Consumer Theory

Mark Dean

Lecture Notes for Fall 2009 Introductory Microeconomics - Brown University

1 Introduction

In this section of the course we will examine the standard methods that economists use to model the behavior of consumers. By a 'consumer' we mean a person who has the opportunity to buy various different commodities at fixed market prices. The big question we are going to ask is "How do consumers choose what to buy given their income and the prices in the economy

2 The Optimization Problem

As we shall see, we are going to assume that a consumer has a well defined set of desires, or 'preferences', which can be represented by a numerical utility function. Furthermore, we will assume that the consumer chooses optimally, in the sense that they choose the option with the highest utility of those available to them. This means that a consumer is solving an **optimization problem.** This is an important class of problems that crop up time and again throughout economics (for example, firms are also assumed to solve optimization problems). It is therefore worth outlining what such problems look like in general terms.

An optimization problem has three key components:

1. **The objects of choice.** What is it that is being chosen? In the case of our consumer, this will be the different bundles of goods that they can buy

2. **The objective function** What is it that the chooser is trying to maximize? In the case of

the consumer this is the utility function

3. **Constraints.** What constraints are there on the choices that can be made? In the case of the consumer, this is the set of goods that they can afford

Optimization problems can always be written in the following form:

Choose <*some object*> **in order to maximize** <*some objective function*> **subject to** <*some constraint*>

One of the most important skills to learn in economics is being able to formulate optimization problems. Here are a couple of examples:

Example 1 *A Brown student is trying to decide what courses they are going to take. They want to get the highest possible grade point average, but they also want to do the economics concentration.*

Choose *courses offered by Brown* ***in order to maximize*** *grade point average* ***subject to*** *satisfying the requirements of the economics concentration*

Example 2 *The British government is signed up to reduce emissions by 25%. However, they want to do so in a way that minimizes the damage to the economy*

Choose *fiscal policy (taxes and subsidies)* ***in order to maximize*** *economic output* ***subject to*** *reducing greenhouse emissions by 25%*

The ability to be able to translate problems into this form is one of the key skills you will need in this course.

We are now going to set up the problem of a consumer in the form of an optimization problem. We are first going to describe what the consumer gets to choose, and the constraints that they face, before discussing what it is that they want to optimize.

3 Consumption Bundles and Budget Constraints

We are going to concentrate on a very simple model of the world in which there are only two possible types of object that our consumer wants to buy. We will call them goods 1 and 2, though you can think of them as any two goods you want. For some reason some economics textbooks think of them as guns and butter. We therefore have the first part of the optimization problem: The objects of choice are the amount of these two good that the consumer wants to buy. We will call these x_1 and x_2. We will call (x_1, x_2) a **consumption bundle.** We will call the set of all feasible bundles the **commodity space.** We can illustrate the commodity space graphically with good 1 on the horizontal axis and good 2 on the vertical axis (see figure 1).

The world we are thinking of has three *parameters* - or things that the consumer cannot control: the price of the two goods (which we call p_1 and p_2) and the amount of money that the consumer has to spend (which we will call M).

The constraints that consumer is operating under should now be obvious. They can only choose bundles of goods that they can afford. This is called the **budget constraint**, and can be written as follows

$$p_1x_1 + p_2x_2 \leq M$$

We can now illustrate the budget constraint in the graph of commodity space. In order to do so we want to graph the line for which the budget constraint holds with equality, which is called the **budget line.**

$$p_1x_1 + p_2x_2 = M$$

Rearranging this equation gives

$$\begin{aligned} p_2x_2 &= M - p_1x_1 \\ x_2 &= \frac{M}{p_2} - \frac{p_1}{p_2}x_1 \end{aligned}$$

The budget line is therefore a straight line in the commodity space, with the slope $-\frac{p_1}{p_2}$. This is the rate at which the consumer can exchange one good for another: If they give up 1 unit of good 2 they will get $\frac{p_1}{p_2}$ units of good 2. The **budget set**, or set of objects that the consumer can afford, are the objects on the 'inside' of the budget line. This is shown in figure 2.

How do changes in the parameters (prices and income) change the budget set? First, think about an increase in price 2. This is going to flatten the slope of the budget line (the rate at which good 1 can be traded for good 2). However, it is not going to change the amount of good 1 that can be bought if all income is spent on that good. Thus, the budget line *pivots* round the point where it touches the horizontal axis (figure 3). Similarly, an increase in price 2 causes the budget line to steepen, and pivot round the point where it touches the vertical axis. Finally, a fall in income doesn't change the slope of the budget line, but does cause it to move inwards - a parallel shift.

One thing that is apparent from the equation of the budget line is that the budget set is invariant to (balanced) inflation: To see this imagine all the prices are income increase by a factor λ. What does this do to the budget line?

$$\begin{aligned} x_2 &= \frac{\lambda M}{\lambda p_2} - \frac{\lambda p_1}{\lambda p_2} x_1 \\ &= \frac{M}{p_2} - \frac{p_1}{p_2} x_1 \end{aligned}$$

So (for example) doubling all the prices *and income* has no effect on what a consumer can buy. This means that one of the prices in the system is *redundant* - we can set one price to whatever we want (for example 1), and then describe the other prices relative to that price. This is called a normalization. So for example, we can set income to 1, then write the other prices as

$$\frac{p_1}{M} x_1 + \frac{p_2}{M} x_2 \leq 1$$

4 Preferences and Indifference Curves

So far we have described two parts of the optimization problem in our simple model: we know what it is that our consumer gets to choose (consumption bundles) and we know what constraints they face (the budget set). We now move onto the final piece of the jigsaw - the objective function, or what it is that consumers maximize.

A moment's reflection will lead you to see that this is by far the most difficult and controversial part of setting up the optimization problem. Human beings are incredibly complex beasts, with all sorts of goals and desires. Moreover, no two people are alike in what they want, so how on earth are we going to come up with a simple, parsimonious mathematical function that captures all of this?

The answer to this is twofold. First, we have simplified the world so that there are only two goods in it. Thus we do not need to understand people's overall psychology, just how they feel about different bundles of these two goods. Second, we are going to make relatively minimal assumptions about what people like or dislike. It is going to turn out that even very minimal assumptions are going to be able to get us quite a long way. Though as ever, there is a trade off between how much we assume, and how much we can predict: the more we assume, the more specific our predictions will be, but the more vulnerable we are to being wrong because our assumptions are not valid.

4.1 Preferences

The starting point of the economic analysis of peoples objectives is the idea of **preferences.**[1] Consider the following question:

> Question: Do you prefer bundle (x_1, x_2) to bundle (y_1, y_2)?
>
> Answers:
>
> (1) I prefer bundle (x_1, x_2) to bundle (y_1, y_2)
>
> (2) I prefer bundle (y_1, y_2) to bundle (x_1, x_2)
>
> (3) I am indifferent between (x_1, x_2) and (y_1, y_2)

[1] For an excellent introduction to this topic, see the first chapter of "Lecture Notes in Economic Microeconomic Theory" by Ariel Rubinstein - it is available free on his website.

We are going to call the answer to such a question a 'preference', and we are going to represent the three answers with the following symbols

1. $(x_1, x_2) \succ (y_1, y_2)$
2. $(y_1, y_2) \succ (x_1, x_2)$
3. $(x_1, x_2) \sim (y_1, y_2)$

These look like the standard mathematical symbols for greater than (>) and equals (=), but are importantly not. That is how we know we are talking about preferences.

The standard assumption in economics is that the consumer has a **preference relation** on the commodity space. By a preference relation we mean a set of preferences that satisfy the following two properties:

1. **Completeness:** For any two objects (x_1, x_2) and (y_1, y_2) in the commodity space, then the consumer can answer one, and exactly one of (1), (2) or (3). In other words, for any two bundles, one and exactly one of the following is true: $(x_1, x_2) \succ (y_1, y_2)$ or $(y_1, y_2) \succ (x_1, x_2)$ or $(x_1, x_2) \sim (y_1, y_2)$
2. **Transitivity:** For any three objects (x_1, x_2), (y_1, y_2) and (z_1, z_2) in the commodity space then
   (a) If $(x_1, x_2) \succ (y_1, y_2)$ and $(y_1, y_2) \succ (z_1, z_2)$ then $(x_1, x_2) \succ (z_1, z_2)$
   (b) If $(x_1, x_2) \sim (y_1, y_2)$ and $(y_1, y_2) \sim (z_1, z_2)$ then $(x_1, x_2) \sim (z_1, z_2)$

It is going to turn out that the existence of a preference relation will be very useful to us in terms of modelling what it is that our consumer likes. This is because we will be able to represent preferences that satisfy completeness and transitivity (and another, technical assumption that we will come back to) with a numerical **utility function.** However, first, it is worth asking whether it is a good assumption that our consumer has a preference relation. In other words, are these good assumptions?

On the face of it, the assumption of completeness may seem to be pretty uncontroversial And in simple settings it is: If I ask you to compare an apple to an orange, or a Ferrari to a Chrysler, or a Miley Cyrus CD to almost anything, then you should be fine giving one and only one answer to the preference question. However, note that we have ruled out the following answers:

(d) I don't know what (x_1, x_2) is.

(e) I can't decide

(f) Sometimes I prefer (x_1, x_2) and sometimes I prefer (y_1, y_2)

All of which may be valid answers in some cases. Some example of cases in which people may find it hard to answer (a), (b) or (c) are the following:

- When alternatives are good along different dimensions (for example a very sporty car that is not very safe to a much safer car that is not very exciting)
- Very emotional choices: Would you prefer your first or second born child to be killed?
- When objects are very complicated: If you are feeling sick, would you rather take Bismuth subsalicylate or 8-methyl-N-vanillyl-6-nonenamide?
- When you are addicted to a substance: Do you always say you prefer smoking to not smoking?

The transitivity axiom can also seem to be pretty reasonable. It would seem strange for someone to say "I prefer steak to cheeseburgers" and "I prefer cheeseburgers to hamburgers" but then claim to prefer hamburgers to steak. Another reason sometimes given to think that transitivity is sensible property is that people whose preferences do not satisfy transitivity may be liable to a 'money pump'. Imagine someone who did indeed prefer steak to cheeseburgers, cheeseburgers to hamburgers and hamburgers to steak., and say that they currently had a hamburger. Presumably you could offer to sell such a person a cheeseburger in exchange for the hamburger and a small amount of money (say 1 cent). Now that they have a cheeseburger, you could sell them a steak in exchange for the cheeseburger and another cent. However, because we know that they prefer hamburger to steak, we could now sell them their original hamburger in exchange for the steak and another cent. The intransitive consumer therefore ends up with the same thing they started with, but is now 3 cents worse off.[2]

[2]Whether the money pump is a good argument for tranitive preferences is a hot topic in philosophy. If you are interested, read "Money Pumps and Diachronic Books" by Isaac Levi in Philosophy of Science, 69 (September 2002) pp. S235–S247

In fact, in many cases people do exhibit intransitive preferences including you. Before the class, you were asked to log on to a website and express preferences between different holidays, described in the form "a holiday in Rome in a 5 star hotel with Zagat food rating 17 for $654". Of the 27 people who did so, the average number of intransitive preferences was 7.8 (out of a possible 70). Only 4 people had no intransitivities. Many of the violations of transitivity involved two alternatives that were actually the same, but differed in the order in which the characteristics appeared in the description: "A weekend in Paris at a 4-star hotel with food quality Zagat 17 for $574," and "A weekend in Paris for $574 with food quality Zagat 17 at a 4-star hotel." All students expressed indifference between the two alternatives, but in a comparison of these two alternatives to a third alternative "A weekend in Rome at a 5-star hotel with food quality Zagat 18 for $612" a quarter of the students gave responses that violated transitivity.

One classic reason for failure of intransitivity is when we are asked to compare things that are very similar. Say, for example that you prefer having sugar in your coffee to not having sugar. However, if I add a single extra grain of sugar to your coffee, you are not going to be able to tell the difference. Thus, if I ask you the question 'do you prefer your coffee with no grains of sugar to one grain of sugar, you will say that you are indifferent. Similarly, if I ask 'do you prefer coffee with one grain or coffee with two grains', you will be indifferent and so on. Yet, if I ask if you prefer 0 grains or 40,000 grains (about a teaspoon) you will say that you prefer the latter. This is a failure of transitivity.

A fun failure of transitivity in the sensory realm is Shepard's scale, which you can listen to at http://asa.aip.org/demo27.html. This is a sequence of tones which sound as if each note is higher than the last, yet in fact it is just the same 10 notes repeated time and again.

4.2 Preferences and Choice

The main reason that we are interested in preferences is that we want to use them to model the way that people make choices: We want to use preferences (and the related concept of utility) to fill in the missing piece of the optimization problem, by acting as the objective function.:

> **Choose:** *a consumption bundle*
> **In order to maximize:** *preferences (or utility)*
> **Subject to:** *the budget constraint.*

A natural question to ask is whether this is a good assumption: do people choose in order to maximize their preferences? Is there any way to test whether this is the case?

In fact, the problem is more difficult than that, as we do not in general get to observe preferences directly - we do not get to see the answer to the preference question we posed above. Instead, economists try to answer a slightly different question: Do consumers look like they are maximizing **any** preference ordering? Put another way, can we find **any** definition of preferences such that our consumer seem to be making choices according to the preferences? We say that such preferences **rationalize** a particular set of choices.

The answer to this is 'it depends' - some patterns of choice can be rationalized, while others cannot. Here is an example of some choice behavior that seems difficult to rationalize with any set of preferences

Example 3 *When the consumer is offered a choice of steak or cheeseburger they choose the cheeseburger. When they are offered a choice of steak, hamburger of cheeseburger they choose the steak.*

Why is this problematic? You might argue that, by choosing the cheeseburger over the steak, then the consumer has **revealed a preference** for the cheeseburger over the steak. So what should this person do when offered a choice between cheeseburger, steak and hamburger? Well, they could well choose the hamburger if they like the hamburger more than the cheeseburger, or they might choose the cheeseburger, if they like the cheeseburger more than the hamburger. What they should *not* choose is the steak, as they have already demonstrated that they prefer the cheeseburger over the steak.

This intuition is captured more formally in a property called the **independence of irrelevant alternatives:**

Definition 1 *Say that we see a consumer choose an option x from a set of alternatives A. We then add some alternatives B and ask the consumer to choose again. The independence of irrelevant alternatives (which I will abbreviate sometimes to IIA) states that the consumer must choose either x or something from B.*

For us to be able to rationalize a particular set of choices, we in fact need two things to be true:

1. Choices are **well behaved:** The consumer always chooses the same thing when offered the same set of alternatives, and they always choose one (and only one) alternative from those available
2. The independence of irrelevant alternatives holds

Under these circumstances, we can find a set of preferences $\succ$ that rationalize the consumers choices - that is the item that is chosen from any set of alternatives is preferred according to $\succ$ to all the other alternatives in that set.

In order to 'prove' this statement, we can construct a candidate set of preferences, and show that it does the job. This candidate set of preferences is going to be very natural:. We will say that a commodity bundle (x_1, x_2) is preferred to a bundle (y_1, y_2) if it is chosen over that bundle in a direct choice. We will use the symbol $\bar{\succ}$ to indicate this is the case: in other words we will write $(x_1, x_2)\bar{\succ}(y_1, y_2)$ if (x_1, x_2) is chosen over (y_1, y_2) when only those two alternatives were available.

In order for me to show you that $\bar{\succ}$ rationalizes the choices of this consumer, I need to show you three things

1. $\bar{\succ}$ **is complete**. This follows directly from the assumption that peoples choices are *well behaved*. This means that, for any two consumption bundles (x_1, x_2) and (y_1, y_2) it must be the case that when offered the choice between the two, the consumer either always chooses (x_1, x_2) or always chooses (y_1, y_2), so, by definition of $\bar{\succ}$, exactly one of $(x_1, x_2)\bar{\succ}(y_1, y_2)$ or $(y_1, y_2)\bar{\succ}(x_1, x_2)$ is true
2. $\bar{\succ}$ **is transitive.** I am going to argue that, if $\bar{\succ}$ is intransitive, then the consumer must have violated the independence of irrelevant alternatives. This is enough to show that, if the consumer has *not* violated the independence of irrelevant alternatives, then $\bar{\succ}$ is not intransitive.

   Imagine that $\bar{\succ}$ is intransitive. This means that we can find bundles such that $(x_1, x_2)\bar{\succ}(y_1, y_2)$, $(y_1, y_2)\bar{\succ}(z_1, z_2)$ but $(z_1, z_2)\bar{\succ}(x_1, x_2)$. In other words, the consumers choices were as follows

   From (x_1, x_2) and (y_1, y_2) they chose (x_1, x_2)

   From (z_1, z_2) and (y_1, y_2) they chose (y_1, y_2)

   From (x_1, x_2) and (z_1, z_2) they chose (z_1, z_2)

Now let's think about their choice when they are offered all three alternatives (x_1, x_2), (y_1, y_2) and (z_1, z_2). I claim that **any** choice they make from this set will violate the independence of irrelevant alternatives. Say they chose (x_1, x_2). But we already know that the chose (z_1, z_2) from (x_1, x_2) and (z_1, z_2), so that is a violation of IIA. Similarly, if (y_1, y_2) is chosen, this violates IIA relative to the choice of (x_1, x_2) from (x_1, x_2) and (y_1, y_2) and the choice of (z_1, z_2) violates IIA relative to the choice of (y_1, y_2) from (z_1, z_2) and (y_1, y_2) Thus any consumer for whom $\bar{\succ}$ is intransitive will have violated IIA

3. $\bar{\succ}$ **rationalizes the consumer's choices**. In other words, I have to show that the consumer always chooses the bundle that is most preferred from any set of alternatives. Again, I am going to show this by demonstrating that a consumer for whom this is not true must have violated IIA. To see this, imagine that a consumer chooses a bundle (x_1, x_2) from some set of alternatives, but in that set there is an alternative (y_1, y_2) such that $(y_1, y_2) \bar{\succ} (x_1, x_2)$. But, by definition, this means that (y_1, y_2) was chosen in a direct choice over (x_1, x_2), so the fact that the consumer is now choosing (x_1, x_2) violates IIA.

The smart ones amongst you may have noticed that the concept of indifference has disappeared from our discussion here. This is because the above discussion can get a lot more complicated if one allows for indifference. For example, you may be able to think of a way that we can rationalize ANY set of choices by assuming that the consumer is indifferent between all outcomes. We will come back to these issues later.

So do people satisfy the independence of irrelevant alternatives? The answer is, as is so often the case with economic, 'yes and no'. One recent study,[3] gives relatively encouraging results. In this experiment, the authors ask subjects to make 50 choices from different budget sets, which are presented to subjects in the form of graphs like those we have already seen (an example of their experimental interface is shown in figure 6). They found that, on average, people's choices were close to satisfying the necessary rationality property.[4] Over 30% of people had no violations of rationality (from 50 choices), over 60% had one or fewer violations and 80% have 2 or fewer

[3] *Consistency and Heterogeneity of Individual Behavior under Uncertainty* by Syngjoo Choi, Douglas Gale, Ray Fisman and Shachar kariv American Economic Review, December 2007, 97(5), pp. 1921-1938. Available on the web at http://emlab.berkeley.edu/~kariv/CFGK_III.pdf

[4] This experiment and the following ones actually test a stronger rationatiality property - the Generalized Axiom of Revealed Preference. We will come back to this in time.

violations. By economic standards, that counts as pretty good. There is also some evidence that people become more rational with practice, in particular exposure to 'market like' conditions.

Other experiments have looked at whether specific sub-populations are rational. One famous study[5] looked at how early rationality developed by getting children of different ages to make choices (between bundles containing boxes of juice and potato chips). The study then compares how many violations of rationality were observed at different age groups. This study suggests that rationality does develop reasonably early. The average number of violations for 7 year olds was 4.3, compared to 2.1 for 7 year olds and 2.0 for 21 year olds. If people had chosen randomly we would have expected to see 8.5 violations per person, so all the groups did better than random. The 11 year olds also did better than the 7 year olds, but there is little change between the 11 and 21 year olds.

Another fun study[6] looks at whether rationality is a property that extends beyond humans. Chen and Santos train monkeys to use money in order to buy different foods from the experimenters. They then changed the prices that the monkeys faced in order to see whether they violated the rationality assumptions. They found that, in general, the monkeys were pretty good at obeying IIA and its extensions. However, they did exhibit some of the same violations of rationality that we discuss below

While the picture painted by the above studies is quite rosy for rationality, there are a number of studies that highlight *particular* cases in which people violate the necessary conditions. Two classic examples are

1. **The Endowment Effect:** First demonstrated by Jack Knetsch, along with Daniel Kahneman and Richard Thayler, the endowment effect is a behavioral phenomena by which people seem to like things they already own, simply because they own them. In the original experiment, two groups of subjects were brought into the laboratory. The first group were given a coffee mug when they came in, and were then asked if they wanted to exchange the coffee

[5] *GARP for Kids: On the Development of Rational Choice Behavior* by William T. Harbaugh & Kate Krause & Timothy R. Berry, 2001. American Economic Review, American Economic Association, vol. 91(5), pages 1539-1545, December.

[6] The Evolution of Rational and Irrational *Economic Behavior: Evidence and Insight from a Non-human Primate Species* by Keith Chen and Laurie Santos - chapter 7 of Neuroeconomics: Decision Making and the Brain,edited by Paul Glimcher, Colin Camerer, Ernst Fehr, and Russell Poldrack. Academic Press: Elsevier, 2009

mug for a chocolate bar. In this group 89% of people chose the coffee mug. A second group were given the chocolate bar when they came in, and were asked if they wanted to exchange it for the coffee mug. In this group, only 10% of people chose the coffee mug. Extrapolating, this suggests that there must be people who want the coffee mug when they initially have the coffee mug and want the chocolate when they initially have the coffee mug. Such behavior violates the assumption that choices are 'well behaved' (why?)

2. **The Decoy Effect:** The decoy effect is a phenomena by which adding an inferior option to the choice set can change the choices that people make between superior alternatives - thus violating the independence of irrelevant alternatives. Consider the following example. A consumer is first offered a choice between two MP3 players

	Player A	Player B
Memory	120 gig	60 gig
Cost	$400	$250

In many cases, people may choose the cheaper, smaller MP3 player. Now imagine a choice between 3 players

	Player A	Player B	Player C
Memory	120 gig	60 gig	100 gig
Cost	$400	$250	$450

Player C is clearly inferior to player A, and so no one will ever choose it. Yet its presence in the choice set can cause people to switch their choice from player B to player A. Why? Well, there is some evidence that the fact that Player A is better than at least one other available alternative can make it more attractive to people. This clearly leads to a violation of IIA (why?)

4.3 Preferences and Welfare

We have now derived conditions under which we can represent people's choices by a preference relation, and I have presented some evidence that these conditions may hold reasonably well in some situations. We have actually done something pretty special here - we have managed to come up with a definition of preferences that capture people's behavior without having to say anything about what people like. We haven't had to wrestle with any difficult questions about whether people will like apples more than oranges. or Ferraris more than Chryslers. We haven't even ruled out the possibility of them liking Miley Cyrus. All we have done is observed their choices, and said that if they satisfy a (fairly weak) consistency condition, then we can come up with a preference relation that captures their behavior. This is a great advantage, because if we needed to decide what it is that people like and dislike, or what it is that makes people happy, before we could do any economic analysis, then we wouldn't get very far.

However, something very important has happened between sections 4.1 and 4.2 that we need to remember. In the first case, we were treating preferences as an expression of what people actually **liked**. In fact, we were treating preferences as responses to questions about what people did like. However, in the second case, all we had to go on were the choices people make. Here, if the necessary conditions were satisfied, we can come up with a **revealed preference** relation that describes people choices. However, unless we are prepared to say that what people choose is tautologically the same as what they like, then these concepts do not necessarily have anything to do with each other.

Consider the following example. A consumer has to choose between different brands of washing powder. They have no idea which brand is best, so they always choose the brand with a name that comes first alphabetically. Such a consumer will satisfy the independence of irrelevant alternatives (check that you agree with this), and so we can generate a revealed preference relation that matches their choice. However, it would be odd to say that the consumer is always choosing the washing powder that they like the most.

Does this distinction matter? Well, it depends what we are using our models for. If we only want to use them for **positive** purposes, then we may not care: these preferences explain the choices that people make, and that is all we want to do. However, in many cases, economists will slip into using their models **normatively** - using them to say that a particular policy is 'good', or 'optimal'. When they do so, they are almost always assuming that the optimal thing for an individual is the thing that maximizes their preferences (or utility) - by which they mean the revealed preferences that come from choice. We are going to fall into this trap ourselves.

Is this a valid thing to do? To me, this is a very difficult question, and one to which I have not seen a compelling answer. Is it meaningful to say that you 'like' option A but choose option B? Even if we do think revealed preference is the same as preference, is this what we should be setting policy to maximize? What about drug addicts, or just people that make mistakes in their reasoning? These are questions you will have to make up your own mind about, but you should bear them in mind as soon as I start describing something as 'optimal'.

4.4 Indifference Curves and the Marginal Rate of Substitution

Having spent some time discussing the nature of preferences, and whether there are such things as well behaved, we are now going to proceed on the assumption that there our consumer does have well behaved preferences. In this section we will learn how to describe preferences in the commodity space that we introduced is section 3

Ideally we would like to be able to use three dimensional graphs in order to represent preferences in the commodity space: For any consumption bundle, we would like to use the third, 'vertical' axis to represent preferences, with higher levels used to represent more preferred options. Figure 7 shows an example of such a graph for a particular set of preferences.

However, three dimensional graphs tend to be difficult to work with, so instead we translate the three dimensions into two in the ways that cartographers have for years: by using **isoquants.** An isoquant is a line on a graph that links points that are equal in some way. On a map, contour lines link areas of equal height. Here, we will use **indifference curves** to link areas of equal preference - or linking bundles that are indifferent. Figure 8 translates the preferences of figure 7 to two dimensions.

What can we learn from indifference curves? Well, the first thing we can read is the **marginal rate of substitution (MRS)** of one good for another. This is defined as the rate at which one good can be traded for another while keeping the consumer indifferent. To understand this concept, look first at figure 9. Think first of a consumer that currently has bundle (x_1, x_2). One question we could ask is the following: Imagine we give the consumer some more of good 1, moving them to y_1 - a change that we will call Δx_1. The how much of good 2 would we have to take away from them in order to keep them just indifferent between the two points. We can read off the answer from the indifference curve: The bundle that is indifferent to (x_1, x_2) with y_1 is (y_1, y_2). Thus, to keep the consumer indifferent, we would have to reduce the amount of consumption of good 2 to y_2, a change that we call Δx_2. The ratio of the increase in good 1 to the decrease in good 2 that is needed to keep the consumer indifferent is given by the red chord in figure 9, and is equal to

$\Delta x_2/\Delta x_1$.

Now imagine that we make Δx_1 smaller and smaller, and keep track of the Δx_2 that keeps the consumer indifferent. This is the process shown in figure 10. In the limit, the slope of the chord connecting the indifferent point will be equal to the slope of the indifference curve at (x_1, x_2). It is this slope that we call the marginal rate of substitution at (x_1, x_2). It is the instantaneous rate at which the two goods can be traded off with each other to keep the consumer indifferent.

This process should look familiar to you all from your calculus class. And for the maths jocks amongst you, we can define the MRS in a similar way.

$$\begin{aligned} MRS(x_1, x_2) &= -\lim_{\Delta x_1 \to 0} \frac{\Delta x_2}{\Delta x_1} \\ \text{such that } (x_1, x_2) &\sim (x_1 + \Delta x_1, x_2 + \Delta x_2) \end{aligned}$$

4.5 Types of Preferences

We now move on to use indifference curves to describe different types of preferences. We begin by describing two further assumptions that economists usually make about preferences. The first is that people's preferences are **monotonic.** This basically means that people like more stuff rather than less. If we compare two bundles (x_1, x_2) and (y_1, y_2) such that x_1 is at least as big as y_1 and x_2 is at least as big as y_2 then (x_1, x_2) should be at least as good as (y_1, y_2). We say that preferences are **strictly monotonic** if in addition it is the case that if one of those relations is *strict* (i.e. either x_1 is bigger than y_1 or x_2 is bigger than y_2) then we assume that (x_1, x_2) is actually preferred to (y_1, y_2). Mathematically, we can write

$$\begin{aligned} \text{if } x_1 &\geq y_1 \text{ and } x_2 \geq y_2 \text{ then} \\ (x_1, x_2) &\succ (y_1, y_2) \text{ or } (x_1, x_2) \sim (y_1, y_2) \end{aligned}$$

for monotonicity, and

$$\begin{aligned} \text{if also either } x_1 &> y_1 \text{ or } x_2 > y_2 \text{ then} \\ (x_1, x_2) &\succ (y_1, y_2) \end{aligned}$$

for strict monotonicity.

Monotonicity implies three things about indifference curves

1. They are downward sloping
2. They are cannot cross
3. Moving in a 'North Easterly' direction in the commodity space moves to more preferred indifference curves.

You should check that you understand all three of these properties. Figure 11 gives an example of monotonic indifference curves.

Is monotonicity a good assumption? Well, one thing you might worry about is things that you actively *dislike* (Miley Cyrus CDs), and so would actually like fewer of. But this is not really a problem - we can just redefine the good as 1-(number of Miley Cyrus CDs), and we have monotonicity again. A more serious problem is satiation. I certainly prefer having 1 cheeseburger to 0 cheeseburgers, and 2 cheeseburgers to 1 cheeseburger, but do I really prefer 11 cheeseburgers to 10?. At best, I am probably indifferent (a violation of strict monotonicity) At worst I may actually prefer 10 cheeseburgers to 11, as 11 make me feel sick (a violation of monotonicity). Thus, it is not always the case that monotonicity is a good assumption.

Another property that economists often assume is **convexity.** To understand convexity. Think of the following question. Say you are in a bar, and you currently have 10 drinks lined up in front of you, but only 2 sliders. How many drinks would you give away in order to get another slider? Now say you had 2 drinks and 10 sliders. How many drinks would you now exchange for a slider? In general, we would expect people to give away more drinks in exchange for a slider when they have a lot of drinks relative to the number of sliders. In other words, we would expect the MRS of sliders to drinks to be higher when you have a lot of drinks than when you have a few drinks.

In terms of our indifference curves, convexity implies that curves get flatter as we move along them: When you have a small amount of good 1, you need a lot of good 2 in return for giving some of it up to keep you indifferent. However, when you have a lot of good 1, you need relatively little of good 2 in return for giving up good 1 to keep you indifferent. Figure 12 shows this property of convex indifference curves, while figure 13 gives an example of non-convex indifference curves.

One implication of convexity[7] is that the average of two bundles of goods between which the consumer is indifferent is at least as good as either of the bundles themselves. We can say that

[7] In fact, this is how convexity is defined

preferences are convex if the following is true

$$\begin{aligned} \text{For any two bundles } (x_1, x_2) &\sim (y_1, y_2) \\ &\quad \text{and number } \lambda \text{ between 0 and 1} \\ (\lambda x_1 + (1-\lambda) y_1, \lambda x_2 + (1-\lambda) y_2) &\succ (x_1, x_2) \text{ or} \\ (\lambda x_1 + (1-\lambda) y_1, \lambda x_2 + (1-\lambda) y_2) &\sim (x_1, x_2) \text{ and} \end{aligned}$$

We will say that preferences are *strictly* convex if the average of the two is preferred to either of the goods.

We conclude this section by introducing two types of preferences which act as the two extremes of the set of monotone, convex preferences. The first is what we call **perfect substitutes**. The easiest way to think about perfect substitutes is with an example. Say that the two goods in our commodity space are different types of bottled water: Dasani and Aquafina. Say that all you care about is the total number of bottles of water you have (this is all you should care about, they are both just tap water). What would your preferences look like? Well, how many bottles of Dasani would I have to give you in exchange for a bottle of Aquafina to keep you indifferent? If all you care about is the total number of bottles, then the answer is 1, regardless of how many bottles of Dasani you already have. Thus, your MRS (and thus the slope of the indifference curve) is -1 everywhere. - your indifference curves are straight lines! This is demonstrated in figure 14. Note that these preferences are convex, but not *strictly* convex.

At the other extreme, we have goods we call **perfect compliments.** Again, we can illustrate perfect compliments with an example. Think of good 1 as being left shoes and good two as being right shoes. Say that all you care about is the total number of pairs of shoes (which is again what you should care about, unless you are a shoe fetishist, or have one or three feet). What do your preferences look like? In order to answer this question, lets think about all the points in the commodity space in which you have exactly 2 pairs of shoes. One obvious point is the one where you have 2 left and 2 right shoes. However, it is also when you have 2 left and 3 right shoes, 2 left and 4 right shoes and so on. Also when you have 2 right and 3 left, 2 right and 4 left and so on. If you graph all these points, you will see you get an 'L' shaped graph like the ones shown in figure 15. These are the indifference curves for perfect compliments.

5 Utility Functions

So far we have been representing the preferences of our consumer with the $\succ$ relation. This is all very well and good, but we want to describe the behavior of a consumer who chooses in order to maximize their preferences. It is much easier to think about maximizing a numerical function than it is to think about maximizing a set of binary relations. Thus, we want to know the answer to the following question:

Question: Under what circumstances can we find a numerical function u that **represents** the preferences of our consumer? By represent we mean that

$$u(x_1, x_2) > u(y_1, y_2) \text{ if and only if } (x_1, x_2) \succ (y_1, y_2)$$

Amazingly, (or perhaps not that amazingly). The answer is that we can find such a function if (and only if) the preferences are complete and transitive.[8]

In order to get an idea why this statement is true, I will give you an example of a procedure that will generate such a utility function. Lets say that we have a set of different goods $a, b, c, d...$, and say that our consumer has preferences that are complete and transitive over these objects. Here is the procedure

Step 1 Set the utility of object a to 0.5, so $u(a) = 0.5$

Step 2 Take object b. if $b \sim a$ then set $u(a) = 0.5$. If $b \succ a$ then set $u(b) = 0.75$. If $a \succ b$ then set $u(a) = 0.25$

Step 3 For object c, if it indifferent to any of the objects already assigned utility, then give it the same utility. If not, let α equal the highest utility of all objects to which we have given utilities such that c is preferred to that object (if there is no such object, let $\alpha = 0$). Let β equal the lowest utility of all objects to which we have given utilities such that that object is preferred to c (if there are no such objects, let $\beta = 1$). Then set the utility of c as $u(c) = \frac{\alpha+\beta}{2}$

Step 4 Repeat step 3 for all remaining objects

[8] For the maths jocks, note that this statement is only true for countable commodity spaces. For uncountable spaces we need another condition - either a continuity condition or an order-denseness condition

It is clear that our procedure works in step 2. To show that it works at step 3, assume that $a \succ b$ so $u(a) = 0.5$ and $u(b) = 0.25$ (you can check that this works in the other cases as well). There are now 5 possibilities

1. $c \succ a$. In this case $\alpha = 1$ and $\beta = 0.5$, so $u(c) = 0.75$. Note that, by transitivity, this implies that $c \succ b$, so our function works

2. $c \sim a$. In this case $u(c) = 0.5$. Note that, by transitivity, this implies that $c \succ b$, so our function works

3. $a \succ c$ and $c \succ b$. In this case $\alpha = 0.5$ and $\beta = 0.25$, so $u(c) = 0.375$

4. $c \sim b$ In this case $u(c) = 0.25$. Note that, by transitivity, this implies that $a \succ c$, so our function works

5. $b \succ c$. In this case $\alpha = 0.25$ and $\beta = 0.$, so $u(c) = 0.125$. Note that, by transitivity, this implies that $a \succ c$, so our function works

You should convince yourself that only one of these possibilities can occur - for example it is impossible for $c \succ a$ and $b \succ c$ by transitivity.

So step 3 works. To show that step 4 also works is beyond the scope of this course, as it relies on a proof by induction (which basically means assuming that it works with n objects and showing it works with $n + 1$ objects).

So, assuming that we have a complete and transitive[9] preference ordering, then we now have our utility function. Great! However, it is very important to remember that the utility function cannot tell us any more than the original preference ordering. By this I mean that the utility function can tell us whether one object is preferred to another, but it cannot tell us any more than that. For example, it cannot tell us things like one good is **twice as** good as another good - even if the utility of one good is twice as large as another.

Why is this? Well, remember, all the utility function is doing is representing the original preferences. And if I currently have a utility function in which $u(y)$ is twice as high as $u(x)$, I can

[9] assuming we have countability

always pick another one that represents the preferences just as well in which the utility is four times as high as $u(x)$. To see this, define another utility function v such that for every object s, $v(s) = u(s) - \frac{2}{3}u(x)$. Does this new function represent the original preferences? Sure! Say $a \succ b$. As the utility function u represented these preferences, then this implies that $u(a) > u(b)$. But now subtract $\frac{2}{3}u(x)$ from each side. The inequality still holds, so

$$u(a) - \frac{2}{3}u(x) > u(b) - \frac{2}{3}u(x)$$

so

$$v(a) > v(b)$$

But now

$$v(x) = u(x) - \frac{2}{3}u(x) = \frac{1}{3}u(x)$$

while

$$\begin{aligned} v(y) &= u(y) - \frac{2}{3}u(x) \\ &= 2u(x) - \frac{2}{3}u(x) \\ &= \frac{4}{3}u(x) \end{aligned}$$

So now the utility of y (as measured by v) is 4 times as high as that for x.

More generally, if we have one utility function that represents a set of preferences, then any **strictly increasing transform** of that utility function will also represent those preferences. A monotonic transformation of a function $f(x)$ is another function $g(f(x))$ such that $g(.)$ is strictly increasing. Thus adding or subtracting things from a utility function, or multiplying it by a positive number, or taking logs, will still mean it is the same utility function!

The next thing we can think about is what sort of utility function represents the preferences we have come across in the above section. Remember, these are functions that will take as their input the amount of good one and amount of good two that the consumer consumes, so $u(x_1, x_2)$ It should be pretty obvious that monotonic preferences should be represented by monotonic utility function: That is that utility should be increasing in both good 1 and good 2. If we add an amount s to good 1 and t to good 2 such that $s \geq 0$ and $t \geq 0$ then utility cannot go down

$$u(x_1 + s, x_2 + t) \geq u(x_1, x_2)$$

Furthermore, for strictly monotonic utility functions, if either $s > 0$ or $t > 0$ then utility has to be strictly higher

$$u(x_1 + s, x_2 + t) > u(x_1, x_2)$$

Convex preferences are (somewhat confusingly) represented by a utility function that is **quasi** concave. This means that, for any x_1, x_2, the set of points y_1, y_2 such that $u(y_1, y_2) \geq u(x_1, x_2)$ is convex (if this definition doesn't mean very much for you, don't worry. That is one for the maths jocks)

What about perfect substitutes? Well, these were goods such that all the consumer cared about was the total number of the two good (remember, the total number of bottles of water). Thus, $(x_1, x_2) \succ (y_1, y_2)$ if and only if $x_1 + x_2 > y_1 + y_2$. We can therefore use this for a utility function for this type of goods[10]

And perfect compliments. Well, here we said that the consumer only cared about the total number of completed pairs of good 1 and good 2 they could make. This implies that their preferences are governed by the smaller of good 1 and good 2. We can therefore represent their preferences by the function $u(x_1, x_2) = \min(x_1, x_2)$

A final class of interesting utility functions are called the Cobb-Douglas utility functions. These are utility functions of the form

$$u(x_1, x_2) = x_1^{\alpha} x_2^{\beta}$$

where α and β are positive numbers less than or equal to 1. You will learn to love Cobb-Douglas utility functions, as they are very easy to work with. For a start, they represent preferences that are monotone and convex (check)! Therefore many economists (and exam setters) choose to work with them. Figure 15 shows a graph of the indifference curves for $u(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$ for different levels of utility.

Now that we have defined the concept of utility, we can define the concept of **marginal utility** This measures the rate at which utility increases as we increase the amount the consumer has of

[10] In fact, we use perfect substitutes more generally to denote any preferences that can be represented by a utility function of the form $\alpha x_1 + \beta x_2$.

a particular good.[11] It is defined as the partial derivative of the utility function with respect to that good. We will use $MU_1(x_1, x_2)$ to denote the marginal utility with resect to good 1 when the consumer has the bundle x_1,x_2

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

For a Cobb Douglas utility function, we can see that

$$\begin{aligned} MU_1(x_1, x_2) &= \frac{\partial u(x_1, x_2)}{\partial x_1} \\ &= \frac{\partial x_1^{\alpha} x_2^{\beta}}{\partial x_1} \\ &= \alpha x_1^{(\alpha-1)} x_2^{\beta} \end{aligned}$$

A few things to note about this

1. It is positive: So as we give the consumer more of good 1, their utility goes up. This shows that the preferences represented are monotonic
2. For a fixed x_2, marginal utility is **decreasing** in x_1. That is, as we increase the amount of x_1 the consumer has, the less happy an additional unit of x_1 will make them
3. For a fixed x_1, marginal utility is **increasing** in x_2. That is, as we increase the amount of x_2 the consumer has, the more happy an additional unit of x_1 will make them

2 and 3 together follow from the fact that Cobb-Douglas preferences are convex.

Marginal utility is related to the concept of the MRS. To see this, remember that the MRS is the slope of an indifference curve. The one thing we know about indifference curves is that every point on it has the same utility (yes?). Therefore, lets take the total derivative of the utility function with respect to both x_1 and x_2

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

But, if we are moving along an indifference curve, then $du(x_1, x_2) = 0$. Thus, we can rearrange the above expression

[11] If you followed the previous discussion, then the concept of marginal utility should be making you a bit nervous. Given that utility functions are not unique, how much can we trust marginal utility? Could we arbitrarily double marginal utility? Could we change its sign?

$$
\begin{aligned}
0 &= \frac{\partial u(x_1, x_2)}{\partial x_1} dx_1 + \frac{\partial u(x_1, x_2)}{\partial x_2} dx_2 \\
\Rightarrow &\quad \frac{\partial u(x_1, x_2)}{\partial x_1} dx_1 = -\frac{\partial u(x_1, x_2)}{\partial x_2} dx_2 \\
\Rightarrow &\quad \frac{dx_2}{dx_1} = -\frac{\frac{\partial u(x_1,x_2)}{\partial x_1}}{\frac{\partial u(x_1,x_2)}{\partial x_2}}
\end{aligned}
$$

Thus, the marginal rate of substitution between two goods is equal to the ratio of the marginal utilities of those two goods.

6 Optimal Choice

We have now set up the three elements of the consumers optimization problem:

Choose: *a consumption bundle*

In order to maximize: *preferences (or utility)*

Subject to: *the budget constraint.*

In this section we will think about the tools that we need to solve the optimization problem, and therefore complete our model of consumer behavior. To start with, we will think about the problem intuitively, and then we will show how we can use the tools of calculus to find the solution.

For the purposes of this section, we will concentrate on preferences that are monotonic. This means we know that indifference curves get 'better' (i.e. represent more preferred objects) as we move in a North-Easterly direction in the commodity space. The problem of choosing a consumption bundle in order to maximize preferences is therefore the same as choosing a consumption bundle on the most North-Easterly indifference curve possible.

What is such an optimal bundle going to look like? There are two possibilities, which we will illustrate by examples. First, let us go back to the example of perfect substitutes. In particular, lets assume that the preferences of consumers can be represented by

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

And let us assume that the price of goods is 1 for good 1 and 1 for good 2, and the consumers income is 10. The consumers problem is therefore

Choose: *a consumption bundle* (x_1, x_2)

In order to maximize: $u(x_1, x_2) = x_1 + 2x_2$

Subject to: $x_1 + x_2 \leq 10$

This budget constrain and indifference curves are illustrated in figure 16. The slope of the budget line is $-\frac{p_1}{p_2} = -1$, while the slope of the indifference curves is always $-\frac{1}{2}$. The indifference curves are therefore flatter than the budget constraint. What is the optimal consumption bundle for this consumer? Well, remembering that the optimal point is that which gets the consumer onto the most North-Easterly indifference curve, it should be clear that it is optimal for the consumer to spend **all** of their money on good two. A quick examination of figure 16 will show that the indifference curve going through the point $(0, 10)$ is the best that the consumer can obtain This indifference curve is clearly obtainable (by consuming $(10, 0)$, which is affordable). Moreover, it is better than all the indifference curves below it, and all the indifference curves above it are unobtainable. We call such a solution a **corner solution** because it occurs at one corner of the budget constraint.

We should also be able to convince ourselves that this is the best bundle for the consumer in the following way: We know that the consumer can trade off goods in the market at a ratio of 1 to 1: If they consume one less unit of good one they can consume one more unit of good two. However, the marginal rate of substitution along the indifference curve is $\frac{1}{2}$ (note that it is only because these goods are perfect substitutes that the MRS is a constant): this means that to keep the consumer indifferent, they would only have to give them half a unit of good 2 in return for a unit of good 1 in order to keep them indifferent. Thus, if the consumer is currently consuming any units of good 1, they could always trade that in for more good two and make themselves happier. It is therefore can only be optimal for the consumer to spend all their money on good 2.

Next we consider the case of perfect compliments. Remember, these are good that can be represented by a utility function such as $u(x_1, x_2) = \min(x_1, x_2)$. In this case, let us assume that the the price of good one is 1, the price of good two is 2 and the consumer again has income 10. The optimization problem is now

Choose: *a consumption bundle* (x_1, x_2)
In order to maximize: $u(x_1, x_2) = \min(x_1, x_2)$
Subject to: $x_1 + 2x_2 \leq 10$

This example is illustrated in figure 17. Now what is the optimal bundle for the consumer to consume? From figure 17, it should be clear that the optimal point is for the consumer to consume an equal proportion of good one and good two: the highest indifference curve that they can get on is the one which touches the budget constraint, but does not cross the budget line. This occurs at the 'kink' in the indifference curves, where $x_1 = x_2$. Knowing this, we can use the budget constraint to solve explicitly for x_1 and x_2

$$\begin{aligned} x_1 + 2x_2 &= 10 \\ x_1 + 2x_1 &= 10 \\ x_1 &= \frac{10}{3} = x_2 \end{aligned}$$

We call this type of solution, where the consumer consumes strictly positive amounts of both good an **interior solution**.

Again, we should be able to see intuitively that for this example the solution is for the consumer to consume the same amount of both goods. Say that the consumer wasn't doing this and was (for example) consuming more of good one than good two. Then the consumer could consume less of good one *without making themselves any worse off.* By consuming less of good one, they could buy more of good two, which would make them better off, as $u(x_1, x_2) = \min(x_1, x_2)$. Thus, the only case where the consumer cannot make themselves better off is where they are consuming exactly the same amount of good one and good two.

6.1 The Tangency Condition

So far we identified the two possible types of solutions we might find for the optimization problem: interior solutions and corner solutions. In general this may not tell us very much. However, this is a very useful differentiation in the case where the consumer's indifference curves are *differentiable.* For the purposes of this course, we can think of a differentiable function as one that is 'smooth', in the sense that it has a unique tangent (and therefore unique slope) at any point. Of all the functions that we have come across so far, the only one that is not differentiable is $\min(x_1, x_2)$. As figure 18 shows, this function does not have a unique slope at its kink.

The interesting property of interior solutions for differentiable functions is that, at any such solution, **the slope of the indifference curve has to be the same as that of the budget line -** in other words they are tangent**.** This is an incredibly useful property, and one that you can see this from figure 19. This shows two points at which the budget line is not tangent to the indifference curves - (y_1, y_2) and (z_1, z_2) - and one where it is - (x_1,x_2). If we look at (y_1, y_2), it is clear that this is not an optimal consumption bundle - the consumer could move onto a higher indifference curve by swapping good 1 for good 2 - i.e. by moving down the budget constraint.

Similarly, at (z_1, z_2), you can move to a higher indifference curve by swapping good 1 for good 2. However, at (x_1, x_2), there is nowhere to go - either moving up or down the budget line will move the consumer to a lower indifference curve. Thus, this is (or at least could be) the optimal consumption bundle.

This tangency condition also makes sense: at the tangency point, we know that the slope of the budget line equals the slope of the indifference curve. The first of these slopes is the rate at which the consumer can trade off good 1 for good 2 in the market, while the second measures the rate at which the consumer trades off good 1 for good 2 while remaining indifferent. If there were a difference between these two ratios, then the consumer could move to a different indifference curve because the market prices would either over- or under- value good 1 relative to good 2 compared to the preferences of the consumer. The consumer could therefore swap good 1 for good 2 (or visa versa) and so move to a higher indifference curve. Thus, it is only when the two are equal that the market 'values' the tradeoff of good 1 to good 2 at the same rate as the consumer, and so they cannot move on to a higher indifference curve.

So we know that any interior solution has to be a tangency point. Is it the case that any tangency point is also an interior solution? The answer is no! In fact, sometimes the tangency point is the **worst** bundle of goods that the consumer could choose while spending all their money. Two examples of tangency points that are not interior solutions are shown in figures 20 and 21. In figure 20 we look at indifference curves for concave preferences - the opposite of convex preferences (i.e. people prefer 'extreme' bundles to average bundles). Here, we can certainly find a point of tangency - (x_1, x_2). However, we are now finding the worst indifference curve that is feasible from the budget line - any other point on the budget line would but the consumer on a higher indifference curve. The optimal bundle here is in fact a corner solution - (y_1, y_2). Figure 21 shows a case where the optimal bundle is in fact a point of tangency - (z_1, z_2). However, other points of tangency - such as (w_1, w_2) are not optimal bundles.

So the tangency condition is *necessary* for a bundle to be optimal, but not *sufficient*. This is just restating that any interior optimum must be a tangency point, but not all tangency points are optimal. Thus, tangency conditions are not 'magic bullets', that give us a solution to the problem for free. However, they do suggest a procedure for finding the optimal bundle. We know that **either** the solution to the optimization problem will be a corner solution (i.e. $x_1 = 0$ or $x_2 = 0$), or it will be a point of tangency. Thus, a sure fire (if brute force) way of finding the optimal bundle is to list all such points (i.e. all points of tangency and all corner solutions) and determine which is the best of those bundles. Unless you have a reason for a priori ruling out some possible solutions to the optimization problem, then this is the approach you should take to solving the problem

One case in which tangency conditions are both necessary and sufficient is when preferences are strictly convex - for example those which are represented in figure 19. An examination of this graph should tell you the following facts[12]

1. There can be at most one tangency point between strictly convex, monotone preferences and the budget constraint
2. Such a point must be a maximum

This is one of the reasons that economists love convex preferences so much! Though note there is still one slight issue - even with a convex function, there may be no points of tangency (this is a homework hint for those who are staying awake), in which case the solution to the problem will be a corner solution.

6.2 Locating Tangency Points

All this is very well and good, but knowing about tangency conditions is not very helpful if we don't know how to locate them. Luckily, this is where we can use some of the machinery that mathematicians have developed - in particular the tools of calculus. This is one of the reasons that we went to all the effort of deriving a numerical utility function to start with.[13]

Remember a point of tangency is a point at which the slope of the budget line equals the slope of the indifference curve. We know exactly what the slope of the budget line is: it is just the ratio $-\frac{p_1}{p_2}$. However, we have also discussed how to find the slope of the indifference curve. Remember that this is marginal rate of substitution of one good for another. Furthermore, we have already established the result that the MRS of two goods is the ratio of the marginal utility with respect to those two goods. This gives us the following result

[12] Which we can also prove, but this goes beyond the scope of this course

[13] Though, maths jocks, note that we have not discussed conditions that guarantee the existence of a differentiable utility function, so I am taking some liberties here.

Result 1 A bundle (x_1, x_2) is a point of tangency between the budget line and an indifference curve if

$$-\frac{p_1}{p_2} = -\frac{\frac{\partial u(x_1,x_2)}{\partial x_1}}{\frac{\partial u(x_1,x_2)}{\partial x_2}} = -\frac{MU_1(x_1,x_2)}{MU_2(x_1,x_2)} = -MRS(x_1,x_2)$$

and so

$$\frac{p_1}{p_2} = \frac{\frac{\partial u(x_1,x_2)}{\partial x_1}}{\frac{\partial u(x_1,x_2)}{\partial x_2}} = \frac{MU_1(x_1,x_2)}{MU_2(x_1,x_2)} = MRS(x_1,x_2)$$

The tools of calculus also give us a way to determine whether people have strictly convex preferences An examination of figure 17 shows us that the defining characteristic of strictly convex preferences is that the marginal rate of substitution *decreases* along the indifference curve as x_1 increases. At low x_1 the consumer will trade a lot of x_2 for a little x_1, while at high x_1 they will trade a lot of x_1 for little x_2. Thus, the slope of the indifference curve must be flattening as x_1 increases - or becoming less negative. This means that the *second derivative* of the indifference curve must be positive. This will become more clear with the following example.

6.3 An Example: Cobb Douglas Preferences

We now put this all together to solve for the optimal bundle in the case of Cobb Douglas preferences. We will assume that the utility of the consumer is given by $u(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$. We will also assume that $p_1 = 3$, $p_2 = 4$ and $M = 15$. The consumer's problem is therefore

Choose: *a consumption bundle* (x_1, x_2)
In order to maximize: $u(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$
Subject to: $3x_1 + 4x_2 \leq 15$

The approach that we are going to take is to start by finding any point of tangency between the budget line and the indifference curve. We know that, from result 1, this will be a point at which

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

So we need to calculate the marginal utility with respect to x_1 and x_2

$$\begin{aligned}\frac{\partial u(x_1,x_2)}{\partial x_1} &= \frac{1}{2} x_1^{\frac{1}{2}-1} x_2^{\frac{1}{2}} = \frac{1}{2} x_1^{-\frac{1}{2}} x_2^{\frac{1}{2}} \\ \frac{\partial u(x_1,x_2)}{\partial x_2} &= \frac{1}{2} x_1^{\frac{1}{2}} x_2^{-\frac{1}{2}}\end{aligned}$$

Now, something quite handy is going to happen when we take the necessary ratio

$$\begin{aligned}\frac{\frac{\partial u(x_1,x_2)}{\partial x_1}}{\frac{\partial u(x_1,x_2)}{\partial x_2}} &= \frac{\frac{1}{2}x_1^{-\frac{1}{2}}x_2^{\frac{1}{2}}}{\frac{1}{2}x_1^{\frac{1}{2}}x_2^{-\frac{1}{2}}} \\ &= \frac{x_2}{x_1}\end{aligned}$$

Another magic property of the Cobb Douglas utility function!.

We therefore know that a point of tangency will occur when

$$\frac{3}{4} = \frac{x_2}{x_1}$$

or

$$x_2 = \frac{3}{4}x_1$$

Well, this is okay but we have only got x_2 as a function of x_1. This is not surprising, as we have not so far 'told' the procedure which budget line we are on - or what the income is. We can therefore substitute this relationship into the budget constraint:

$$\begin{aligned}3x_1 + 4x_2 &= 15 \\ &\Rightarrow 3x_1 + 4\left(\frac{3}{4}x_1\right) = 15 \\ &\Rightarrow 3x_1 + 3x_1 = 15 \\ &\Rightarrow x_1 = \frac{15}{6}\end{aligned}$$

and so $x_2 = \frac{15}{8}$.

But is this the solution to the optimization problem? Well, we know that this is the only tangency point, as it is the only solution to the above system of equations. But what about corner solutions? Well, looking at the utility function, we can see that spending all the money on good 1 or all the money on good 2 is a bad idea:

$$u(0, \frac{15}{4}) = 0^{\frac{1}{2}}\left(\frac{15}{4}\right)^{\frac{1}{2}} = 0 = u(5,0)$$

so it must be that the tangency point is the optimum bundle.

In fact, we didn't need to do this last step, as Cobb Douglas preferences are strictly convex. To see this, take some level of utility $\bar{u}$, and note that the indifference curve related to this utility level is given by

$$\bar{u} = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$$

or

$$\begin{aligned} \frac{\bar{u}}{x_1^{\frac{1}{2}}} &= x_2^{\frac{1}{2}} \\ x_2 &= \frac{\bar{u}^2}{x_1} \end{aligned}$$

The first derivative of this function is

$$\frac{dx_2}{dx_1} = -\frac{\bar{u}^2}{x_1^2}$$

While the second derivative is

$$\frac{d^2x_2}{dx_1^2} = 2\frac{\bar{u}^2}{x_1^3}$$

Which is positive. The slope of the indifference curve is therefore increasing (becoming less negative) as x_1 increases, which is how we characterized convex preferences!

Congratulations! You have just solved your first 'standard' economics optimization problem. You will be absolutely amazed how many more problems of this type I will force you to solve.

Just to reiterate, the steps we went through to solve this problem are as follows:

1. Solve for the marginal utilities $\frac{\partial u(x_1,x_2)}{\partial x_1}$ and $\frac{\partial u(x_1,x_2)}{\partial x_2}$
2. Use the tangency condition $\frac{p_1}{p_2} = \frac{\frac{\partial u(x_1,x_2)}{\partial x_1}}{\frac{\partial u(x_1,x_2)}{\partial x_2}}$ to get an expression for x_1 as a function of x_2
3. Substitute this expression into the budget constraint to get an explicit solution for x_1 and x_2 (as a function of the parameters of the problem: p_1, p_2 and M)
4. Check that the tangency solutions are better than the corner solutions

For completeness, it is worth noting that there is another way of solving this problem. Seeing as preferences are monotonic, we know that the consumer will always spend all their money. This means that if we know what x_1 is we know what x_2 is, as

$$\begin{aligned} 3x_1 + 4x_2 &= 15 \\ &\Rightarrow \quad x_2 = \frac{15}{4} - \frac{3}{4}x_1 \end{aligned}$$

Thus, we can substitute this into the utility function, which is now only a function of x_1

$$\begin{aligned} u(x_1, x_2) &= x_1^{\frac{1}{2}} x_2^{\frac{1}{2}} \\ &= x_1^{\frac{1}{2}} \left(\frac{15}{4} - \frac{3}{4} x_1 \right)^{\frac{1}{2}} \end{aligned}$$

This is now an unconstrained optimization problem, so we can differentiate with respect to x_1 and set that equal to zero to find the critical points

$$\begin{aligned} & \frac{d}{dx_1} x_1^{\frac{1}{2}} \left(\frac{15}{4} - \frac{3}{4} x_1 \right)^{\frac{1}{2}} \\ = & \frac{1}{2} x_1^{-\frac{1}{2}} \left(\frac{15}{4} - \frac{3}{4} x_1 \right)^{\frac{1}{2}} - \frac{3}{8} x_1^{\frac{1}{2}} \left(\frac{15}{4} - \frac{3}{4} x_1 \right)^{-\frac{1}{2}} = 0 \end{aligned}$$

so

$$\frac{1}{2} x_1^{-\frac{1}{2}} x_1^{\frac{1}{2}} \left(\frac{15}{4} - \frac{3}{4} x_1 \right)^{\frac{1}{2}} = \frac{3}{8} x_1^{\frac{1}{2}} \left(\frac{15}{4} - \frac{3}{4} x_1 \right)^{-\frac{1}{2}}$$

which gives

$$\begin{aligned} \frac{15}{4} - \frac{3}{4} x_1 &= \frac{6}{8} x_1 \\ x_1 &= \frac{15}{6} \end{aligned}$$

The same answer as the other approach! Of course, this approach doesn't excuse you for checking that this solution is better than the corner solutions.

6.4 The Dual Problem

Before we go on, we need to introduce a problem closely related to the consumer's optimization problem: This is the expenditure minimization problem. So far we have asked the question: Assume that the consumer has a fixed budget constraint: What is the maximum utility they can achieve given their income:

> **The consumer's optimization problem**
>
> **Choose:** *a consumption bundle* (x_1, x_2)
>
> **In order to maximize:** $u(x_1, x_2)$
>
> **Subject to:** $p_1 x_1 + p_2 x_2 \leq M$

A closely related problem is the following: Imagine that the consumer has to achieve a fixed level of utility: What is the least expensive way for them to do this? The dual problem is defined as

The consumer's dual problem

Choose: *a consumption bundle* (x_1, x_2)

In order to minimize: $p_1x_1 + p_2x_2$

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

For the consumer's problem, we fixed a budget line and asked what the highest indifference curve that the consumer could obtain. The dual problem fixes an indifference curve, and asks what the smallest budget constraint can obtain that utility (assuming fixed prices).

It seems like the solution to these two problems should be related: in particular it seems like the tangency results we derived for the consumer's optimization problem should also hold for the dual problem. And in fact this is the case - *assuming we have monotonic preferences*[14]*.* In this case, the following is true:

Result 2 Say a bundle (x_1, x_2) solves the consumer's problem: it maximizes $u(x_1, x_2)$ subject to $p_1x_1 + p_2x_2 \leq M$, and utility $u(x_1, x_2) = u^*$. If preferences are monotonic, then (x_1, x_2) also solves the dual problem: it minimizes $p_1x_1 + p_2x_2$ subject to $u(x_1, x_2) = u^*$. Moreover, the minimum expenditure is M

This somewhat pointless-seeming exercise is going to come in very useful in the next section.

[14] Maths jocks: we also need continuity here

7 The Demand Function

Now we know how to solve the consumer's problem, which is great: If we know (or are prepared to make assumptions about) a consumer's preferences (represented by a utility function), we can figure out what consumption bundle they will choose from a given budget set. In this section we will look at how this optimal bundle - the bundle that the consumer chooses - depends on the properties of that budget set.

Have another look at the budget set

$$p_1x_1 + p_2x_2 \leq M$$

The set is defined by three *parameters*: p_1, p_2 and M.[15] These are things that the consumer cannot control (they do not get to choose the prices or income they receive), but they clearly can affect what it is that the consumer will choose. In other words, the bundle that the consumer chooses will depend of the values of these things. This relationship is captured by what we call the **demand function**, which we write as

$$x_1^*(p_1, p_2, M)$$
$$x_2^*(p_1, p_2, M)$$

This is the amount of good 1 and good 2 that the consumer will choose if they face prices p_1, p_2 and M. In other words, $(x_1^*(p_1, p_2, M), x_2^*(p_1, p_2, M))$ is the solution of the consumer's optimization problem:

> The consumption bundle $(x_1^*(p_1, p_2, M), x_2^*(p_1, p_2, M))$ maximize $u(x_1, x_2)$ subject to the budget constraint $p_1x_1 + p_2x_2 \leq M$

Clearly, firms, policy maker and so on are going to want to know about the properties of these functions: how does demand for a particular good change as the price of that goes up? What about when people get richer? These are the questions we are going to ask in this section: how do the consumer's optimal bundles change as the parameters of the consumer's problem change?

[15] Though remember that we have said that we can forget about one of these parameters - look at the end of section 3

7.1 The Relationship between Demand and Income.

We are first going to think about how demand for a particular good depends on M, or income - assuming that prices stay constant. As has been our tradition, we will start out by thinking about two specific examples - perfect substitutes and perfect compliments.

First, let's derive the demand function for perfect substitute preferences. Remember that the consumer's problem for perfect substitutes is

Choose: *a consumption bundle* (x_1, x_2)

In order to maximize: $u(x_1, x_2) = x_1 + x_2$

Subject to: $p_1 x_1 + p_2 x_2 = M$

Recall from section 6 that the optimal thing for the consumer to do is to spend all of their money on good one if the price of good one is less that the price of good two, and spend all of their money on good two otherwise. What about the case where the price of good one is equal to good two? Well, in the case the consumer's utility will be the same whatever they choose to consume - as long as they spend all their money (check that you agree with this statement). We can therefore only make very weak predictions in this case: we know that the consumer cannot spend less that 0 on either good, and we know that they can't spend more that $\frac{M}{p_1}$ on good one or $\frac{M}{p_2}$ on good two, but beyond that we don't know very much. Thus, we have to think about three cases

$$\begin{aligned} x_1^*(p_1, p_2, M) &= 0 \text{ if } p_1 > p_2 \\ &= \frac{M}{p_1} \text{ if } p_2 > p_1 \\ 0 &\leq x_1^*(p_1, p_2, M) \leq \frac{M}{p_1} \text{ if } p_2 = p_1 \end{aligned}$$

Using this information, we can graph how demand for good one changes with income for perfect substitutes. Such graphs are, in general, called Engle curves. Figure 22 shows the Engle curves for perfect substitutes under two conditions: where $p_1 > p_2$ and where $p_2 < p_1$.

Mathematically, we describe how demand changes with income using the concept of the **income elasticity of demand (IED).** We do this in a slightly more complicated way than might seem obvious. Rather than just measure the rate of change of demand with respect to income - i.e.

$$\frac{\partial x_1^*(p_1, p_2, M)}{\partial M}$$

We measure how a *percentage* change in income leads to a *percentage* change in demand. The income elasticity of demand is given by

$$\varepsilon_1^m = \frac{\partial x_1^*(p_1, p_2, M)}{\partial M} . \frac{M}{x_1^*(p_1, p_2, M)}$$

We will use ε_1^m to refer to the income elasticity of demand of good 1. This formula can look a bit mysterious. The way to think about it is as follows. Say that an $z\%$ change in income leads to a $y\%$ change in demand. Then we want the IED to be $\frac{y\%}{z\%}$ - in other words, the percentage change in demand is the percentage change in income multiplied by the elasticity. But

$$z\% = \frac{\Delta M}{M} \text{ and } y\% = \frac{\Delta x_1^*}{x_1^*}$$

Thus

$$IED = \frac{\frac{\Delta x_1^*}{x_1^*}}{\frac{\Delta M}{M}}$$

rearranging this gives

$$IED = \frac{\Delta x_1^*}{\Delta M} \frac{M}{x_1^*}$$

Taking the limit as ΔM goes to zero gives the formula for income elasticity for demand.

The reason for that we measure elasticity in this way is that we don't want arbitrary changes of units to affect elasticity. So for example, if we were to suddenly decide to measure x as 'ounces of potatoes' rather than 'pounds of potatoes' then this would affect $\frac{\partial x_1^*(p_1,p_2,M)}{\partial M}$ but not ε_1^m

For our perfect substitutes, we can calculate IED for the case where $p_1 < p_2$

$$\begin{aligned}
\varepsilon_1^m &= \frac{\partial x_1^*(p_1, p_2, M)}{\partial M} . \frac{M}{x_1^*(p_1, p_2, M)} \\
&= \frac{1}{P_1} . \frac{M}{x_1^*(p_1, p_2, M)} \\
&= \frac{1}{P_1} . \frac{M}{M/P_1} = \frac{1}{P_1} . P_1 = 1
\end{aligned}$$

So a 1% increase in income leads to a 1% increase in consumption of good 1.

We can perform the same exercise for perfect compliments In this case, the consumer's problem is

Choose: *a consumption bundle* (x_1, x_2)

In order to maximize: $u(x_1, x_2) = \min(x_1, x_2)$

Subject to: $p_1 x_1 + p_2 x_2 = M$

As we have demonstrated before, the best thing for the consumer to do is to buy exactly the same amount of good 1 as of good 2, so $x_1 = x_2$. Substituting this into the budget constraint gives $p_1 x_1 + p_2 x_1 = M$ and so $x_1 = \frac{M}{p_1 + p_2} = x_2$. Thus the demand function is given by

$$x_1^*(p_1, p_2, M) = \frac{M}{p_1 + p_2}$$

Figure 23 shows the Engle curve for perfect compliments. Once again, the Engle curve is a straight line, but here the slope is flatter than the curve for perfect substitutes (in the case where $p_1 < p_2$). The IED for perfect compliments

$$\begin{aligned}
\varepsilon_1^m &= \frac{\partial x_1^*(p_1, p_2, M)}{\partial M} . \frac{M}{x_1^*(p_1, p_2, M)} \\
&= \frac{1}{p_1 + p_2} \frac{M}{x_1^*(p_1, p_2, M)} \\
&= \frac{1}{p_1 + p_2} \frac{M}{\frac{M}{p_1 + p_2}} = 1
\end{aligned}$$

So again the income elasticity of demand is equal to 1.

In both these examples an increase in income leads to an increase in the amount consumed of both goods. But does this have to be the case? The answer is no! Even with monotonic, convex preferences, an *increase* in income can lead to a *decrease* in the consumption of one of the goods. We can illustrate such a case graphically in figure 24. In this case, income increases from M to M', but the tangency conditions imply that this leads to a *fall* in the demand for good one. Moreover, the indifference curves we have drawn are those of convex, monotonic preferences. Over this range of income changes, the Engle curve for this good would be downward sloping.

Is such a situation likely? I would argue that it is pretty easy to think of goods that we consume less of as we get richer. Ramen noodles spring to mind. Such goods (i.e. those with a negative IED) are known as inferior goods. All other goods are known as normal goods.

7.2 The Relationship between Demand and a Good's Own Price

Next, we would like to analyze how the demand for a good changes with the price of that good, keeping the price of the other good, and income, constant. This relationship is the **demand curve** that you may have met in Econ 011. Again, let's begin with our old friends perfect compliments and perfect substitutes. Recall, that for perfect substitutes, the demand function is given by

$$\begin{aligned} x_1^*(p_1,p_2,M) &= 0 \text{ if } p_1 > p_2 \\ &= \frac{M}{p_1} \text{ if } p_2 < p_1 \\ 0 &\leq x_1^*(p_1,p_2,M) \leq \frac{M}{p_1} \text{ if } p_2 = p_1 \end{aligned}$$

The resulting demand curve (or actually the inverse demand curve, as we have switched the axes) is shown in figure 25. For prices such that $p_1 < p_2$, the relationship between p_1 and x_1^* is given by $\frac{M}{p_1}$, which looks like a downward sloping convex curve. At the point at which $p_1 = p_2$, demand can be anywhere between $\frac{M}{p_2}$ and 0, so at this point the demand curve is a vertical line. For any prices $p_1 > p_2$, demand for good one is zero. Note that for prices around $p_1 = p_2$, tiny changes in prices can make demand 'jump' from 0 to $\frac{M}{p_1}$, so in this region, demand is very sensitive to changes in prices.

Just as we defined the concept of the income elasticity of demand, we can define the **price elasticity of demand (PED)** as

$$\varepsilon_1 = -\frac{\partial x_1^*(p_1,p_2,M)}{\partial p_1}.\frac{p_1}{x_1^*(p_1,p_2,M)}$$

We use ε_1 to refer to the price elasticity of good one. Note that, as $\frac{\partial x_1^*(p_1,p_2,M)}{\partial p_1}.\frac{p_1}{x_1^*(p_1,p_2,M)}$ is (almost) always negative, we define the PED as the negative of this value so we don't have to carry a minus sign around with us everywhere.

The price elasticity of demand of a good is one of its most important properties from the point of view of an economist (and businessmen). If a good has a high price elasticity of demand, consumers will respond a lot to changes in price, while if it has a low elasticity they will respond only a little. Goods with $\varepsilon > 1$ we call elastic, while those with $\varepsilon < 1$ we call inelastic. Why is 1 important? Think about how the amount of money, $p_1 x_1^*(p_1,p_2,M)$ the consumer will spend on a good will changes as prices change.

$$\frac{\partial}{\partial p_1}\left(p_1 x_1^*(p_1,p_2,M)\right) = x_1^*(p_1,p_2,M) + p_1\frac{\partial x_1^*(p_1,p_2,M)}{\partial p_1}.$$

This will be greater that zero if

$$\begin{aligned}
x_1^*(p_1,p_2,M) + p_1\frac{\partial x_1^*(p_1,p_2,M)}{\partial p_1} &> 0 \\
&\Rightarrow -1 < \frac{p_1}{x_1^*(p_1,p_2,M)}\frac{\partial x_1^*(p_1,p_2,M)}{\partial p_1} \\
&\Rightarrow \varepsilon < 1
\end{aligned}$$

Thus, if a good is price inelastic, consumers will spend more on it if its price goes up, whereas if it is elastic they will spend less on it as the price goes up.

For perfect substitutes (again, if $p_2 < p_1$) the price elasticity of demand for good 1 is given by

$$\begin{aligned}
\varepsilon_1 &= -\frac{\partial x_1^*(p_1,p_2,M)}{\partial p_1}.\frac{p_1}{x_1^*(p_1,p_2,M)} \\
&= \frac{M}{p_1^2}\frac{p_1}{\frac{M}{p_1}} = 1
\end{aligned}$$

Perfect substitutes therefore have unit elasticity when $p_1 < p_2$. This makes sense - as the consumer is spending all their money on good 1, a change in the price of that good will not affect how much money they spend on that good.

Moving on to perfect compliments, here we remember that

$$x_1^*(p_1,p_2,M) = \frac{M}{p_1+p_2}$$

Figure 27 shows the inverse demand curve, while the elasticity is given by

$$\begin{aligned}
\varepsilon_1 &= -\frac{\partial x_1^*(p_1,p_2,M)}{\partial p_1}.\frac{p_1}{x_1^*(p_1,p_2,M)} \\
&= \frac{M}{(p_1+p_2)^2}\frac{p_1}{\frac{M}{(p_1+p_2)}} = \frac{p_1}{p_1+p_2}
\end{aligned}$$

So good 1 will always be price inelastic, as $\frac{p_1}{p_1+p_2} < 1$. This also makes sense. As the consumer will always buy the same amount of good 1 and good 2, an increase in the price of good 1 will lead to an increase in spending on that good.

Again, in both these examples we have had demand falling as price increases. Do our assumptions guarantee that? Interestingly, the answer is again no! Figure 28 demonstrates a case in which monotonic, convex indifference curves can lead to demand for good 1 *increasing* as the price goes up. So what the hell is going on here?

On way to think about this is that the change from p_1 to p'_1 actually has two effects. It changes the rate at which good one can be traded off for good two in the market place (the change in the price ratio). But it *also* makes the consumer worse off, in the sense that they can no longer afford to be on the same indifference curve. We can therefore separate out these two effects: We can ask what the effect of the price change would be if consumers could stay on the same indifference curve, then we could ask what effect the change in income has, resulting from the price change.

Figure 29 makes these two effects explicit, by doing the following thought experiment: First, imagine that we change prices (increased p_1) *but gave the consumer an increase in income*, so they could keep on the same indifference curve. This would lead the consumer to move from A to B, and decrease the amount they consume of good 1. We call this the **substitution effect**. You should convince yourself that the substitution effect must always be negative for convex preferences.

However, we also have to deal with the fact that the price increase makes the consumer poorer - they do not in fact, get the income increase to offset the price rise, so they will end up on a lower indifference curve. This is the pure **income effect** and moves them from B to C. This effect could go either way. If the good is normal, then the income effect will also act to reduce demand for good 1. However, if the good is inferior, then the effective reduction in income could lead to an *increase* in demand for good 1. In extreme cases, this income effect could offset the substitution effect, and an increase in the price of good 1 could lead to an increase in the demand for that good.

To make this formal, we need to think again about the dual problem, which (as you remember) is

> **The consumer's dual problem**
>
> **Choose:** *a consumption bundle* (x_1, x_2)
>
> **In order to minimize:** $p_1x_1 + p_2x_2$
>
> **Subject to:** $u(x_1, x_2) = \bar{u}$

Here, we can define a demand function which measures how demand changes as price changes, *utility is fixed,* but income may change (in contrast to the standard demand curve, which keeps income fixed, and allows utility to change). This demand is called **compensated** (or Hicksian) demand, which we will define as the solution to the dual problem.

$$(h_1^*(p_1, p_2, \bar{u}), h_2^*(p_1, p_2, M)) \text{ minimizes } p_1x_1 + p_2x \text{ subject to } u(x_1, x_2) = \bar{u}$$

The compensated demand is what we calculate when we go from point A to point B in figure 29. We can also define what we are going to call the **expenditure function:** the cost of obtaining utility $\bar{u}$ with prices p_1 and p_2

$$e(p_1, p_2, \bar{u}) = p_1 h_1^*(p_1, p_2, \bar{u}) + p_2 h_2^*(p_1, p_2, \bar{u})$$

We can use these two concepts to formally break down the change in demand due to price changes into the income and substitution effects. To see this, note that we claimed that the bundle that solves the consumer's problem is the same as that which solves the dual problem if we are talking about the same indifference curve. Say we are on the indifference curve that gives utility $\bar{u}$. Then

$$x_1^*(p_1, p_2, e(p_1, p_2, \bar{u})) = h_1^*(p_1, p_2, \bar{u})$$

In other words, the bundle that minimizes expenditure while obtaining utility $\bar{u}$ is the same as the bundle that maximizes utility if income is equal to $e(p_1, p_2, \bar{u})$, or the amount needed to obtain utility $\bar{u}$.

Differentiating both sides of this with respect to p_1 gives[16]

$$\begin{aligned} \frac{\partial x_1^*}{\partial p_1} + \frac{\partial x_1^*}{\partial M}\frac{\partial e}{\partial p_1} &= \frac{\partial h_1^*}{\partial p_1} \\ &\Rightarrow \frac{\partial x_1^*}{\partial p_1} = \frac{\partial h_1^*}{\partial p_1} - \frac{\partial x_1^*}{\partial M}\frac{\partial e}{\partial p_1} \end{aligned}$$

Which already looks like we have split the effect between the substitution effect (the first term) and the income effect (the second term). In fact, we can do better than this, because we know something about $\frac{\partial e}{\partial p_1}$

$$\begin{aligned} \frac{\partial e(p_1, p_2, \bar{u})}{\partial p_1} &= \frac{\partial}{\partial p_1}(p_1 h_1^*(p_1, p_2, \bar{u}) + p_2 h_2^*(p_1, p_2, \bar{u}) \\ &= h_1^*(p_1, p_2, \bar{u}) + p_1\frac{\partial h_1^*}{\partial p_1} + p_2\frac{\partial h_2^*}{\partial p_1} \end{aligned}$$

[16] For convenience, we will now write x_1^* for $x_1^*(p_1, p_2, e(p_1, p_2, \bar{u}))$ and h_1^* for $h_1^*(p_1, p_2, \bar{u})$, but you should remember that these are still functions.

The change in the expenditure function can be split into the direct effect of the change in prices $(h_1^*(p_1, p_2, \bar{u}))$ and the indirect effect due to the change in the optimal bundle $\left(p_1 \frac{\partial h_1^*}{\partial p_1} + p_2 \frac{\partial h_2^*}{\partial p_1}\right)$. However, it turns out that, because this is the solution of an optimization problem we can ignore the indirect effect. This is an application of the **envelope theorem.** To see this, notice that, changes in compensated demand due to the change in price cannot lead to changes in utility

$$\frac{\partial u}{\partial p_1} = \frac{\partial u}{\partial h_1^*} \frac{\partial h_1^*}{\partial p_1} + \frac{\partial u}{\partial h_2^*} \frac{\partial h_2^*}{\partial p_1} = 0$$

But, from the tangency condition, we know that

$$\begin{aligned} \frac{\frac{\partial u}{\partial h_2^*}}{\frac{\partial u}{\partial h_1^*}} &= \frac{p_2}{p_1} \\ &\Rightarrow \frac{\partial u}{\partial h_2^*} = \frac{p_2}{p_1} \frac{\partial u}{\partial h_1^*} \end{aligned}$$

So we get

$$\begin{aligned} \frac{\partial u}{\partial h_1^*} \frac{\partial h_1^*}{\partial p_1} + \frac{\partial u}{\partial h_2^*} \frac{\partial h_2^*}{\partial p_1} &= 0 \\ &\Rightarrow \frac{\partial u}{\partial h_1^*} \frac{\partial h_1^*}{\partial p_1} + \frac{p_2}{p_1} \frac{\partial u}{\partial h_1^*} \frac{\partial h_2^*}{\partial p_1} = 0 \\ &\Rightarrow \frac{\partial u}{\partial h_1^*} \frac{1}{p_1} \left(p_1 \frac{\partial h_1^*}{\partial p_1} + p_2 \frac{\partial h_2^*}{\partial p_1} \right) = 0 \\ &\Rightarrow p_1 \frac{\partial h_1^*}{\partial p_1} + p_2 \frac{\partial h_2^*}{\partial p_1} = 0 \end{aligned}$$

and so

$$\frac{\partial e(p_1, p_2, \bar{u})}{\partial p_1} = h_1^*(p_1, p_2, \bar{u}) = x_1^*(p_1, p_2, e(p_1, p_2, \bar{u}))$$

Substituting back into the equation above gives us

$$\frac{\partial x_1^*}{\partial p_1} = \frac{\partial h_1^*}{\partial p_1} - \frac{\partial x_1^*}{\partial m} x_1^*$$

This is the (rightfully) famous **Slutsky Equation.**

The change in demand with respect to a price change is the sum of the change in compensated demand (the substitution effect), plus the income effect $\frac{\partial x_1^*}{\partial m} x_1^*$.

For an increase in prices to lead to an increase in demand it has to be the case that this expression is positive. As $\frac{\partial h_1^*}{\partial p_1}$ is negative it must be the case that $\frac{\partial x_1^*}{\partial m} x_1^*$ is both negative and large enough to overcome the substitution effect. This gives us a hint of the type of goods for which this might be true. First, such goods have to be inferior (as $\frac{\partial x_1^*}{\partial m}$ has to be negative). Second, they have to make up a large proportion of consumption (x_1^* has to be large). Such goods are called **Giffen goods.**

Do such goods exist in practice? Evidence is mixed. The standard example given is that of staple foods. Imagine that you are an ancient British peasant, and you are choosing between eating gruel and beef. If the price of gruel goes up, then you may not be able to afford to buy beef - but in order to keep yourself alive you need a certain amount of calories, so you buy more gruel. Economists have spent some time looking for examples such as this, and some recent evidence from consumption patterns in China have suggested that rice may indeed be a Giffen good under some circumstances.

7.3 The Relationship between Demand and the Other Good's Price

The final relationship we are interested in is the relationship between the price of good 2 and the amount demanded of good 1: How does demand for good one change as we change the price of good two? Intuitively, we can think of this relationship as going either way. Consider the case of Pepsi and Coke. Here, you might expect an increase in the price of pepsi to lead to an *increase* in the demand for coke. This is the case of **substitutes** - or goods that can be substituted for each other. On the other hand, consider the case of golf bags and golf clubs. Here, you might expect a rise in the price of golf clubs to lead to a *fall* in the demand for golf bags - if you are buying fewer golf clubs, then you are going to need fewer golf bags as well. Such goods are called **compliments.**

These two words should sounds familiar - we have been banging on about perfect substitutes and perfect compliments for this entire chapter. Figure 30 shows the relationship between $x_1^*(p_1, p_2, M)$ and p_2 for perfect substitutes. For $p_2 < p_1$ demand for good 1 is zero. For $p_2 > p_1$ demand for good 1 is $\frac{M}{p_1}$. For prices such that p_1 is close to p_2, then a tiny increase in the price of good two can lead to a 'jump' in demand for good 1 from 0 to $\frac{M}{p_1}$.

Figure 31 shows the same relationship for perfect compliments. Here the relationship is smoothly downward sloping for all values of p_2. In fact, the curve is exactly the same shape as the demand curve for good 1: For perfect compliments, $x_1^*(p_1, p_2, M)$ responds in exactly the same way to a change in the price of good 1 and a change in the price of good 2!

We can also define the elasticity of demand for good one with respect to the price of good 2. We call the the **cross elasticity of demand**

$$\varepsilon_1^2 = \frac{\partial x_1^*(p_1, p_2, M)}{\partial p_2} . \frac{p_2}{x_1^*(p_1, p_2, M)}$$

A positive number indicates that the two goods are substitutes - an increase in the price of good 2 leads to an increase in the demand for good 1 - while a negative number means that they are compliments.

If we look at perfect substitutes, then we will see that the cross elasticity of demand for good 1 is 0 everywhere apart from where $p_1 = p_2$ where it is (effectively) infinity! An arbitrarily small increase in the price of good 2 will cause demand for good 1 to leap from 0 to $\frac{M}{p_1}$. This is why these goods are called perfect substitutes. In contrast, for perfect compliments, the cross elasticity of demand is given by

$$\begin{aligned} \varepsilon_1^2 &= \frac{\partial x_1^*(p_1, p_2, M)}{\partial p_2} . \frac{p_2}{x_1^*(p_1, p_2, M)} \\ &= -\frac{M}{(p_1 + p_2)^2} \frac{p_1}{\frac{M}{(p_1+p_2)}} = -\frac{p_2}{p_1 + p_2} \end{aligned}$$

which is (unsurprisingly) always negative.

7.4 Consumer Surplus

We now know how to start off with assumptions about a consumer's preferences and make predictions about their behavior in the form of a demand function. We are now breifly going to discuss what we can learn going in the other direction - in other words what we can learn about a consumer's 'welfare' from their demand function. In particular, we are going to develop some measure of how much benefit, or 'consumer surplus', a consumer gets from buying a particular good at a particular price. So, for example, we want to know how much benefit does a consumer get from being able to buy bananas for 50c each, compared to the case when they are not allowed to buy bananas.

Informally, we want to develop the idea that the area under the demand curve for bananas tells us something about the benefit a consumer gets from buying bananas. To get the idea, look at the graph shown in figure 32, which shows the demand curve, but in now with price on the vertical axis and amount on the horizontal axis. What the demand curve tells us is that, if the consumer were to be only offered x_a bananas, then she would happily pay price p_a per banana.In other words, p_a is her reservation price for x_a bananas. If she were offered x_b bananas, then here reservation price is p_b and so on. However, she only actually has to pay p^* per unit. Thus, we might get the idea that, for the first x_a bananas the consumer consumes,. the benefit of consumption outweights the price. This is the idea that consumer surplus tries to capture. as shown in figure 33. The triangle A between the price line and the demand curve is what we refer to as the consumer surplus (note that the box B is total expenditure on the good). We want to think of this as the benefit that the consumer gets from being allowed to buy this good.

How does this idea of consumer surplus link in with the carefully constructed work on preferences and so on that we have already done? Well, unfortunately the answer is 'not perfectly'. However, one case in which the two ideas do link up is in the case of quasi-linear preferences. Remember, these are preferences are of the form

$$u(x, y) = v(x) + y$$

so, in some sense, the consumer's preference for x is separated from there preference for y: their marginal utility with respect to x is independent of the amount of y they consume We can think of this as the consumer's preference with respect to a particular good, x, (bananas), and 'everything else' y. In fact, we will think of y as just amounts of money that the consumer can spend on other goods, so the consumer's budget constraint is

$$p_x x + y \leq M$$

How much does the consumer benefit from being able to consume good x? Well, if the consumer buys x units of good x at price p_x, their utility is

$$v(x) + M - xp_x$$

whereas, if they did not consume any of good x, their utility would just be (setting $v(0) = 0$)

$$M$$

Thus, the added utility of being able to consume x is $v(x) - xp_x$, the utility given by x minus income spent on x

How does this relate to the area under the demand curve? To see this, remember that the demand function records the consumers optimal consumption of good x at each price, which, assuming we have an interior solution, is given by

$$MRS_{x,y} = \frac{p_x}{p_y}$$

or, in this case

$$\frac{\partial v(x)}{\partial x} = p_x$$

Thus, if we integrate the demand function between 0 and x^* we get

$$\begin{aligned} & \int_0^{x^*} p_x(x)dx \\ = & \int_0^{x^*} \frac{\partial v(x)}{\partial x}dx \\ = & \ v(x^*) \end{aligned}$$

So the area under the demand cuve is equal to the total utility from consuming x^*. Thus, the consumer surplus equals

$$\int_0^{x^*} p_x(x)dx - xp_x$$

As xp_x is the area of the box B, the consumer surplus is given by the triangle A.

Of course, this result is specific to our choice of quazi-linear preferences, and particular the lack of income effect. However, more generally, we can hope that the area under the demand curve gives approximately the same result.

7.5 Rationality Revisited

Before closing the chapter on consumer demand, I want to return to the question of what sort of behavior is and is not in line with our assumptions. Consider the pattern of behavior exhibited in figure 33. When choosing from budget set A, the consumer chooses bundle z_1, z_2, but when the choose from budget set B they chose w_1, w_2. Is this behavior a violation of the independence of irrelevant alternatives? No it isn't, because IIA only deals with cases where one choice set can be created by adding objects to the other choice set, which is not the case here. However, this behavior still looks very funny, for a couple of reasons. Firstly, z_1, z_2 was chosen when w_1, w_2 was available and visa versa. This is bad, but could (just about) be explained by the fact that the consumer is indifferent between the two bundles. But, what about if the consumer had preferences that are strictly monotonic? In that case we really are in trouble: We can find a bundle z_1^*, z_2^* that was both available when w_1, w_2 was chosen and (by strict monotonicity) is preferred to z_1, z_2 - see figure 34 for an example. Thus by transitivity, it must be the case that $(w_1, w_2)\bar{\succ}(z_1, z_2)$ or, under the assumption of strict monotonicity, (w_1, w_2) is revealed preferred to (z_1, z_2). However, it is also the case that we can find a bundle w_1^*, w_2^* that must b e strictly preferred to w_1, w_2 and was available when z_1, z_2 was chosen. By the same logic, $(z_1, z_2)\bar{\succ}(w_1, w_2)$. Clearly, this type of behavior is not consistent with the optimization of strictly monotonic preferences.

In light of this, we clearly need to sharpen up our test of what is and is not consistent with rational behavior. In order to do so, we need to redefine what we mean by revealed preference. With strictly monotonic preferences, we can say that one bundle of goods (x_1, x_2) is strictly revealed preferred to another (y_1, y_2) if, for some set of prices, (x_1, x_2) is chosen when (y_1, y_2) was available for a lower total expenditure:

$$\begin{aligned} (x_1, x_2) &\bar{\succ} (y_1, y_2) \\ &\text{if, for some budget set } (p_1, p_2, M) \\ &(x_1, x_2) \text{ is chosen and} \\ p_1 y_1 + p_2 y_2 &< M \end{aligned}$$

It turns out (though we will not prove it in this course) that we can think of a particular set of choices as being generated by the optimization of a complete, transitive, strictly monotonic preferences if and only if the following **generalized axiom of revealed preference** holds.

Greneralized Axiom of Revealed Preference We cannot find any sequence of bundles (x_1^1, x_2^1), (x_1^2, x_2^2),....,(x_1^n, x_2^n) such that

$$\begin{aligned}(x_1^1, x_2^1) &\bar{\succ} (x_1^2, x_2^2)\\(x_1^2, x_2^2) &\bar{\succ} (x_1^3, x_2^3)\\&...\\(x_1^{n-1}, x_2^{n-1}) &\bar{\succ} (x_1^n, x_2^n)\\(x_1^n, x_2^n) &\bar{\succ} (x_1^1, x_2^1)\end{aligned}$$

In other words, it must be the case that we cannot find a cycle of bundles such that the first is revealed preferred to the second, the second to the third and so on until the last bundle is revealed to the first.

This axiom is clearly violated by our example, in which we had $(w_1, w_2)\bar{\succ}(z_1, z_2)$ and $(z_1, z_2)\bar{\succ}(w_1, w_2)$, forming a (very short) cycle.

It is in fact this condition (GARP) that are tested in the studies we discussed in section 4.2.

7.6 Do we need Rationality?

The take home message of the above section is that we can use rational choice theory to determine how people's choices are affected by the parameters of the consumer problem, and that the resulting properties are 'sensible'; in particular in most cases demand is going to be negatively related to price. However, in a famous paper in 1962, Gary Becker pointed out that we do not *need* rationality to get downward sloping demand curves. In fact, he pointed out that demand curves are likely to be downward sloping for almost any decision rule. To see this, have a look at figure 35. The point of this figure is to think about what the demand function would look like for a consumer who chose *randomly.* Well, not quite at random - we will assume that they are minimally rational in that they do choose a bundle on the budget line, but beyond that we assume that they are equally likely to pick any point on the budget curve. In this case, on average, they will spend half their money on good one, and half their money on good two. In other words, the average demand for good one is $\frac{M}{2p_1}$. Thus, if the price of good one increases (to p_1^* in the diagram), then the average demand for good one will fall to $\frac{M}{2p_1^*}$: on average, demand will be downward sloping.

What does this mean? Well, on the one hand, it is good for us, in that it means that the assumptions we need to get downward sloping demand curves are weaker than we though. Thus, if all we want to take away from this section is that demand curves should be downward sloping, we are happy. However, we are also going to want to interpret people's demand as saying something about their preferences. This we clearly cannot do if people are just choosing randomly. Moreover, this analysis suggests that, just because we see a downward sloping demand curve, doesn't mean that we can assume we are witnessing rational consumers!

8 Choice Under Uncertainty

Up until now, we have thought of the objects between which our consumers are choosing as being physical items - chairs, tables, apples, brandy etc. We pretty much know what will happen when we buy such things. However, we can also think of cases where the outcomes of the choices we make are *uncertain* - we don't know exactly what will happen when we buy a particular object. Think of the following examples:

- You are deciding whether or not to buy a share in AIG
- You are deciding whether or not to put your student loan on black at the roulette table
- You are deciding whether or not to buy a house that straddles the San Andreas fault line

In each case, while you may understand exactly what it is that you are buying, (or choosing between), the outcomes, in terms of the things that you care, about are uncertain. Here we are going to think about how to model a consumer who is making such choices.

In order to be concrete, let's think about a specific example. You are in a fairground, and come across a (very boring) game of chance. For an amount of money $\$x$, you can flip a coin. If it comes down as heads, you get \$10. If it comes down tails, you get nothing (let's assume that you get to choose the coin, so you are pretty sure that there is a 50% chance of a head and a 50% chance of tails). The question is, for what price x would you choose to play the game. In other words, you have a choice between the following two options.

1. Not play the game and get nothing
2. Play the game, and get $-x$ for sure, plus a 50% chance of getting \$10.

How would you make a decision like this? The earliest thinkers on the subject suggested the following strategy: Figure out the **expected value** (or average pay-out) of playing the game, and see if it is bigger than 0. If it is, then play the game, if not, then don't.

So what is the expected value of the game? With a 50% chance you will get $\$10-x$, while with a 50% chance you will get $-x$. Thus, the average payoff is going to be

$$\begin{aligned} & 0.5(10-x)+0.5(-x) \\ = & \ 5-x \end{aligned}$$

Thus the value of the game is $\$5-x$. In other words, following this strategy, you should play the game if the cost of playing is less than \$5.

Does this sound sensible? People though so until Daniel Bernoulli (the Dutch-Swiss maths superstar) came up with the following example:

Example 4 (The St. Petersburg Paradox) *Imagine that the fairground guy offers you a different game. Now, you first of all flip the coin. If it comes down heads, then you get \$2. If it comes down tails, you flip again. If you get heads on that go, you get \$4, otherwise you flip again. If it comes down heads* ***then*** *you get \$8, otherwise you flip again, and so on.*

What is the expected value of this game? Well, there is a $\frac{1}{2}$ probability that you will get heads on the first trial, and so get \$2. But there is a $\frac{1}{2}$ chance that you will get tails and flip again. There is then a $\frac{1}{2}$ chance that you will get heads on that go, and so get \$4. For that to happen, you would have to get tails on the first go (probability $\frac{1}{2}$) and heads on the second go (probability $\frac{1}{2}$). Thus, there is a $\frac{1}{4}$ probability that you will get \$4. *Using the same logic, there is a $\frac{1}{8}$ chance you will get* \$8 *and so on. The expected value of the game is therefore*

$$\begin{aligned} & \frac{1}{2}\$2+\frac{1}{4}\$4+\frac{1}{8}\$8+\frac{1}{16}\$16+... \\ = & \ \$1+\$1+\$1+\$1+...... \\ = & \ \infty \end{aligned}$$

The expected value of the game is ∞, and so that is how much you should be willing to pay. In other words, however much the fairground guy is prepared to charge you, you should be willing to pay it.

Assuming that you are not one of the people that is prepared to pay ∞ to play this game, what has gone wrong? Bernoulli suggested one solution: Perhaps the difference in 'happiness' brought about by getting extra money decreases as the amount of money you have increases. In other words getting \$1 extra if you only have \$1 means a lot more than getting \$1 extra if you have \$1 million. This is what we would (these days) call the decreasing marginal utility of wealth.

Example 5 *Say a pauper finds a magic lottery ticket, that has a 50% chance of \$1 million and a 50% chance of nothing. A rich person offers to buy the ticket off him for \$499,999 for sure. According to our 'expected value' method', the pauper should refuse the rich person's offer!*

Bernoulli argued that this is ridiculous. For the pauper, the difference in quality of life between getting nothing and \$499,999 is massive, while the difference between \$499,999 and \$1 million is relatively small. Thus, by turning down the rich persons offer, they are gaining relatively little (a 50% chance of getting \$1 million rather than \$499,999) and loosing an awful lot (a 50% chance of getting 0 rather than \$499,999). Moreover, Bernoulli argued. if this is the case, what we should be maximizing is **expected utility**, rather than expected value. In other words, if $u(x)$ is the utility of getting an amount x then, the pauper should choose to accept the rich guy's offer if

$$\frac{1}{2}u(\$1,000,000)+\frac{1}{2}u(\$0)<u(\$499,999)$$

The idea is that the utility gap between 0 and $\$499,999$ is larger than the gap between $\$499,999$ and $\$1,000,000$. For example, it could be that

$$\begin{aligned} u(\$0) &= 0 \\ u(\$499,999) &= 10 \\ u(\$1,000,000) &= 16 \end{aligned}$$

If this were the case, then Bernoulli suggests that the pauper should accept the offer, as the expected utility of the lottery ticket is 8, while the expected utility of the rich man's offer is 10. In fact he proposed that the utility of getting an amount x could be approximated by the function $\ln(x)$ (note that this exhibits decreasing marginal utility. If this is right, then the most that you should pay for the St. Petersburg game is about \$60.

By and large, most of modern economics agrees with Bernoulli's assessment of how choice under uncertainty should work. However, if you remember back to choice under *certainty,* we in general don't like the idea of utility functions coming out of nowhere. When we were talking about choice under certainty, we were very careful to ask the question: what has to be true about a person's preferences for us to be able to represent them with a utility function? The answer was that preferences had to be complete and transitive. Here we want to be equally careful: Under what circumstances can people's preferences be represented by an *expected* utility function? Do they just have to be complete and transitive? Or do we need something more?

To answer this question, we need to be a little bit more formal about what we are doing here. Before we were thinking about people's preferences over *objects* (chairs, lamps, Miley Cyrus CD's etc.). Now we are thinking about peoples preferences over what we will call *lotteries.* What do we mean by a lottery?. In order to fix ideas, lets think of games of chance in which you can win one of four prizes: an apple, a banana, a cockatiel and a nothing. A lottery is just a list of four numbers indicating the probability of winning each of the prizes. so for example

$$\begin{matrix} 0.15 \\ 0.35 \\ 0.5 \\ 0 \end{matrix}$$

is a lottery. This gives a 15% chance that you win the apple, a 35% chance you win the banana, a 50% chance that you win the cockatiel and a 0% chance you win nothing. Another lottery is

$$\begin{matrix} 0.1 \\ 0.1 \\ 0.7 \\ 0.1 \end{matrix}$$

The only thing that has to be true for such a list to be a lottery is (a) each of the numbers has to be greater than or equal to zero (you can't have a negative probability) and (b) they have to add to 1 (the probability of winning one of the prizes is 100%). In general, we will write

$$p = \begin{pmatrix} p_a \\ p_b \\ p_c \\ p_n \end{pmatrix}$$

for such a lottery.

We will think about a consumer that has complete and transitive preference over such lotteries. In other words,

1. **Completeness:** For any two lotteries p and q one and exactly one of the following is true: $p \succ q$ or $q \succ q$ or $p \sim q$

2. **Transitivity:** For any three lotteries p, q and r

   (a) If $p \succ q$ and $q \succ r$ then $p \succ r$

   (b) If $p \sim q$ and $q \sim r$ then $p \sim r$

Is this enough to guarantee an expected utility representation? In other words, can we find utility numbers for each of the prizes $u(a)$, $u(b)$ $u(c)$ and $u(n)$ such that

$$p \succ q$$

if and only if

$$\begin{aligned} & p_a u(a) + p_b u(b) + p_c u(c) + p_n u(n) \\ > \quad & q_a u(a) + q_b u(b) + q_c u(c) + q_n u(n) \end{aligned}$$

The answer is no: we need one more axiom - which is called the independence axiom. To understand this axiom, think of the following question:

Question: Think of two different lotteries, p and q. Just for concreteness, let;s say that p is a 25% chance of winning the apple and a 75% chance of winning the banana, while q is a 75% chance of winning the apple and a 25% chance of winning the banana. Say you prefer the lottery p to the lottery q. Now I offer you the following choice between option 1 and 2

1. I flip a coin. If it comes up heads, then you get p. Otherwise you get the lottery that gives you the cockatiel for sure

2. I flip a coin. If it comes up heads, you get q. Otherwise you get the lottery that gives you the cockatiel for sure

Which do you prefer?

The independence axiom basically states that if you prefer p to q, then you have to prefer option 1 to option 2. This seems intuitively plausible. After all, in the choice between 1 and 2, then if the coin comes up tails, then you get the same thing in both cases. If it comes up heads then for 1 your get p and for 2 you get q. If you prefer p to q, then it seems natural that you should prefer 1 to 2. In fact, the independence axiom says slightly more than this:

Axiom 1 (The Independence Axiom) *Say a consumer prefers lottery p to lottery q. Then, for any other lottery r and number $0 < \alpha \leq 1$ they must prefer*

$$\alpha p + (1 - \alpha) r$$

to

$$\alpha q + (1 - \alpha) r$$

Does this axiom hold in practice? This has been one of the biggest question in behavioral economics and economic theory over the past 50 years. The answer (as usual) seems to be - quite a lot of the time yes, sometimes no. To see where it might fail, consider the following choices

A	100% chance of $1 million
B	89% chance of $1 million, 1% chance of nothing, 10% chance of $5 million
C	89% chance of nothing, 11% chance of $1 million
D	90% chance of nothing 10% of $5 Million

Ask yourself the following: Which do you prefer - A or B. Also, which do you prefer - C to D? Most people prefer A to B and D to C. As you will show for homework, this violates the independence axiom.

Be that as it may, we can show that if the independence axiom is satisfied (and one other technical condition, that we will not discuss)[17], then we can find utility numbers $u(a)$, $u(b)$ $u(c)$ and $u(n)$ such that our consumer behaves like a utility maximizer. A rigorous proof of this lies beyond the scope of this course, but I will sketch the procedure here

1. Find the best prize - in other words the prize such that getting that prize for sure is preferred to all other lotteries. Give that prize utility 1 (for convenience, let's say that apples are the best prize)

[17] It's a continuity condition, maths fans

2. Find the worst prize - in other words the prize such that all lotteries are preferred to getting that prize for sure. Give that prize utility 0 (for convenience, let's say that getting nothing is the worse prize)

3. For other prizes (e.g. bananas), find the probability λ such that the consumer is indifferent between getting apples with probability λ and nothing with probability $(1-\lambda)$, and bananas for sure. Let $u(b) = \lambda$. i.e.

$$\begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \sim u(b) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + (1 - u(b)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

4. Do the same for cockatiels, so

$$\begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \sim u(c) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + (1 - u(c)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

Now, we need to show that

$$p \succ q$$

if and only if

$$\begin{aligned} & p_a u(a) + p_b u(b) + p_c u(c) + p_n u(n) \\ > \quad & q_a u(a) + q_b u(b) + q_c u(c) + q_n u(n) \end{aligned}$$

To see why this is true, lets think of a simple example: again let's say that p is a 25% chance of winning the banana and a 75% chance of winning the cockatiel, while q is a 75% chance of winning the banana and a 25% chance of winning the cockatiel. Now, note that

$$\begin{aligned} p &= \begin{pmatrix} 0 \\ 0.25 \\ 0.75 \\ 0 \end{pmatrix} \\ &= 0.25 \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} + 0.75 \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \end{aligned}$$

But we know that

$$\begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \sim u(b) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + (1 - u(b)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

and

$$\begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \sim u(c) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + (1 - u(c)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

and so, by the independence axiom

$$p \sim 0.25\left(u(b)\begin{pmatrix}1\\0\\0\\0\end{pmatrix}+(1-u(b))\begin{pmatrix}0\\0\\0\\1\end{pmatrix}\right)$$

$$+0.75\left(u(c)\begin{pmatrix}1\\0\\0\\0\end{pmatrix}+(1-u(c))\begin{pmatrix}0\\0\\0\\1\end{pmatrix}\right)$$

$$= (0.25u(b)+0.75u(c))\begin{pmatrix}1\\0\\0\\0\end{pmatrix}+(1-0.25u(b)+0.75u(c))\begin{pmatrix}0\\0\\0\\1\end{pmatrix}$$

or p is indifferent to a mixture between the best and the worse lotteries, with the weight on the best lottery given by $0.25u(b) + 0.75u(c)$. But this is the expected utility of p. Similarly, q is indifferent to such a mixture, with the weight on the best lottery equal to the expected utility of q. In other words

$$q \sim (0.75u(b)+0.25u(c))\begin{pmatrix}1\\0\\0\\0\end{pmatrix}+(1-0.75u(b)+0.25u(c))\begin{pmatrix}0\\0\\0\\1\end{pmatrix}$$

Thus, if the expected utility of p is higher than the expected utility of q, this tells us that the lottery to which p is indifferent puts a higher weight on the best prize than does the lottery to which q is indifferent. Thus, p must be preferred to q

How does this relate back to the St. Petersburg paradox and the pauper? Well, what we were basically saying there is that it seemed sensible for the pauper to turn down a 'fair gamble': Even though the lottery ticket gave a higher average payoff than $499,999 for sure, they would still prefer the sure thing. We call such people **risk averse**

Definition 2 *A person is risk averse if they always prefer an amount x for sure to a lottery that* *has an expected value of x*

Are people risk averse? In general, yes, but not always (after all, people *do* by lottery tickets!)

Bernoulli's explanation was that risk aversion comes about if people have decreasing marginal utility of wealth. This may seem to be a bit of a mysterious statement. How are these two linked? Well, it turns out they they do indeed imply each other if the person is an expected utility maximizer. We will show this graphically now, and you will show it formally for homework.

In order to answer this question, we first need to remember what a utility function that exhibits decreasing marginal utility looks like. The answer should come quickly if you think about the work we have done on production: such a utility function will be concave, as shown in figure 36. Figure 37 shows how to calculate the expected utility of a gamble that gives an amount z with probability 0.5 and y with utility 0.5 - you can read it off the line segment that connects the two. However, for a concave function, this will always be below the utility of getting $0.5z + 0.5y$ for sure. Thus, an individual who exhibits decreasing marginal utility will have a concave utility function and so exhibit risk aversion.