet certain utility level also goes up. Now to find out the curvature of the expenditure function we take the second order derivative:

$$\frac{\partial^2 E}{\partial p_i^2} = \frac{\partial h_i^*}{\partial p_i} < 0$$

This implies that the expenditure function is concave in prices.

Definition 16. A **Giffen good** is one whose Marshallian demand is positively related to its price. Therefore good i is Giffen if,

$$\frac{\partial x_i^*}{\partial p_i} > 0$$

But from the Hicksian demand we know that,

$$\frac{\partial h_i^*}{\partial p_i} < 0$$

Hence from the Slutsky equation,

$$\frac{\partial x_i^*}{\partial p_i} = \frac{\partial h_i^*}{\partial p_i} - x_1^* \frac{\partial x_i^*}{\partial I}$$

we get that for a good to be Giffen we must have,

$$\frac{\partial x_i^*}{\partial I} < 0$$

and x_i^* needs to be large to overcome the effect of substitution effect.

Definition 17. A **luxury good** is defined as one for which the elasticity of income is greater than one. Therefore for a luxury good i,

$$\epsilon_{i,I} = \frac{\frac{dx_i^*}{dI}}{\frac{x_i^*}{I}} > 1$$

We can also define luxury good in the following alternative way.

Definition 18. If the budget share of a good is increasing in income then it is a **luxury good**.

Before we explain the equivalence of the two definitions let us first define the concept of budget share. Budget share of good i, denoted by $s_i(I)$, is the fraction of income I that is devoted to the expenditure on that good. Therefore,

$$s_i(I) = \frac{p_i x_i^*(p, I)}{I}$$

Now to see how the two definitions are related we take the derivative of $s_i(I)$ w.r.t. I.

$$\frac{ds_i(I)}{dI} = \frac{\frac{dx_i^*}{dI} p_i I - p_i x_i^*}{I^2}$$

Now if good i is luxry then we know that,

$$\frac{ds_i(I)}{dI} > 0$$

$$\iff \frac{\frac{dx_i^*}{dI} p_i I - p_i x_i^*}{I^2} > 0$$

$$\iff \frac{dx_i^*}{dI} p_i I - p_i x_i^* > 0$$

$$\iff \frac{\frac{dx_i^*}{dI}}{\frac{x_i^*}{I}} > 1$$

$$\iff \epsilon_{i,I} > 1$$

Therefore we see that the two definitions of luxury good are equivalent. Hence a luxury good is one which a consumer spends more, proportionally, as her income goes up.

8.2 Elasticities

Definition 19. Revenue from a good i is defined as the following:

$$R_i(p_i) = p_i x_i^*$$

Differentiating $R_i(p_i)$ w.r.t. p_i we get,

$$\begin{aligned} R_i'(p_i) &= x_i^* + p_i \frac{dx_i^*}{dp_i} \\ &= x_i^* \left[1 + \frac{p_i}{x_i^*} \frac{dx_i^*}{dp_i}\right] \\ &= x_i^* \left[1 + \epsilon_{i,i}\right] \end{aligned}$$

We say that:

Demand is inelastic if,

$$R_i^{'}(p_i) > 0 \Rightarrow \epsilon_{i,i} \in (-1, 0)$$

Demand is elastic if,

$$R_i^{'}(p_i) < 0 \Rightarrow \epsilon_{i,i} \in (-\infty, -1)$$

Demand is unit-elastic if,

$$R_i^{'}(p_i) = 0 \Rightarrow \epsilon_{i,i} = -1$$

9 Welfare Measurement

In order to do welfare comparison of different price situations it is important that we move out of the utility space and deal with money as then we would have an objective measure that we can compare across individual unlike utility. Let the initial price vector be given by,

$$p^0 = (p_1^0, p_2^0, \ldots, p_n^0)$$

and the new price vector be,

$$p^1 = (p_1^1, p_2^1, \ldots, p_n^1)$$

9.1 Compensating Variation

The notion of compensating variation asks how much additional amounts of income is required to maintain the initial level of utility under new prices.

$$CV = E(p^1, u_0) - E(p^0, u_0)$$

where $E(p^0, u_0)$ is the expenditure function evaluated at price p^0 and utility level u_0. This gives us a measure of loss or gain of welfare of one individual in terms of money due to change in prices.

9.2 Equivalent Variation

The notion of equivalent variation asks how much additional amounts of income is required to raise the level of utility from the initial level to a specified new level given the same prices.

$$EV = E(p^0, u_1) - E(p^0, u_0)$$

Let the change in price from p^0 to p^1 is only through the change in price of commodity 1. Let $p_1^1 > p_1^0$ and $p_i^1 = p_i^0$ for all other $i = 2, 3, \ldots, n$. Then we can write,

$$\begin{aligned} CV &= E(p^1, u_0) - E(p^0, u_0) \\ &= \int_{p_1^0}^{p_1^1} \frac{\partial E(p_1, u_0)}{\partial p_1} dp_1 \\ &= \int_{p_1^0}^{p_1^1} h_1^*(p_1, u_0) dp_1 \end{aligned}$$

9.3 Introduction of New Product

Let's think of a scenario where a new product is introduced. Let that be commodity k. This can be thought of as a reduction of the price of that product from $p_k = \infty$ to $p_k = \bar{p}$ where $\bar{p}$ is the price of the new product. Then one can measure the welfare gain of introducing a new product by calculating the CV with the change in price of the new product from infinity to $\bar{p}$.

$$CV = -\int_{\bar{p}}^{\infty} h_k^*(p_k, u_0) dp_k$$

9.4 Inflation Measurement

Let the reference consumption bundle be denoted by,

$$x^0 = (x_1^0, x_2^0, \ldots, x_n^0)$$

and the reference price be,

$$p^0 = (p_1^0, p_2^0, \ldots, p_n^0)$$

Then one measure of inflation is the Laspayers Price Index,

$$I_L = \frac{p^1 \cdot x^0}{p^1 \cdot x^0}$$

The other measure is Paasche Price Index,

$$I_P = \frac{p^1 \cdot x^1}{p^1 \cdot x^1}$$

where x^1 is the consumption bundle purchased at the new price p^1. Here the reference bundle x^0 is the optimal bundle under the price situation p^0. Therefore we can say,

$$p^0 \cdot x^0 = E(p^0, u_0)$$

where u_0 is the utility level achieved with price p^0 and income $p^0 \cdot x^0$. Now given the new price situation p^1 we know that,

$$p^1 \cdot x^0 \geq E(p^1, u_0)$$

$$\Rightarrow I_L = \frac{p^1 \cdot x^0}{p^1 \cdot x^0} \geq \frac{E(p^1, u_0)}{E(p^0, u_0)}$$

Hence we see that the Laspayers Price Index is an overestimation of price change.

10 Pareto Efficiency and Competitive Equilibrium

We now consider a model with many agents where we make prices endogenous (initially) and later incomes as well. Let there be I individuals, each denoted by i,

$$i = 1, 2, \ldots, I$$

K commodities, each denoted by k,

$$k = 1, 2, \ldots., K$$

a consumption bundle of agent i be denoted by x^i,

$$x^i = (x_1^i, x_2^i, \ldots, x_K^i)$$

and the utility function of individual i be denoted by,

$$u^i \colon \mathbb{R}_+^K \to \mathbb{R}$$

and society has endowment of commodities denoted by e,

$$e = (e_1, e_2, \ldots, e_K)$$

A social allocation is a vector of consumption bundles for all the individuals,

$$x = (x^1, x^2, \ldots, , x^i, \ldots, x^I)$$

The total consumption of commodity k by all the individuals can not exceed the endowment of that commodity, which is referred to as the feasibility constraint. We say a social allocation is feasible if,

$$\sum_{i=1}^{I} x_k^i \leq e_k \;\; \forall \; k = 1, 2, \ldots, K$$

which represent the K feasibility constraints.

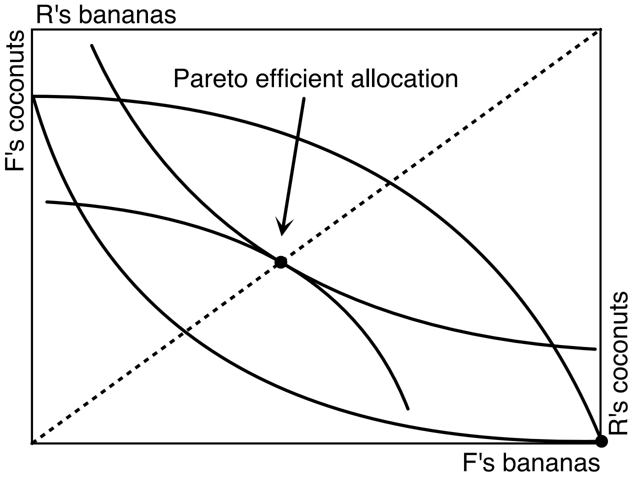

Definition 20. An allocation x is **Pareto efficient** if it is feasible and there exists no other feasible allocation y such that nobody is worse off and at least one individual is strictly better off, i.e. there is no y such that for all i:

$$u^i(y^i) \geq u^i(x^i)$$

and at for some i':

$$u^{i'}(y^{i'}) > u^{i'}(x^{i'})$$

We say that an allocation y is **Pareto superior** to another allocation x if for all i:

$$u^i(y^i) \geq u^i(x^i)$$

and at for some i':

$$u^{i'}(y^{i'}) > u^{i'}(x^{i'})$$

and we say that y Pareto dominates x if, for all i:

$$u^i(y^i) > u^i(x^i)$$

In a 2 agent (say, Robinson and Friday), 2 goods (say, coconuts and bananas) economy we can represent the allocations in an Edgeworth box. Note that we have a total of four axes in the Edgeworth box. The origin for Friday is in the south-west corner and the amount of bananas he consumes is measured along the lower horizontal axis whereas his amount of coconuts is measured along the left vertical axis. For Robinson, the origin is in the north-east corner, the upper horizontal axis depicts Robinson's banana consumption, and the right vertical axis measures his coconut consumption. The height and width of the Edgeworth box are one each since there are one banana and one coconut in this economy. Hence, the endowment bundle is the south-east corner where the amount of Friday's bananas and Robinson's coconuts are both equal to one. This also implies that Friday's utility increases as he moves up and right in the Edgeworth box, whereas Robinson is better off the further down and left he gets. Any point inside the the two ICs is an allocation that gives both Robinson and Friday higher utility. Hence any point inside is a Pareto superior allocation than the initial one. The point where the two ICs are tangent to each other is a Pareto efficient point as starting from that point or allocation, it is not possible to raise one individual's utility without reducing other's. Hence the set of Pareto efficient allocations in this economy is the set of points in the Edgeworth box where the two ICs are tangent to each other. This is depicted as the dotted line in the box. It is evident from the picture that there can be many Pareto efficient allocations. Specifically, allocations that give all the endowment of the society to either Robinson or Friday are also Pareto efficient as as any other allocation would reduce that person's utility.

10.1 Competitive Equilibrium

A competitive equilibrium is the pair (p, x), where p is the price vector for the K commodities:

$$p = (p_1, \ldots, p_k, \ldots, p_K)$$

and x is the allocation:

$$x = (x^1, x^2, \ldots, , x^i, \ldots, x^I),$$

such that markets clear for all commodities k:

$$\sum_{i=1}^{I} x_k^i \leq e_k,$$

allocation is affordable for each individual i:

$$p \cdot x^i \leq p \cdot e^i,$$

and for each individual i there is no y^i such that

$$p \cdot y^i \leq p \cdot e^i$$

and

$$u^i(y^i) > u^i(x^i)$$

11 Social Welfare

We here are trying to formalize the problem from the point of view of a social planner. The social planner has endowments given by the endowment vector $e = (e_1, e_2, \ldots, e_K)$ and attaches weight α^i to individual i's utility. So for him the optimization problem is given by:

$$\max_{x^1, x^2, \ldots, x^I} \sum_{i=1}^{I} \alpha^i u^i(x^i) \quad \alpha^i \geq 0, \ \sum_{i}^{I} \alpha^i = 1$$

$$\text{subject to} \ \sum_{i=1}^{I} x_k^i \leq e_k \ \forall \ k = 1, 2, \ldots, K$$

The Lagrange is given by,

$$L(x, \lambda) = \sum_{i=1}^{I} \alpha^i u^i(x^i) + \sum_{k=1}^{K} \lambda_k \left(e_k - \sum_{i=1}^{I} x_k^i \right)$$

The first order conditions for individual i and for any two goods k and l are:

$$\begin{aligned} x_k^i &: \quad \alpha^i \frac{\partial u^i\left(x^i\right)}{\partial x_k^i} - \lambda_k = 0, \\ x_l^i &: \quad \alpha^i \frac{\partial u^i\left(x^i\right)}{\partial x_l^i} - \lambda_l = 0. \end{aligned}$$

If we consider the ratio for any two commodities, we get for all i and for any pair k, l of commodities:

$$\frac{\frac{\alpha_i \partial u^i\left(x^i\right)}{\partial x_k^i}}{\frac{\alpha_i \partial u^i\left(x^i\right)}{\partial x_l^i}} = \frac{\lambda_k}{\lambda_l}$$

This means that the MRS between two goods k and l is same across individuals which is the condition for Pareto Optimality. Hence a specific profile of weights $\alpha = (\alpha^1, \alpha^2, \ldots, \alpha^I)$ will give us a specific allocation among the set of Pareto efficient allocations. Therefore we have the following powerful theorem:

Theorem 1. *The set of Pareto efficient allocations and the set of welfare maximizing allocations across all possible vectors of weights are identical.*

Below we solve a particular example with Cobb-Douglas preferences.

Example 3. Let Ann and Bob have the following preferences:

$$u^A\left(x_1^A, x_2^A\right) = \alpha \ln x_1^A + (1-\alpha) \ln x_2^A$$

$$u^B\left(x_1^B, x_2^B\right) = \beta \ln x_1^B + (1-\beta) \ln x_2^B$$

Let the weight on Ann's utility function be γ and therefore the weight on Bob's utility function is $(1-\gamma)$. The Lagrange expression is then given by,

$$L(x,\lambda) = \gamma u^A + (1-\gamma)u^B + \lambda_1[e_1 - x_1^A - x_1^B] + \lambda_2[e_2 - x_2^A - x_2^B]$$

The F.O.C.s are then given by,

$$\begin{aligned}
\frac{\partial L(x,\lambda)}{\partial x_1^A} &= \frac{\gamma\alpha}{x_1^A} - \lambda_1 = 0 \\
\frac{\partial L(x,\lambda)}{\partial x_2^A} &= \frac{\gamma(1-\alpha)}{x_2^A} - \lambda_2 = 0 \\
\frac{\partial L(x,\lambda)}{\partial x_1^B} &= \frac{(1-\gamma)\beta}{x_1^B} - \lambda_1 = 0 \\
\frac{\partial L(x,\lambda)}{\partial x_2^B} &= \frac{(1-\gamma)\beta}{x_2^B} - \lambda_2 = 0
\end{aligned}$$

Hence we get,

$$(1-\gamma)\beta x_1^A = \gamma\alpha x_1^B$$

and

$$(1-\gamma)\beta x_2^A = \gamma(1-\alpha) x_2^B$$

Thus from the feasibility conditions we get that the allocations would be,

$$x_1^A = \frac{\gamma\alpha}{\gamma\alpha + (1-\gamma)\beta} e_1, x_2^A = \frac{\gamma(1-\alpha)}{(1-\gamma)\beta + \gamma(1-\alpha)} e_2$$

$$x_1^B = \frac{(1-\gamma)\beta}{\gamma\alpha + (1-\gamma)\beta} e_1, x_2^B = \frac{(1-\gamma)\beta}{(1-\gamma)\beta + \gamma(1-\alpha)} e_2$$

Here by varying the value of γ in its range [0,1] we can generate the whole set of Pareto efficient allocations.

12 Competitive Equilibrium Continued

Example 4. We now consider a simple example, where Friday is endowed with the only (perfectly divisible) banana and Robinson is endowed with the only coconut. That is $e^F = (1,0)$ and $e^R = (0,1)$. To keep things simple suppose that both agents have the same utility function

$$u(x_B, x_C) = \alpha\sqrt{x_B} + \sqrt{x_C}$$

and we consider the case where $\alpha > 1$, so there is a preference for bananas over coconuts that both agents share. We can determine the indifference curves for both Robinson and Friday that correspond to the same utility level that the initial endowments provide. The indifference curves are given by

$$\begin{aligned} u^F\left(e^F_{B,}e^F_C\right) &= \alpha\sqrt{e^F_B} + \sqrt{e^F_C} = \alpha = u^F(1,0) \\ u^R\left(e^R_{B,}e^R_C\right) &= \alpha\sqrt{e^R_B} + \sqrt{e^R_C} = 1 = u^R(0,1) \end{aligned}$$

All the allocations between these two indifference curves are Pareto superior to the initial endowment.
We can define the net trade for Friday (and similarly for Robinson) by

$$\begin{aligned} z^F_B &= x^F_B - e^F_B \\ z^F_C &= x^F_C - e^F_C \end{aligned}$$

Notice that since initially Friday had all the bananas and none of the coconuts

$$\begin{aligned} z^F_B &\leq 0 \\ z^F_C &\geq 0 \end{aligned}$$

There could be many Pareto efficient allocations (e.g., Friday gets everything, Robinson gets everything, etc.), but we can calculate which allocations are Pareto optimal. If the indifference curves at an allocation are tangent then the marginal rates of substitution must be equated. In this case, the resulting condition is

$$\frac{\frac{\partial u^F}{\partial x^F_B}}{\frac{\partial u^F}{\partial x^F_C}} = \frac{\frac{\alpha}{2\sqrt{x^F_B}}}{\frac{1}{2\sqrt{x^F_C}}} = \frac{\frac{\alpha}{2\sqrt{x^R_B}}}{\frac{1}{2\sqrt{x^R_C}}} = \frac{\frac{\partial u^R}{\partial x^R_B}}{\frac{\partial u^R}{\partial x^R_C}}$$

which simplifies to

$$\frac{\sqrt{x^F_C}}{\sqrt{x^F_B}} = \frac{\sqrt{x^R_C}}{\sqrt{x^R_B}}$$

and, of course, since there is a total of one unit of each commodity, for market clearing we must have

$$\begin{aligned} x_C^R &= 1 - x_C^F \\ x_B^R &= 1 - x_B^F \end{aligned}$$

so

$$\frac{\sqrt{x_C^F}}{\sqrt{x_B^F}} = \frac{\sqrt{1 - x_C^F}}{\sqrt{1 - x_B^F}}$$

and squaring both sides

$$\frac{x_C^F}{x_B^F} = \frac{1 - x_C^F}{1 - x_B^F}$$

which implies that

$$x_C^F - x_C^F x_B^F = x_B^F - x_C^F x^B$$

and so

$$\begin{aligned} x_C^F &= x_B^F \\ x_C^R &= x_B^R. \end{aligned}$$

What are the conditions necessary for an equilibrium? First we need the conditions for Friday to be optimizing. We can write Robinson's and Friday's optimization problems as the corresponding Lagrangian, where we generalize the endowments to any $e^R = (e_B^R, e_C^R)$ and $e^F = (e_B^F, e_C^F)$:

$$\mathcal{L}\left(x_B^F, x_C^F, \lambda^F\right) = \alpha\sqrt{x_B^F} + \sqrt{x_C^F} + \lambda\left(p_B e_B^F + e_C^F - p_B x_B^F - x_C^F\right), \qquad (15)$$

where we normalize $p_C = 1$ without loss of generality. A similar Lagrangian can be set up for Robinson's optimization problem. The first-order conditions for (15) are

$$\frac{\partial \mathcal{L}}{\partial x_B^F} = \frac{\alpha}{2\sqrt{x_B^F}} - \lambda^F p_B = 0 \qquad (16)$$

$$\frac{\partial \mathcal{L}}{\partial x_C^F} = \frac{1}{2\sqrt{x_C^F}} - \lambda^F = 0 \qquad (17)$$

$$\frac{\partial \mathcal{L}}{\partial \lambda^F} = p_B e_B^F + e_C^F - p_B x_B^F - x_C^F = 0. \qquad (18)$$

Solving as usual by taking the ratio of equations (16) and (17) we get the following expression for the relative (to coconuts) price of bananas

$$p_B = \alpha\frac{\sqrt{x_C^F}}{\sqrt{x_B^F}}$$

so that we can solve for x_C^F as a function of x_B^F

$$x_C^F = \left(\frac{p_B}{\alpha}\right)^2 x_B^F.$$

Plugging this into the budget constraint from equation (18) we get

$$p_B x_B^F + \left(\frac{p_B}{\alpha}\right)^2 x_B^F = p_B e_B^F + e_C^F.$$

Then we can solve for Friday's demand for bananas

$$x_B^F = \frac{p_B e_B^F + e_C^F}{p_B + \left(\frac{p_B}{\alpha}\right)^2}$$

and for coconuts

$$x_C^F = \left(\frac{p_B}{\alpha}\right)^2 \frac{p_B e_B^F + e_C^F}{p_B + \left(\frac{p_B}{\alpha}\right)^2}.$$

The same applies to Robinson's demand functions, of course.

Now we have to solve for the equilibrium price p_B. To do that we use the market clearing condition for bananas, which says that demand has to equal supply (endowment):

$$x_B^F + x_B^R = e_B^F + e_B^R.$$

Inserting the demand functions yields

$$\frac{p_B e_B^F + e_C^F}{p_B + \left(\frac{p_B}{\alpha}\right)^2} + \frac{p_B e_B^R + e_C^R}{p_B + \left(\frac{p_B}{\alpha}\right)^2} = e_B^F + e_B^R = e_B,$$

where e_B is the social endowment of bananas and we define $e_C = e_C^F + e_C^R$. We solve this equation to get the equilibrium price of bananas in the economy:

$$p_B^* = \alpha\sqrt{\frac{e_C}{e_B}}.$$

So we have solved for the equilibrium price in terms of the primitives of the economy. This price makes sense intuitively. It reflects relative scarcity in the economy (when there are relatively more bananas than coconuts, bananas are cheaper) and preferences (when consumers value bananas more, i.e., when α is larger, they cost more). We can then plug this price back into the previously found equations both for agents' consumption and have an expression for consumption in terms of the primitives.

Now we mention the two fundamental welfare theorems which lay the foundation for taking competitive markets as the benchmark for any study of markets and prices. The first one states that competitive equilibrium allocations are always Pareto efficient and the second one states that any Pareto efficient allocation can be achieved as an outcome of competitive equilibrium.

Theorem 2. *(First Welfare Theorem) Every Competitive Equilibrium allocation* x^* *is Pareto Efficient.*

Proof. Suppose not. Then there exists another allocation y, which is feasible, such that

$$\text{for all } i\text{: } u^i(y) \geq u^i(x^*)$$
$$\text{for some } i'\text{: } u^{i'}(y) > u^{i'}(x^*).$$

If $u^i(y) \geq u^i(x^*)$, then the budget constraint (and monotone preferences) implies that

$$\sum_{k=1}^{K} p_k y_k^i \geq \sum_{k=1}^{K} p_k x_k^{*k} \tag{19}$$

and for some i'

$$\sum_{k=1}^{K} p_k y_k^{i'} > \sum_{k=1}^{K} p_k x_k^{*i'}. \tag{20}$$

Equations (19) and (20) imply that

$$\sum_{i=1}^{I} \sum_{k=1}^{K} p_k y_k^i > \sum_{i=1}^{I} \sum_{k=1}^{K} p_k x_k^{*i} = \sum_{k=1}^{K} p_k e_k,$$

where the left-most term is the aggregate expenditure and the right-most term the social endowment. This is a contradiction because feasibility of y means that

$$\sum_{i=1}^{I} y_k^i \leq \sum_{i=1}^{I} e_k^i = e_k$$

for any i and hence

$$\sum_{i=1}^{I} \sum_{k=1}^{K} p_k y_k^i \leq \sum_{k=1}^{K} p_k e_k.$$

□

Theorem 3. *(Second Welfare Theorem) Every Pareto efficient allocation can be decentralized as a competitive equilibrium. That is, every Pareto efficient allocation is the equilibrium for some endowments.*

13 Decision Making under Uncertainty

So far, we have assumed that decision makers have all the needed information. This is not the case in real life. In many situations, individuals or firms make decisions before knowing what the consequences will be. For example, in financial markets investors buy stocks without knowing future returns. Insurance contracts exist because there is uncertainty. If individuals were not uncertain about the possibility of having an accident in the future, there would be no need for car insurance.

Definition 21. $\pi = (\pi_1, \pi_2, \ldots, \pi_N)$ represents a **probability distribution** if

$$\pi_n \geq 0 \;\; \forall\, n = 1, 2, \ldots, N, \quad \text{and} \quad \sum_{n=1}^{N} \pi_n = 1$$

Now to conceptualize uncertainty we define the concept of lottery.

Definition 22. A **lottery** L is defined as follows:

$$L = (x; \pi) = (x_1, x_2, \ldots, x_N; \pi_1, \pi_2, \ldots, \pi_N)$$

where $x = (x_1, x_2, \ldots, x_N) \in \mathbb{R}^N$ is a profile of money awards (positive or negative) to be gained in N different states and $\pi = (\pi_1, \pi_2, \ldots, \pi_N)$ is the probability distribution over the N states.

13.1 St. Petersberg Paradox

Here we talk about a well-known lottery known as the *St. Petersburg paradox* which was proposed by Daniel Bernoulli in 1736. A fair coin is tossed until head comes up for the first time. Then the reward paid out is equal to 2^{n-1}, where n is the number of coin tosses that were necessary for head to come up once. This lottery is described formally as

$$L_{SP} = \left\{ 1, 2, 4, \ldots, 2^{n-1}, \ldots; \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \ldots, \frac{1}{2^n}, \ldots \right\}.$$

Its expected value is

$$E\left[L_{SP}\right] = \sum_{n=1}^{\infty} \pi_n x_n = \sum_{n=1}^{\infty} \frac{1}{2^n} 2^{n-1} = \sum_{n=1}^{\infty} \frac{1}{2} = \infty.$$

Hence the expected payoff from this lottery is infinitely large and an individual offered this lottery should be willing to pay an infinitely large amount for the right to play this lottery. This is not, however, what people do and hence the paradox.

13.2 Expected Utility

The St. Petersberg paradox emphasized the fact that expected value may not be the right way to describe individual's preferences over lotteries. In general utility over lotteries is a function:

$$U\colon \mathbb{R}^N \times [0,1]^N \to \mathbb{R}$$

Expected Utility is a particular formulation that says that there is another utility function defined over money

$$u\colon \mathbb{R} \to \mathbb{R}$$

such that the utility over the lottery is of the following form:

$$U(x_1, x_2, \ldots, x_N; \pi_1, \pi_2, \ldots, \pi_N) = \sum_{n=1}^{N} u(x_n)\pi_n$$

Definition 23. A decision maker is called *risk averse* if the utility function $u : \mathbb{R} \longrightarrow \mathbb{R}_+$ is concave and she is called *risk loving* or a risk seeker if u is convex.

Suppose a lottery be given by

$$L = (x_1, x_2; \pi_1, \pi_2)$$

Then the individual is risk averse if

$$\pi_1 u(x_1) + \pi_2 u(x_2) < u(\pi_1 x_1 + \pi_2 x_2)$$

and the individual is risk loving if

$$\pi_1 u(x_1) + \pi_2 u(x_2) > u(\pi_1 x_1 + \pi_2 x_2)$$

and the individual is risk neutral if

$$\pi_1 u(x_1) + \pi_2 u(x_2) = u(\pi_1 x_1 + \pi_2 x_2)$$

13.3 Risky Investment

Let an individual with wealth w deciding how much to invest in a risky asset which pays return r_1 in state 1 and return r_2 in state 2 such that

$$(1 + r_1) < 1$$

$$(1 + r_2) > 1$$

Therefore state 1 is the bad state which gives a negative return while state 2 is the good state which gives a positive return. If z is the amount that's invested in this risky asset then the individual's expected utility is given by,

$$U(z) = \pi_1 u((1+r_1)z - z + w) + \pi_2 u((1+r_2)z - z + w)$$

So the individual solves the following problem:

$$\max_z \pi_1 u((1+r_1)z - z + w) + \pi_2 u((1+r_2)z - z + w)$$

The marginal utility of investment is given by,

$$\frac{dU(z)}{dz} = \pi_1 u'((1+r_1)z - z + w).r_1 + \pi_2 u'((1+r_2)z - z + w).r_2$$

Therefore the MU at $z = 0$ is given by,

$$\left[\frac{dU(z)}{dz}\right]_{z=0} = \pi_1 r_1 u'(w) + \pi_2 r_2 u'(w) = [\pi_1 r_1 + \pi_2 r_2]u'(w)$$

Therefore,

$$\left[\frac{dU(z)}{dz}\right]_{z=0} \gtreqless 0 \quad \text{according as} \quad \pi_1 r_1 + \pi_2 r_2 \gtreqless 0$$

Hence whether the individual will invest anything in the asset or not will depend on the expected value of the asset. If the expected value is positive he will invest a positive amount irrespective of his degree of risk aversion. The actual value of z will of course depend on the concavity of his utility function.

14 Theory of Production

We can use tools similar to those we used in the consumer theory section of the class to study firm behaviour. In that section we assumed that individuals maximize utility subject to some budget constraint. In this section we assume that firms will attempt to maximize their profits given a demand schedule and production technology.

Firms use inputs or commodities $x_1, \ldots, x_I$ to produce an output y. The amount of output produced is related to the inputs by the production function $y = f(x_1, \ldots, x_I)$, which is formally defined as follows:

Definition 24. A *production function* is a mapping $f : \mathbb{R}^I_+ \longrightarrow \mathbb{R}_+$.

The prices of the inputs/commodities are $p_1, \ldots, p_I$ and the output price is p_y. The firm takes prices as given and independent of its decisions.

Firms maximize their profits by choosing the optimal amount and combination of inputs.

$$\max_{x_1,\ldots,x_I} p_y f(x_1, \ldots, x_I) - \sum_{i=1}^{I} p_i x_i. \tag{21}$$

Another way to describe firms' decision making is by minimizing the cost necessary to produce an output quantity $\bar{y}$.

$$\min_{x_1,\ldots,x_I} \sum_{i=1}^{I} p_i x_i \text{ s.t. } f(x_1, \ldots, x_I) \geq \bar{y}.$$

The minimized cost of production, $C(\bar{y})$, is called the cost function.

We make the following assumption for the production function: positive marginal product

$$\frac{\partial f}{\partial x_i} \geq 0$$

and declining marginal product

$$\frac{\partial^2 f}{\partial x_i^2} \leq 0.$$

The optimality conditions for the profit maximization problem (21) and the FOCs for all i

$$p_y \frac{\partial f}{\partial x_i} - p_i = 0.$$

In other words, optimal production requires equality between marginal benefits and marginal cost of production. The solution to the profit maximization problem then is

$$x_i^*(p_1, \ldots, p_I, p_y), i = 1, \ldots, I$$
$$y^*(p_1, \ldots, p_I, p_y),$$

i.e., optimal demand for inputs and optimal output/supply.

The solution of the cost minimization problem (), on the other hand is

$$x_i^*(p_1, \ldots, p_I, \bar{y}), i = 1, \ldots, I,$$

where $\bar{y}$ is the firm's production target.

Example 5. One commonly used production function is the Cobb-Douglas production function where

$$f(K, L) = K^\alpha L^{1-\alpha}$$

The interpretation is the same as before with α reflecting the relative importance of capital in production. The marginal product of capital is $\frac{\partial f}{\partial K}$ and the marginal product of labor is $\frac{\partial f}{\partial L}$.

In general, we can change the scale of a firm by multiplying both inputs by a common factor: $f(tK, tL)$ and compare the new output to $tf(K, L)$. The firm is said to have *constant returns to scale* if

$$tf(K, L) = f(tK, tL),$$

it has *decreasing returns to scale* if

$$tf(K, L) > f(tK, tL),$$

and *increasing returns to scale* if

$$tf(K, L) < f(tK, tL).$$

Example 6. The Cobb-Douglas function in our example has constant returns to scale since

$$f(tK, tL) = (tK)^{\alpha}(tL)^{1-\alpha} = tK^{\alpha}L^{1-\alpha} = tf(K, L).$$

Returns to scale have an impact on market structure. With decreasing returns to scale we expect to find many small firms. With increasing returns to scale, on the other hand, there will be few (or only a single) large firms. No clear prediction can be made in the case of constant returns to scale. Since increasing returns to scale limit the number of firms in the market, the assumption that firms are price takers only makes sense with decreasing or constant returns to scale.

15 Imperfect Competition

15.1 Pricing Power

So far, we have considered market environments where single agent cannot control prices. Instead, each agent was infinitesimally small and firms acted as price takers. This was the case in competitive equilibrium. There are many markets with few (oligopoly) or a single firms, (monopoly) however. In that case firms can control prices to some extent. Moreover, when there are a few firms in a market, firms make interactive decisions. In other words, they take their competitors' actions into account. In Section **??**, we will use game theory to analyse this type of market structure. First, we cover monopolies, i.e., markets with a single producer.

15.2 Monopoly

If a firm produces a non-negligible amount of the overall market then the price at which the good sells will depend on the quantity sold. Examples for firms that control the overall market include the East India Trading Company, Microsoft (software in general because of network externalities and increasing returns to scale), telecommunications and utilities (natural monopolies), Standard Oil, and De Beers.

For any given price there will be some quantity demanded by consumers, and this is known as the demand curve $x : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$ or simply $x(p)$. We assume that consumers demand less as the price increases: the demand function is downward sloping or $x'(p) < 0$. We can invert this relationship to get the inverse demand function $p(x)$ which reveals the price that will prevail in the market if the output is x.

If the firm is a monopolist that takes the demand data $p(x)$ as given then its goal is to maximize

$$\pi(x) = p(x)x - c(x) \tag{22}$$

by choosing the optimal production level. For the cost function we assume $c'(x) > 0$ and $c''(x) \geq 0$, i.e., we have positive and weakly increasing marginal costs. For example, $c(x) = cx$ satisfies these assumptions (a Cobb-Douglas production function provides this for example). The monopolist maximizes its profit function (22) over x, which leads to the following FOC:

$$p(x) + xp'(x) - c'(x) = 0. \tag{23}$$

Here, in addition to the familiar $p(x)$, which is the marginal return from the marginal consumer, the monopolist also has to take the $xp'(x)$ into account, because a change in quantity also affects the inframarginal consumers. For example, when it increases the quantity supplies, the monopolist gets positive revenue from the marginal consumer, but the inframarginal consumers pay less due to the downward sloping demand function. At the optimum, the monopolist equates marginal revenue and marginal cost.

Example 7. A simple example used frequently is $p(q) = a - bq$, and we will also assume that $a > c$ since otherwise the cost of producing is higher than any consumer's valuation so it will never be profitable for the firm to produce and the market will cease to exist. Then the firm want to maximize the objective

$$\pi(x) = (a - bx - c)x.$$

The efficient quantity is produced when $p(x) = a - bx = c$ because then a consumer buys an object if and only if they value it more than the cost of producing, resulting

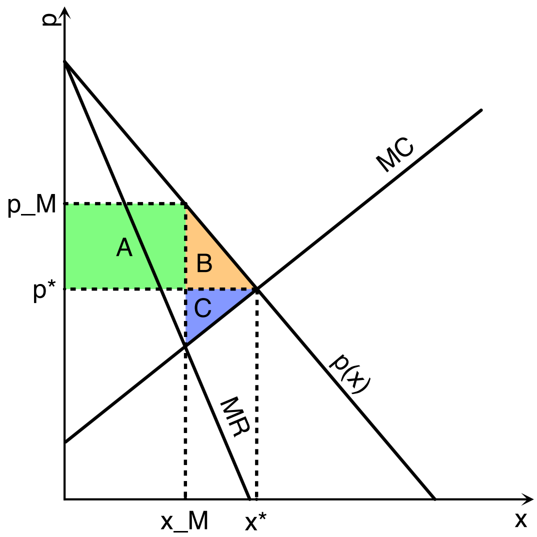

Figure 5: Monopoly

in the highest possible total surplus. So the efficient quantity is

$$x^* = \frac{a-c}{b}.$$

The monopolist's maximization problem, however, has FOC

$$a - 2bx - c = 0$$

where $a - 2bx$ is the marginal revenue and c is the marginal cost. So the quantity set by the monopolist is

$$x^M = \frac{a-c}{2b} < x^*.$$

The price with a monopoly can easily be found since

$$\begin{aligned} p^M &= a - bx^M \\ &= a - \frac{a-c}{2} \\ &= \frac{a+c}{2} \\ &> c. \end{aligned}$$

Figure 5 illustrates this.

A monopoly has different welfare implications than perfect competition. In Figure 5, consumers in a monopoly lose the areas A and B compared to perfect competition. The monopolist loses area C and wins area A. Hence, there are distributional implications (consumers lose and the producer gains) as well as efficiency implications (overall welfare decreases).

We can write the monopolist's FOC (23) in terms of the demand elasticity introduced in Section **??** as follows:

$$\begin{aligned} p(x^*) + x^* p'(x^*) &= c'(x^*) \Longleftrightarrow \\ p(x^*)\left[1 + \frac{x^* p'(x^*)}{p(x^*)}\right] &= c'(x^*) \Longleftrightarrow \\ p(x^*) &= \frac{c'(x^*)}{1 + \epsilon_p^{-1}}. \end{aligned}$$

Since $\epsilon_p < 0$, we have that $p(x^*) > c'(x^*)$, in other words, the monopolist charges more than the marginal cost. This also means that if demand is very elastic, $\epsilon_p \to \infty$, then $p(x^*) \approx c'(x^*)$. On the other hand, if demand is very inelastic, $\epsilon_p \approx -1$, then $p(x^*) \gg c'(x^*)$.

16 Imperfectly Competitive Market

16.1 Price Discrimination

In the previous section we saw that the monopolist sets an inefficient quantity and total welfare is decreased. Is there a mechanism, which allows the monopolist to offer the efficient quantity and reap the entire possible welfare in the market? The answer is yes if the monopolist can set a two-part tariff, for example. In general, the monopolist can extract consumer rents by using price discrimination.

First degree price discrimination (perfect price discrimination) means discrimination by the identity of the person or the quantity ordered (non-linear pricing). It will result in an efficient allocation. Suppose there is a single buyer and a monopoly seller where the inverse demand is given by $p = a - bx$. If the monopolist were to set a single price it would set the monopoly price. As we saw in the previous section, however, this does not maximize the joint surplus, so the monopolist can do better. Suppose instead that the monopolist charges a fixed fee F that the consumer has to pay to be allowed to buy any positive amount at all, and then sells the good at a price p, and suppose the monopolist sets the price $p = c$. The fixed fee will not affect the quantity that a participating consumer will choose, so if the consumer participates then they will choose quantity equal to x^*. The firm can then set the entry fee to extract all the consumer surplus and the consumer will still be willing to participate. This maximizes the joint surplus, and gives the entire surplus to the firm, so the firm is doing as well as it could under any other mechanism. Specifically, using the functional form from Example 7 the firm sets

$$F = \frac{(a-c)x^*}{2} = \frac{(a-c)^2}{2b}$$

In integral notation this is

$$F = \int_0^{x^*} (p(x) - c)dx.$$

This pricing mechanism is called a two-part tariff, and was famously used at Disneyland (entry fee followed by a fee per ride), greatly increasing revenues.

Now, let's assume that there are two different classes of consumers, type A with utility function $u(x)$ and type B with $\beta u(x), \beta > 1$, so that the second class of consumers has a higher valuation of the good. If the monopolist structures a two-part tariff $(F, p = c)$ to extract all surplus from type B consumers, type A consumers would not pay the fixed fee F since they could not recover the utility lost from using the service. On the other hand, if the firm offers two two-part tariffs $(F_A, p = c)$ and $(F_B, p = c)$ with $F_A < F_B$, all consumers would pick the cheaper contract $(F_A, p = c)$. A solution to this problem would be to offer the contracts $(F_A, p_A > c)$ and $(F_B, p = c)$. Type A consumers pick the first contract and consume less of the good and type B consumers pick the second contract, which allows them to consume the efficient quantity. This is an example for *second degree price discrimination*, which means that the firm varies the price by quantity or quality only. It offers a menu of choices and lets the consumers self-select into their preferred contract.

In addition, there is *third degree price discrimination*, in which the firm varies the price by market or identity of the consumers. For example, Disneyland can charge different prices in different parks. Let's assume there are two markets, $i = 1, 2$. The firm is a monopolist in both markets and its profit maximization problem is

$$\max_{x_1, x_2} x_1 p_1(x_1) + x_2 p_2(x_2) - c(x_1 + x_2).$$

The FOC for each market is

$$p_i(x_i) + x_i p_i'(x_i) = c'(x_1 + x_2),$$

which leads to optimal solution

$$p_i(x_i^*) = \frac{c'(x_1^* + x_2^*)}{1 + \frac{1}{\epsilon_p^i}} \text{ for } i = 1, 2.$$

Hence, the solution depends on the demand elasticity in market i. The price will be different as long as the structure of demand differs.

16.2 Oligopoly

Oligopoly refers to environments where there are few large firms. These firms are large enough that their quantity influences the price and so impacts their rivals. Consequently each firm must condition its behavior on the behavior of the other firms. This strategic interaction is modeled with game theory. The most important model of oligopoly is the Cournot model or the model of quantity competition. The general model is described as follows: Let there be I firms denoted by,

$$i = 1, 2, \ldots, I$$

each producing one homogenous good, where each firm produces q_i amount of that good. Each firm i has cost function

$$c_i(q_i)$$

The total production is given by,

$$q = \sum_{i=1}^{I} q_i$$

We also denote total production by firms other than i by,

$$q_{-i} = \sum_{j \neq i} q_j$$

The profit of firm i is given by,

$$\pi_i(q_i, q_{-i}) = p(q_i, q_{-i})q_i - c_i(q_i)$$

So firm i solves,

$$\max_{q_i} \pi_i(q_i, q_{-i})$$

Hence the F.O.C. is given by,

$$p(q_i, q_{-i}) + \frac{\partial p(q_i, q_{-i})}{\partial q_i} q_i - c_i'(q_i) = 0$$

It is important to note that the optimal production of firm i is dependent on the production by other firms i.e. q_{-i}. This the strategic aspect of this model. Therefore in order to produce any amount of the good firm i must anticipate what others might be doing and every firm thinks the same way. So we need an equilibrium concept here which would tell us that given the production level of every firm no firm wants to move away from its current production.

16.2.1 Example

We here consider the duopoly case, where there are only two firms. Suppose the inverse demand function is given by $p(q) = a - bq$, and the cost of producing is constant and the same for both firms $c_i(q) = cq$. The quantity produced in the market is the sum of what both firms produce $q = q_1 + q_2$. The profits for each firm is then a function of the market price and their own quantity,

$$\pi_i(q_i, q_j) = q_i \left(p\left(q_i + q_j\right) - c\right).$$

The strategic variable that the firm is choosing is the quantity to produce q_i.

Suppose that the firms' objective was to maximize their joint profit

$$\pi_1(q_1, q_2) + \pi_2(q_1, q_2) = \left(q_1 + q_2\right)\left(p\left(q_1 + q_2\right) - c\right)$$

then we know from before that this is maximized when $q_1 + q_2 = q^M$. We could refer to this as the collusive outcome. One way the two firms could split production would be $q_1 = q_2 = \frac{q^M}{2}$.

If the firms could write binding contracts then they could agree on this outcome. However, that is typically not possible (such an agreement would be price fixing), so we would not expect this outcome to occur unless it is stable/self-enforcing. If either firm could increase its profits by setting another quantity, then they would have an incentive to deviate from this outcome. We will see below that both firms would in fact have an incentive to deviate and increase their output.

Suppose now that firm i is trying to choose q_i to maximize its own profits, taking the other firm's output as given. Then firm i's optimization problem is

$$\max_{q_i} \pi_i(q_i, q_j) = q_i \left(a - b\left(q_i + q_j\right) - c\right),$$

which has the associated FOC

$$\frac{\partial \pi_i(q_i, q_j)}{\partial q_i} = a - b\left(2q_i + q_j\right) - c = 0.$$

Then the optimal level q_i^* given any level of q_j is

$$q_i^*\left(q_j\right) = \frac{a - bq_j - c}{2b}.$$

This is firm i's *best response* to whatever firm j plays. In the special case when $q_j = 0$ firm i is a monopolist, and the observed quantity q_i corresponds to the monopoly case. In general, when the rival has produced q_j we can treat the firm as a monopolist facing a "residual demand curve" with intercept of $a - bq_j$. We can write firm i's best response function as

$$q_i^*(q_j) = \frac{a - c}{2b} - \frac{1}{2}q_j.$$

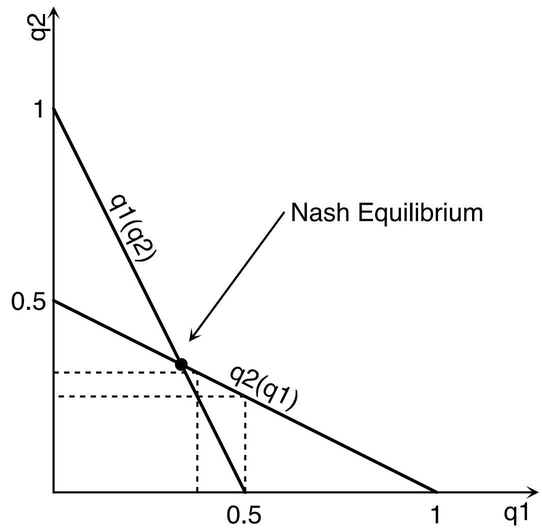

Figure 6: Cournot equilibrium

Hence,

$$\frac{dq_i}{dq_j} = -\frac{1}{2}.$$

This has two important implications. First, the quantity player i chooses is decreasing in its rival's quantity. This means that quantities are *strategic substitutes*. Second, if player j increases their quantity player i decreases their quantity by less than player j increased their quantity (player i decreases his quantity by exactly $\frac{1}{2}$ for every unit player j's quantity is increased). So we would expect that the output in a duopoly would be higher than in a monopoly.

We can depict the best response function graphically. Setting $a = b = 1$ and 0, Figure 6 shows the best response functions. Here, the best response functions are $q_i^*(q_j) = \frac{1-q_j}{2}$.

We are at a "stable" outcome if both firms are producing a best response to their rivals' production. We refer to such an outcome as an equilibrium. That is, when

$$q_i = \frac{a - bq_j - c}{2b} \quad (24)$$

$$q_j = \frac{a - bq_i - c}{2b}. \quad (25)$$

Since the best responses are symmetric we will have $q_i = q_j$ and so we can calculate the equilibrium quantities from the equation

$$q_i = \frac{a - bq_i - c}{2b}$$

and so

$$q_i = q_j = \frac{a-c}{3b}$$

and hence

$$q = q_i + q_j = \frac{2(a-c)}{3b} > \frac{a-c}{2b} = q^M.$$

There is a higher output (and hence lower price) in a duopoly then a monopoly.

More generally, both firms are playing a best response to their rival's action because for all i

$$\pi_i(q_i^*, q_j^*) \geq \pi_i(q_i, q_j^*) \text{ for all } q_i$$

That is, the profits from the quantity are (weakly) higher then the profits from any other output. This motivates the following definition for an equilibrium in a strategic setting.

Definition 25. A *Nash Equilibrium in the duopoly game* is a pair (q_i^*, q_j^*) such that for all i

$$\pi_i(q_i^*, q_j^*) \geq \pi_i(q_i, q_j^*) \text{ for all } q_i.$$

This definition implicitly assumes that agents hold (correct) expectations or beliefs about the other agents' strategies.

A Nash Equilibrium is ultimately a stability property. There is no profitable deviation for any of the players. In order to be at equilibrium we must have that

$$\begin{aligned} q_i &= q_i^*(q_j) \\ q_j &= q_j^*(q_i) \end{aligned}$$

and so we must have that

$$q_i = q_i^*(q_j^*(q_i))$$

so equilibrium corresponds to a fixed-point of the mapping $q_1^*(q_2^*(\cdot))$. This idea can also be illustrated graphically. In Figure 6, firm 1 initially sets $q_1 = \frac{1}{2}$, which is not the equilibrium quantity. Firm 2 then optimally picks $q_2 = q_2^*\left(\frac{1}{2}\right) = \frac{1}{4}$ according to its best response function. Firm 1, in turn, chooses a new quantity according to its best response function: $q_1 = q_1^*\left(\frac{1}{4}\right) = \frac{3}{8}$. This process goes on and ultimately converges to $q_1 = q_2 = \frac{1}{3}$.

16.3 Oligopoly: General Case

Now, we consider the case with I competitors. The inverse demand function (setting $a = b = 1$) is

$$p(q) = 1 - \sum_{i=1}^{I} q_i$$

and firm i's profit function is

$$\pi(q_i, q_{-i}) = \left(1 - \sum_{i=1}^{I} q_i - c\right) q_i, \tag{26}$$

where the vector q_{-i} is defined as $q_{-i} = (q_1, \ldots, q_{i-1}, q_{i+1}, \ldots, q_I)$, i.e., all quantities excluding q_i.

Again, we can define an equilibrium in this market as follows:

Definition 26. A *Nash Equilibrium in the oligopoly game* is a vector $q^* = (q_1^*, \ldots, q_I^*)$ such that for all i

$$\pi_i(q_i^*, q_{-i}^*) \geq \pi_i(q_i, q_{-i}^*) \text{ for all } q_i.$$

We simply replaced the quantity q_j by the vector q_{-i}.

Definition 27. A Nash equilibrium is called *symmetric* if $q_i^* = q_j^*$ for all i and j.

The FOC for maximizing the profit function (26) is

$$1 - \sum_{j \neq i} q_j - 2q_i - c = 0$$

and the best response function for all i is

$$q_i = \frac{1 - \sum_{j \neq i} q_j - c}{2}. \tag{27}$$

Here, only the aggregate supply of firm i's competitors matters, but not the specific amount single firms supply. It would be difficult to solve for I separate values of q_i, but due to symmetry of the profit function we get that $q_i^* = q_j^*$ for all i and j so that equation (27) simplifies to

$$q_i^* = \frac{1 - (I-1)q_i^* - c}{2},$$

which leads to the solution

$$q_i^* = \frac{1 - c}{I + 1}.$$

As I increases (more firms), the market becomes more competitive. Market supply is equal to

$$\sum_{i=1}^{I} q_i^* = I q_i^* = \frac{I}{I+1}(1 - c).$$

As the number of firms becomes larger, $I \longrightarrow \infty$, $q_i^* \longrightarrow 0$ and

$$\sum_{i=1}^{I} q_i^* \longrightarrow 1 - c,$$

which is the supply in a competitive market. Consequently,

$$p^* \longrightarrow c$$

As each player plays a less important strategical role in the market, the oligopoly outcome converges to the competitive market outcome.

Note that we used symmetry in deriving the market outcome from the firms' best response function. We cannot invoke symmetry when deriving the FOC. One might think that instead of writing the profit function as (26) one could simplify it to

$$\pi(q_i, q_{-i}) = (1 - Iq_i - c)\, q_i.$$

This is wrong, however, because it implies that firm i controls the entire market supply (acts as a monopolist). Instead, in an oligopoly market, firm i takes the other firms' output as given.

17 Game Theory

17.1 Basics

Game theory is the study of behavior of individuals in a strategic scenario, where a strategic scenario is defined as one where the actions of one individual affects the payoff or utility of other individuals. In the previous section we introduced game theory in the context of firm competition. In this section, we will generalize the methods used above and introduce some specific language. The specification of (static) game consists of three elements:

1. The players, indexed by $i = 1, \ldots, I$. In the duopoly games, for example, the players were the two firms.

2. The strategies available: each player chooses strategy a_i from the available strategy set A_i. We can write $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_I)$ to represent the strategies of the other $I - 1$ players. Then, a strategy profile of all players is defined by $a = (a_1, \ldots, a_I) = (a_i, a_{-i})$. In the Cournot game, the player's strategies were the quantities chose, hence $A_i = \mathbb{R}_+$.

3. The payoffs for each player as a function of the strategies of the players. We use game theory to analyze situations where there is strategic interaction so the payoff function will typically depend on the strategies of other players as well. We write the payoff function for player i as $u_i(a_i, a_{-i})$. The payoff function is the mapping

$$u_i : A_1 \times \cdots \times A_I \longrightarrow \mathbb{R}.$$

Therefore we can define a game the following way:

Definition 28. A game (in normal form) is a triple,

$$\Gamma = \Big\{\{1, 2, \ldots, I\}, \{A_i\}_{i=1}^{I}, \{u_i(\cdot)\}_{i=1}^{I}\Big\}$$

We now define the concept of best response, i.e. the action for a player i which is best for him (in the sense of maximizing the payoff function). But since we are studying a strategic scenario, what is best for player i potentially depends on what others are playing, or what player i believes others might be playing.

Definition 29. An action a_i is a **best response** for player i against a profile of actions of others a_{-i} if

$$u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i}) \quad \forall\, a'_i \in A_i$$

We say that,

$$a_i \in BR_i(a_{-i})$$

Now we define the concept of Nash Equilibrium for a general game.

Definition 30. An action profile

$$a^* = (a_1^*, a_2^*, \ldots, a_I^*)$$

is a **Nash equilibrium** if,

$$\text{for all } i, \quad u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \forall\, a_i \in A_i$$

or, stated otherwise,

$$\text{for all } i, \quad a_i^* \in BR_i(a_{-i}^*)$$

We know that

$$BR_i\colon \times_{j \neq i} A_j \to A_i$$

Now let's define the following function

$$BR\colon \times_{i=1}^{I} A_i \to \times_{i=1}^{I} A_i$$

as

$$BR = (BR_1, BR_2, \ldots, BR_I)$$

Then we can redefine Nash equilibrium as follows:

Definition 31. An action profile

$$a^* = (a_1^*, a_2^*, \ldots, a_I^*)$$

is a **Nash equilibrium** if,

$$a^* \in BR(a^*)$$

17.2 Pure Strategies

We can represent games (at least those with a finite choice set) in normal form. A normal form game consists of the matrix of payoffs for each player from each possible strategy. If there are two players, 1 and 2, then the normal form game consists of a matrix where the (i, j)th entry consists of the tuple (player 1's payoff, player2's payoff) when player 1 plays their ith strategy and player 2 plays their jth strategy. We will now consider the most famous examples of games.

Example 8. (Prisoner's Dilemma) Suppose two suspects, Bob and Rob are arrested for a crime and questioned separately. The police can prove the committed a minor crime, and suspect they have committed a more serious crime but can't prove it. The police offer each suspect that they will let them off for the minor crime if they confess and testify against their partner for the more serious crime. Of course, if the other criminal also confesses the police won't need his testimony but will give him a slightly reduced sentence for cooperating. Each player then has two possible strategies: Stay Quiet (Q) or Confess (C) and they decide simultaneously. We can represent the game with the following payoff matrix:

		Rob	
		Q	C
Bob	Q	3, 3	−1, 4
	C	4, −1	0, 0

Each entry represents (Bob, Rob)'s payoff from each of the two strategies. For example, if Rob stays quiet while Bob confesses Bob's payoff is 4 and Rob's is -1. Notice that both players have what is known as a dominant strategy; they should confess regardless of what the other player has done. If we consider Bob, if Rob is Quiet then confessing gives payoff $4 > 3$, the payoff from staying quiet. If Rob confesses, then Bob should confess since $0 > -1$. The analysis is the same for Rob. So the only stable outcome is for both players to confess. So the only Nash Equilibrium is (Confess, Confess). Notice that, from the perspective of the prisoners this is a bad outcome. In fact it is Pareto dominated by both players staying quiet, which is not a Nash equilibrium.

The above example has a dominant strategy equilibrium, where both players have a unique dominant strategy.

Definition 32. A strategy is a_i is **dominant** if

$$u_i(a_i, a_{-i}) > u_i(a'_i, a_{-i}) \text{ for all } a'_i \in A_i, a_{-i} \in A_{-i}.$$

If each player has a dominant strategy, then the only rational thing for them to do is to play that strategy no matter what the other players do. Hence, if a dominant strategy equilibrium exists it is a relatively uncontroversial prediction of what will happen in the game. However, it is rare that a dominant strategy will exist in most strategic situations. Consequently, the most commonly used solution concept is Nash Equilibrium, which does not require dominant strategies.

Note the difference between Definitions **??** and 32: A Nash Equilibrium is only defined for the best response of the other players, s^*_{-i}, whereas dominant strategies have to hold for strategies $s_{-i} \in S_{-i}$. A strategy profile is a Nash Equilibrium if each player is playing a best response to the other players' strategies. So a Nash Equilibrium is a stable outcome where no player could profitably deviate. Clearly when dominant strategies exist it is a Nash Equilibrium for all players to play a dominant strategy. However, as we see from the Prisoner's Dilemma example the outcome is not necessarily efficient. The next example shows that the Nash Equilibrium may not be unique.

Example 9. (Coordination Game) We could represent a coordination game where Bob and Ann are two researcher both of whose input is necessary for a project. They decide simultaneously whether to do research (R) or not (N).

		Bob	
		R	N
Ann	R	3, 3	−1, 0
	N	0, −1	1, 1

Here (R,R) and (N,N) are both equilibria. Notice that the equilibria in this game are Pareto ranked with both players preferring to coordinate on doing research. Both players not doing research is also an equilibrium, since if both players think the other will play N they will play N as well.

A famous example of a coordination game is from traffic control. It doesn't really matter if everyone drives on the left or right, as long as everyone drives on the same side.

Example 10. Another example of a game is a "beauty contest." Everyone in the class picks a number on the interval $[1, 100]$. The goal is to guess as close as possible to $\frac{2}{3}$ the class average. An equilibrium of this game is for everyone to guess 1. This is in fact the only equilibrium. Since no one can guess more than 100, $\frac{2}{3}$ of the mean cannot be higher than $66\frac{2}{3}$, so all guesses above this are dominated. But since no one will guess more than $66\frac{2}{3}$ the mean cannot be higher than $\frac{2}{3}(66\frac{2}{3}) = 44\frac{4}{9}$, so no one should guess higher than $44\frac{4}{9}$. Repeating this n times no one should guess higher than $\left(\frac{2}{3}\right)^n 100$ and taking $n \longrightarrow \infty$ all players should guess 1. Of course, this isn't necessarily what will happen in practice if people solve the game incorrectly or expect others too. Running this experiment in class the average guess was approximately 12.

17.3 Mixed Strategy

So far we have considered only pure strategies: strategies where the players do not randomize over which action they take. In other words, a pure strategy is a deterministic choice. The following simple example demonstrates that a pure strategy Nash Equilibrium may not always exist.

Example 11. (Matching Pennies) Consider the following payoff matrix:

		Bob H	Bob T
Ann	H	1, −1	−1, 1
	T	−1, 1	1, −1

Here Ann wins if both players play the same strategy, and Bob wins if they play different ones. Clearly there cannot be pure strategy equilibrium, since Bob would have an incentive to deviate whenever they play the same strategy and Ann would have an incentive to deviate if they play the differently. Intuitively, the only equilibrium is to randomize between H and T with probability $\frac{1}{2}$ each.

While the idea of a matching pennies game may seem contrived, it is merely the simplest example of a general class of zero-sum games, where the total payoff of the players is constant regardless of the outcome. Consequently gains for one player can only come from losses of the other. For this reason, zero-sum games will rarely have a pure strategy Nash equilibrium. Examples would be chess, or more relevantly, competition between two candidates or political parties. Cold War power politics between the US and USSR was famously (although probably not accurately) modelled as a zero-sum game. Most economic situations are not zero-sum since resources can be used inefficiently.

Example 12. A slight variation is the game of Rock-Paper-Scissors.

			Bob	
		R	P	S
Ann	R	0, 0	−1, 1	1, −1
	P	1, −1	0, 0	−1, 1
	S	−1, 1	1, −1	0, 0

Definition 33. A *mixed strategy* by player i is a probability distribution $\sigma_i = \left(\sigma_i\left(S_i^1\right), \ldots, \sigma_i\left(S_i^K\right)\right)$ such that

$$\sigma_i\left(s_i^k\right) \geq 0$$

$$\sum_{k=1}^{K} \sigma_i\left(s_i^k\right) = 1.$$

Here we refer to s_i as an action and to σ_i as a strategy, which in this case is a probability distribution over actions. The action space is $S_i = \left\{s_i^1, \ldots, s_i^K\right\}$.

Expected utility from playing action s_i when the other player plays strategy σ_j is

$$u_i\left(s_i, \sigma_j\right) = \sum_{k=1}^{K} \sigma_j\left(s_j^k\right) u_i\left(s_i, s_j^k\right).$$

Example 13. Consider a coordination game (also knows as "battle of the sexes" similar to the one in Example 9 but with different payoffs

			Bob	
			σ_B	$1-\sigma_B$
			O	C
Ann	σ_A	O	1, 2	0, 0
	$1-\sigma_A$	C	0, 0	2, 1

Hence Bob prefers to go to the opera (O) and prefers to go to a cricket match (C), but both players would rather go to an event together than alone. There are two pure strategy Nash Equilibria: (O, O) and (C, C). We cannot make a prediction, which equilibrium the players will pick. Moreover, it could be the case that there is a third Nash Equilibrium, in which the players randomize.

Suppose that Ann plays O with probability σ_A and C with probability $1 - \sigma_A$. Then Bob's expected payoff from playing O is

$$2\sigma_A + 0(1 - \sigma_A) \tag{28}$$

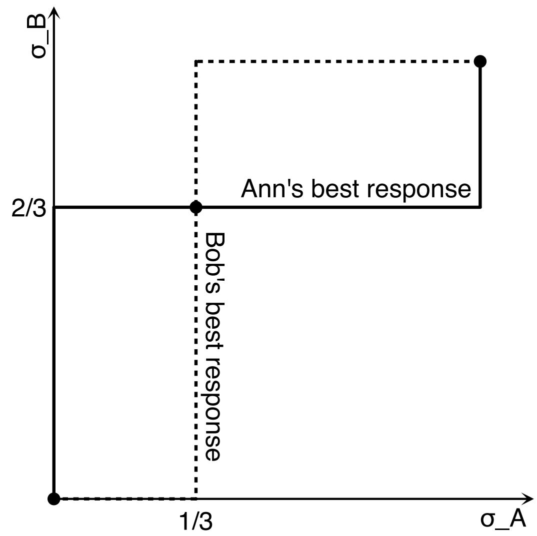

Figure 7: Three Nash Equilibria in battle of the sexes game

and his expected payoff from playing C is

$$0\sigma_A + 1(1 - \sigma_A). \tag{29}$$

Bob is only willing to randomize between his two pure strategies if he gets the same expected payoff from both. Otherwise he would play the pure strategy that yields the highest expected payoff for sure. Equating (28) and (29) we get that

$$\sigma_A^* = \frac{1}{3}.$$

In other words, Ann has to play O with probability $\frac{1}{3}$ to induce Bob to play a mixed strategy as well. We can calculate Bob's mixed strategy similarly to get

$$\sigma_B^* = \frac{2}{3}.$$

Graphically, we can depict Ann's and Bob's best response function in Figure 7. The three Nash Equilibria of this game are the three intersections of the best response functions.

18 Asymmetric Information: Adverse Selection and Moral Hazard

Asymmetric information simply refers to situations where some of the players have relevant information that other players do not. We consider two types of asymmetric information: adverse selection, also known as hidden information, and moral hazard or hidden action.

A leading example for adverse selection occurs in life or health insurance. If an insurance company offers actuarially fair insurance it attracts insurees with above average risk whereas those with below average risk decline the insurance. (This assumes that individuals have private information about their risk.) In other words, individuals select themselves into insurance based on their private information. Since only the higher risks are in the risk pool the insurance company will make a loss. In consequence of this adverse selection the insurance market breaks down. Solutions to this problems include denying or mandating insurance and offering a menu of contracts to let insurees self-select thereby revealing their risk type.

Moral hazard is also present in insurance markets when insurees' actions depend on having insurance. For example, they might exercise less care when being covered by fire or automobile insurance. This undermines the goal of such insurance, which is to provide risk sharing in the case of property loss. With moral hazard, property loss becomes more likely because insurees do not install smoke detectors, for example. Possible solutions to this problem are copayments and punishment for negligence.

18.1 Adverse Selection

The following model goes back to George Akerlof's 1970 paper on "The market for lemons." The used car market is a good example for adverse selection because there is variation in product quality and this variation is observed by sellers, but not by buyers.

Suppose there is a potential buyer and a potential seller for a car. Suppose that the quality of the car is denoted by $\theta \in [0, 1]$. Buyers and sellers have different valuations/willingness to pay v_b and v_s, so that the value of the car is $v_b\theta$ to the buyer and $v_s\theta$ to the seller. Assume that $v_b > v_s$ so that the buyer always values the car more highly then the seller. So we know that trade is always efficient. Suppose that both the buyer and seller know θ, then we have seen in the bilateral trading section that trade can occur at any price $p \in [v_s\theta, v_b\theta]$ and at that price the efficient allocation (buyer gets the car) is realized (the buyer has a net payoff of $v_b\theta - p$ and the seller gets $p - v_s\theta$, and the total surplus is $v_b\theta - v_s\theta$).

The assumption that the buyer knows the quality of the car may be reasonable in some situations (new car), but in many situations the seller will be much better informed about the car's quality. The buyer of a used car can observe the age, mileage, etc. of a car and so have a rough idea as to quality, but the seller has presumably been driving the car and will know more about it. In such a situation we could consider the quality θ as a random variable, where the buyer knows only the distribution but the seller knows the realization. We could consider a situation where the buyer knows the car is of a high quality with some probability, and low quality otherwise, whereas the seller knows whether the car is high quality. Obviously the car could have a more complicated range of potential qualities. If the seller values a high quality car more, then their decision to participate in the market potentially reveals negative information about the quality, hence the term adverse selection. This is because if the car had higher quality the seller would be less willing to sell it at any given price. How does this type of asymmetric information change the outcome?

Suppose instead that the buyer only knows that $\theta \sim U[0,1]$. That is that the quality is uniformly distributed between 0 and 1. The the seller is willing to trade if

$$p - v_s\theta \geq 0 \tag{30}$$

and the buyer, who does not know θ, but forms its expected value, is willing to trade if

$$E[\theta]v_b - p \geq 0. \tag{31}$$

However, the buyer can infer the car's quality from the price the seller is asking. Using condition (30), the buyer knows that

$$\theta \leq \frac{p}{v_s}$$

so that condition (31) becomes

$$E\left[\theta \,\middle|\, \theta \leq \frac{p}{v_s}\right] v_b - p = \frac{p}{2v_s} v_b - p \geq 0, \tag{32}$$

where we use the conditional expectation of a uniform distribution:

$$E[\theta|\theta \leq a] = \frac{a}{2}.$$

Hence, simplifying condition (32), the buyer is only willing to trade if

$$v_b \geq 2v_s.$$

In other words, the buyer's valuation has to exceed twice the seller's valuation for a trade to take place. If

$$2v_s > v_b > v_s$$

trade is efficient, but does not take place if there is asymmetric information.

In order to reduce the amount of private information the seller can offer a warranty of have a third party certify the car's quality.

If we instead assumed that neither the buyer or the seller know the realization of θ then the high quality cars would not be taken out of the market (sellers cannot condition their actions on information they do not have) and so we could have trade. This indicates that it is not the incompleteness of information that causes the problems, but the asymmetry.

18.2 Moral Hazard

Moral hazard is similar to asymmetric information except that instead of considering hidden information, it deals with hidden action. The distinction between the two concepts can be seen in an insurance example. Those who have pre-existing conditions that make them more risky (that are unknown to the insurer) are more likely, all else being equal, to buy insurance. This is adverse selection. An individual who has purchased insurance may become less cautious since the costs of any damage are covered by insurance company. This is moral hazard. There is a large literature in economics on how to structure incentives to mitigate moral hazard. In the insurance example these incentives often take the form of deductibles and partial insurance, or the threat of higher premiums in response to accidents. Similarly an employer may structure a contract to include a bonus/commission rather then a fixed wage to induce an employee to work hard. Below we consider an example of moral hazard, and show that a high price may signal an ability to commit to providing a high quality product.

Suppose a cook can choose between producing a high quality meal ($q = 1$) and a low quality meal ($q = 0$). Assume that the cost of producing a high quality meal is strictly higher than a low quality meal ($c_1 > c_0 > 0$). For a meal of quality q, and price p the benefit to the customer is $q - p$ and to the cook is $p - c_i$. So the total social welfare is

$$q - p + p - c_i = q - c_i$$

and assume that $1 - c_1 > 0 > -c_0$ so that the high quality meal is socially efficient. We assume that the price is set beforehand, and the cook's choice variable is the quality of the meal. Assume that fraction α of the consumers are repeat clients who are informed about the meal's quality, whereas $1 - \alpha$ of the consumers are uniformed (visitors to the city perhaps) and don't know the meal's quality. The informed customers will only go to the restaurant if the meal is good (assume $p \in (0, 1)$). These informed customers allow us to consider a notion of reputation even though the model is static.

Now consider the decision of the cook as to what quality of meal to produce. If they produce a high quality meal then they sell to the entire market so their profits (per customer) are

$$p - c_1$$

Conversely, by producing the low quality meal, and selling to only $1 - \alpha$ of the market they earn profit

$$(1 - \alpha)(p - c_0)$$

and so the cook will provide the high quality meal if

$$p - c_1 \geq (1 - \alpha)(p - c_0)$$

or

$$\alpha p \geq c_1 - (1 - \alpha)c_0$$

where the LHS is the additional revenue from producing a high quality instead of a low quality meal and the RHS is the associated cost. This corresponds to the case

$$\alpha \geq \frac{c_1 - c_0}{p - c_0}.$$

So the cook will provide the high quality meal if the fraction of the informed consumers is high enough. So informed consumers provide a positive externality on the uninformed, since the informed consumers will monitor the quality of the meal, inducing the chef to make a good meal.

Finally notice that price signals quality here: the higher the price the smaller the fraction of informed consumers necessary ensure the high quality meal. If the price is low ($p \approx c_1$) then the cook knows he will lose $p - c_1$ from each informed consumer by producing a low quality meal instead, but gains $c_1 - c_0$ from each uninformed consumer (since the cost is lower). So only if almost every consumer is informed will the cook have an incentive to produce the good meal. As p increases so does $p - c_1$, so the more is lost for each meal not sold to an informed consumer, and hence the lower the fraction of informed consumers necessary to ensure that the good meal will be provided. An uninformed consumer, who also may not know α, could then consider a high price a signal of high quality since it is more likely that the fraction of informed consumers is high enough to support the good meal the higher the price.

18.3 Second Degree Price Discrimination

In Section **??** we considered first and third degree price discrimination where the seller can identify the type of potential buyers. In contrast, second degree price discrimination occurs when the firm cannot observe to consumer's willingness to pay directly. Consequently they elicit these preferences by offering different quantities or qualities at different prices. The consumer's type is revealed through which option they choose. This is known as screening.

Suppose there are two types of consumers. One with high valuation of the good θ_h, and one with low valuation θ_l. θ is also called the buyers' marginal willingness to pay. It tells us how much a buyer what be willing to pay for an additional unit of the good. Each buyer's type is his private information. That means the seller does not know ex ante what type a buyer he is facing is. Let α denote the fraction of consumers who have the high valuation. Suppose that the firm can produce a product of quality q at cost $c(q)$ and assume that $c'(q) > 0$ and $c''(q) > 0$.

First, we consider the efficient or first best solution, i.e., the case where the firm can observe the buyers' types. If the firm knew the type of each consumer they could offer a different quality to each consumer. The condition for a consumer of type $i = h, l$ buying an object of quality q for price p voluntarily is

$$\theta_i q - p(q) \geq 0$$

and for the firm to participate in the trade we need

$$p(q) - c(q) \geq 0.$$

Hence maximizing joint payoff is equivalent to

$$\max_q \theta_i q - p(q) + p(q) - c(q)$$

or

$$\max_q \theta_i q - c(q).$$

The FOC for each quality level is

$$\theta_i - c'(q) = 0,$$

from which we can calculate the optimal level of quality for each type, $q^*(\theta_i)$. Since marginal cost is increasing by assumption we get that

$$q^*(\theta_l) < q^*(\theta_h),$$

i.e., the firm offers a higher quality to buyers who have a higher willingness to pay in the first best case. In the case of complete information we are back to first degree price discrimination and the firm sets the following prices to extract the entire gross utility from both types of buyers:

$$p_h^* = \theta_h q^*(\theta_h) \quad \text{and} \quad p_h^* = \theta_l q^*(\theta_l)$$

so that buyers' net utility is zero. In Figure 8, the buyers' gross utility, which is equal to the price charged, is indicated by the rectangles $\theta_i q_i^*$.

In many situations, the firm will not be able to observe the valuation/willingness to pay of the consumers. That is, the buyers' type is their private information. In

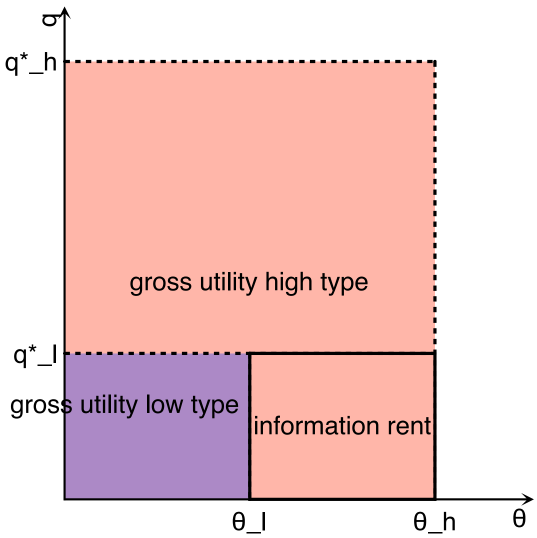

Figure 8: Price discrimination when types are known to the firm

such a situation the firm offers a schedule of price-quality pairs and lets the consumers self-select into contracts. Thereby, the consumers reveal their type. Since there are two types of consumers the firm will offer two different quality levels, one for the high valuation consumers and one for the low valuation consumers. Hence there will be a choice of two contracts (p_h, q_h) and (p_l, q_l) (also called a menu of choices). The firm wants high valuation consumers to buy the first contract and low valuation consumers to buy the second contract. Does buyers' private information matter, i.e., do buyers just buy the first best contract intended for them? High type buyers get zero net utility from buying the high quality contract, but positive net utility of $\theta_h q^*(\theta_l) - p_l > 0$. Hence, high type consumers have an incentive to pose as low quality consumers and buy the contract intended for the low type. This is indicated in Figure 8 as "information rent," i.e., an increase in high type buyers' net utility due to asymmetric information.

The firm, not knowing the consumers' type, however, can make the low quality bundle less attractive to high type buyers by decreasing q_l or make the high quality contract more attractive by increasing q_h or decreasing p_h. The firm's profit maximization problem now becomes

$$\max_{p_h, p_l, q_h, q_l} \alpha \left(p_h - c(q_h)\right) + (1 - \alpha) \left(p_l - c(q_l)\right). \tag{33}$$

There are two type of constraints. The consumers have the option of walking away, so the firm cannot demand payment higher than the value of the object. That is, we must have

$$\begin{aligned} \theta_h q_h - p_h &\geq 0 && (34) \\ \theta_l q_l - p_l &\geq 0. && (35) \end{aligned}$$

These are known as the individual rational (IR) or participation constraints that guarantee that the consumers are willing to participate in the trade. The other type of constraints are the self-selection or incentive compatibility (IC) constraints

$$\begin{aligned} \theta_h q_h - p_h &\geq \theta_h q_l - p_l && (36) \\ \theta_l q_l - p_l &\geq \theta_l q_h - p_h, && (37) \end{aligned}$$

which state that each consumer type prefers the menu choice intended for him to the other contract. Not all of these four constraints can be binding, because that would determine the optimal solution of prices and quality levels. The IC for low type (37) will not binding because low types have no incentive to pretend to be high types: they would pay a high price for quality they do not value highly. On the other hand high type consumers' IR (34) will not be binding either because we argued above that the firm has to incentivize them to pick the high quality contract. This leaves constraints (35) and (36) as binding and we can solve for the optimal prices

$$p_l = \theta_l q_l$$

using constraint (35) and

$$p_h = \theta_h(q_h - q_l) + \theta_l q_l$$

using constraints (35) and (36). Substituting the prices into the profit function (33) yields

$$\max_{q_h, q_l} \alpha \left[\theta_h(q_h - q_l) + \theta_l q_l - c(q_h)\right] + (1 - \alpha)\left(\theta_l q_l - c(q_l)\right).$$

The FOC for q_h is simply

$$\alpha\left(\theta_h - c'(q_h)\right) = 0,$$

which is identical to the FOC in the first best case. Hence, the firm offers the high type buyers their first best quality level $q_R^*(\theta_h) = q^*(\theta_h)$. The FOC for q_l is

$$\alpha(\theta_l - \theta_h) + (1 - \alpha)\left(\theta_l - c'(q_l)\right) = 0,$$

which can be rewritten as

$$\theta_l - c'(q_l) - \frac{\alpha}{1 - \alpha}(\theta_l - \theta_h) = 0.$$

The third term on the LHS, which is positive, is an additional cost that arises because the firm has to make the low quality contract less attractive for high type buyers. Because of this additional cost we get that $q_R^*(\theta_l) < q^*(\theta_l)$: the the quality level for low types is lower than in the first best situation. This is depicted in Figure . The low type consumers' gross utility and the high type buyers' information rent are decreased, but The optimal level of quality offered to low type buyers is decreasing in the fraction of high type consumer α:

$$\frac{dq_R^*(\theta_l)}{d\alpha} < 0$$

since the more high types there are the more the firm has to make the low quality contract unattractive to them.

This analysis indicates some important results about second degree price discrimination:

1. The low type receives no surplus.
2. The high type receives a positive surplus of $q_l(\theta_h - \theta_l)$. This is known as an information rent, that the consumer can extract because the seller does not know his type.
3. The firm should set the efficient quality for the high valuation type.
4. The firm will degrade the quality for the low type in order to lower the rents the high type consumers can extract.

19 Auctions

Auctions are an important application of games of incomplete information. There are many markets where goods are allocated by auctions. Besides obvious examples such as auctions of antique furniture there are many recent application. A leading example is Google's sponsored search auctions. Google matches advertiser to readers of websites and auctions advertising space according to complicated rules.

Consider a standard auction with I bidders, and each bidder i from 1 to I has a valuation v_i for a single object which is sold by the seller or auctioneer. If the bidder wins the object at price p_i then he receives utility $v_i - p_i$. Losing bidders receive a payoff or zero. The valuation is often the bidder's private information so that we have to analyze the uncertainty inherent in such auctions. This uncertainty is captured by modelling the bidders' valuations as draws from a random distribution:

$$v_i \sim F(v_i).$$

We assume that bidders are symmetric, i.e., their valuations come from the same distribution, and we let b_i denote the bid of player i.

There are many possible rules for auctions. They can be either sealed bid or open bid. Examples of sealed bid auctions are the first price auction (where the winner is the bidder with the highest bid and they pay their bid), and the second price auction (where the bidder with the highest bid wins the object and pays the second highest bid as a price). Open bid auctions include English auctions (the auctioneer sets a low price and keeps increasing the price until all but one player has dropped out) and the Dutch auction (a high price is set and the price is gradually lowered until someone accepts the offered price). Another type of auction is the Japanese button auction, which resembles an open bid ascending auction, but every time the price is raised all bidders have to signal their willingness to increase their bid. Sometimes, bidders hold down a button as long as they want to increase their bid and release when they want to exit the auction.

Let's think about the optimal bidding strategy in a Japanese button auction, denotes by $b_i(v_i) = t_i$, where $t_i = p_i$ is the price the winning bidder pays for the good. At any time, the distribution of valuations, F, the number of remaining bidders are known to all players. As long as the price has not reached a bidder's valuation it is optimal for him to keep the button pressed because he gets a positive payoff if all other players exit before the price reaches his valuation. In particular, the bidder with the highest valuation will wait longest and therefore receive the good. He will only have to pay the second highest bidder's valuation, however, because he should release the button as soon as he is the only one left. At that time the price will have exactly reached the second highest valuation. Hence, it is optimal for all bidders to bid their true valuation. If the price exceeds v_i they release the button and get 0 and the highest valuation bidder gets a positive payoff. In other words, the optimal strategy is

$$b_i^*(v_i) = v_i.$$

What if the button auction is played as a descending auction instead? Then it is no longer optimal to bid one's own valuation. Instead, $b_i^*(v_i) < v_i$ because only waiting until the price reaches one's own valuation would mean that there might be a missed chance to get a strictly positive payoff.

In many situations (specifically when the other players' valuations does not affect your valuation) the optimal behavior in a second price auction is equivalent to an English auction, and the optimal behaviour in a first price auction is equivalent to a Dutch auction. This provides a motivation for considering the second price auction which is strategically very simple, since the English auction is commonly used. It's the mechanism used in the auction houses, and is a good first approximation how auctions are run on eBay.

How should people bid in a second price auction? Typically a given bidder will not know the bids/valuations of the other bidders. A nice feature of the second price auction is that the optimal strategy is very simple and does not depend on this information: each bidder should bid their true valuation.

Proposition 3. *In a second price auction it is a Nash Equilibrium for all players to bid their valuations. That is $b_i^* = v_i$ for all i is a Nash Equilibrium.*

Proof. Without loss of generality, we can assume that player 1 has the highest valuation. That is, we can assume $v_1 = \max_i\{v_i\}$. Similarly, we can assume without loss of generality that the second highest valuation is $v_2 = \max_{i>1}\{v_i\}$. Define

$$\mu_i(v_i, b_i, b_{-i}) = \left\{ \begin{array}{c} v_i - p_i, \text{ if } b_1 = \max_j\{b_j\} \\ 0, \text{ otherwise} \end{array} \right\}$$

to be the surplus generated from the auction for each player i. Then under the given strategies ($b = v$)

$$\mu_i(v_i, v_i, v_{-i}) = \left\{ \begin{array}{c} v_1 - v_2, \ i = 1 \\ 0, \text{ otherwise} \end{array} \right\}$$

So we want to show that no bidder has an incentive to deviate.

First we consider player 1. The payoff from bidding b_1 is

$$\mu_1(v_1, b_1, v_{-1}) = \left\{ \begin{array}{c} v_1 - v_2, \text{ if } b_1 > v_2 \\ 0, \text{ otherwise} \end{array} \right\} \leq v_1 - v_2 = \mu_1(v_1, v_1, v_{-1})$$

so player 1 cannot benefit from deviating.

Now consider any other player $i > 1$. They win the object only if they bid more than v_1 and would pay v_1. So the payoff from bidding b_1 is

$$\mu_i(v_i, b_i, v_{-i}) = \left\{ \begin{array}{c} v_i - v_1, \text{ if } b_i > v_1 \\ 0, \text{ otherwise} \end{array} \right\} \leq 0 = \mu_i(v_i, v_i, v_{-i})$$

since $v_i - v_1 \leq 0$. So player i has no incentive to deviate either.

We have thus verified that all players are choosing a best response, and so the strategies are a Nash Equilibrium. □

Note that this allocation is efficient. The bidder with the highest valuation gets the good.

Finally, we consider a first price sealed bid auction. There, we will see that it is optimal for bidders to bid below their valuation, $b_i^*(v_i) < v_i$, a strategy called bid shedding. Bidder i's expected payoff is

$$\max_{b_i} (v_i - b_i) \Pr(b_i > b_j \text{ for all } j \neq i) + 0 \Pr(b_i < \max\{b_j\} \text{ for all } j \neq i). \quad (38)$$

Consider the bidding strategy

$$b_i(v_i) = cv_i$$

i.e., bidders bid a fraction of their true valuation. Then, if all players play this strategy,

$$\Pr(b_i > b_j) = \Pr(b_i > cv_j) = \Pr\left(v_j < \frac{b_i}{c}\right). \tag{39}$$

With valuations having a uniform distribution on $[0, 1]$, (39) becomes

$$\Pr\left(v_j < \frac{b_i}{c}\right) = \frac{b_i}{c}$$

and (38) becomes

$$\max_{b_i} (v_i - b_i) \frac{b_i}{c} + 0$$

with FOC

$$\frac{v_i - 2b_i}{c} = 0$$

or

$$b_i^* = \frac{v_i}{2}.$$

Hence, we have verified that the optimal strategy is to bid a fraction of one's valuation, in particular, $c = \frac{1}{2}$.