Intermediate Microeconomics

A Handbook

Author: Robert Pettis

Institute: University of Texas at Arlington

Date: February 12, 2021

Version: 1.22

"I'm told it makes sense, but I seriously have my doubts."

- John Oliver, on Economics

Contents

Chapter Demand

1.1 Introduction

Learning Objectives

- *Utility*
- *Indifference Curves*
- *Budget Constraint*
- *Optimal Consumption Choices*
- *Consumer Demand*

We live in a world where resources are scarce.[1] Aside from natural resources (such as gold, lumber, etc.), this includes intangible resources, such as our time spent working and thus, our own income. Given this, individuals need to make decisions on how to best allocate their consumption of resources to make themselves as best off as they can possibly be. In economics, we call the satisfaction gained from such decisions utility.

Definition 1.1. Utility

Utility is a numeric value indicating the consumer's relative well-being. Higher utility indicates greater satisfaction than lower utility. ♣

Our goal, as rational people, is to maximize our own utility. With this goal in mind, if we are able to quantify utility, perhaps we can use it as a tool to learn more about our behavior. At first, this may seem silly. There is no way we can accurately come up with a number for how happy people are, let alone compare this number across people. . . But what we can observe are people's behavior in regard to consumption. From there, with some simple assumptions, we can try to quantify the optimal levels of consumption for a person. Our goal of this chapter will be to determine optimal consumption bundles for goods. Through this, we will calculate the demand for goods.

1.2 Utility

First, let's learn a bit more about utility. We will measure utility as the output of some function of some good x, $u(x) = f(x)$. Some assumptions we make about this function are as follows:

1. Consumers are utility maximizers
2. More is better
3. Diminishing marginal utility

[1] Thankfully, otherwise there wouldn't be a need for Economists!

4. We can compare all bundles of goods
5. Preferences are transitive

Let's look into each of these a bit more carefully.

Consumers Are Utility Maximizers Our *first assumption* is that consumers are utility maximizers. This assumption may give you some pause. You may ask about a number of cases where people may seem to not be rational. "What about drug addicts?" you may give as an example. My response to this would be that these agents ARE acting rationally. . . according to their utility function, that is. Perhaps their utility function puts the greatest weight on a short term consumption of the drug they are addicted to. If this is the case, then consuming more, with little regard for other consumption, is maximizing their utility function. Perhaps they would be better off if they did not use drugs...but this would require a change in their utility function. A person checking in to rehabilitation, for example, is trying to change their utility function. With this view, this particular assumption does not seem like much of a stretch.

More Is Better Our *second assumption* is that more is better in regard to consumption. In other words, I will always be better off if I consume more of a good. This assumption is somewhat toned down by our next assumption.

Diminishing Marginal Utility While the previous assumption may state that we will always be better off consuming more, our *third assumption* states that the amount of new utility we get for consuming one more unit of a good, or marginal utility, decreases the more of the good we consume.

Definition 1.2. Marginal Utility

The Marginal Utility of consuming a particular good [undergoing a particular activity] is the amount of utility gained by the consumer gained by consuming one more unit of the good [performing one more unit of the activity]. ♣

Let's explain this concept with the example in Table 1.1 below. Imagine a person who gets their utility solely from chocolate. Suppose that this person eats one piece of a chocolate bar. This affords them 70 units of utility.[2] The amount of utility gained from consuming no chocolate to consuming one unit of chocolate is 70. Thus, we would say that the marginal utility at one unit of chocolate is 70, or $MU_{C=1} = 70$. We may also notice that the amount of utility gained from additional consumption is always positive, thus we do not violate our "More Is Better" assumption. However, the amount of happiness we get for each additional bite diminishes. Has this ever happened to you? Perhaps the first bite of food is satisfying, but by the time you are full, the dish may not taste as good to you as it did before.

[2] In general, we either say that utility is unitless or in a fictional unit called 'utils.'

Chocolate Consumed	Total Utility	Marginal Utility
0	0	-
1	70	70
2	80	10
3	85	5
4	88	3

Table 1.1: Bites of chocolate consumed and the associated gain in utility.

Technical Explanation 1.1. Marginal Utility

We can say then, that the marginal utility of consuming a good is the rate of change of utility from consuming that good. If I had some utility function that expressed my utility as a function of my consumption, I could find the rate of change of that function, and that would be my marginal utility. We could do this by finding the first derivative of the function: $MU_A = \frac{\partial U(x_A)}{\partial X_A} = f'(x_A)$. ♡

Comparability of Goods In our analysis of utility, we must be able to compare all bundles of goods. If I have some choice between two bundles of goods, Bundle A and Bundle B, our *fourth assumption* is that a consumer would either: prefer A to B; prefer B to A; or be indifferent between them.

Technical Explanation 1.2. Preference Notation

We could also say this mathematically. I could say that an individual prefers bundle A to B in the following way: $A \succ B$, *where the* $\succ$ *symbol indicates preference.* ♡

Transitive Preferences If a consumer prefers bundle X to bundle Y ($X \succ Y$) and bundle Y to bundle Z ($Y \succ Z$), we make our *fifth assumption* that he/she prefers X to Z ($X \succ Z$). In a simplified example, if I prefer pizza to apples, and I prefer apples over black-eyed peas, then it should make sense that I prefer pizza to black-eyed peas as well.

1.3 Indifference Curves

In this course, we will primarily investigate relationships between two goods.[3] In these cases, our utility becomes a function of two goods, e.g. $u(X, Y) = f(X, Y)$. This is represented graphically in figure 1.1. Notice that our assumptions still hold. For example, the utility value (labeled "U" on the z-axis) increases for each increase in either X or Y. We also notice from the shape of the function that the more a single good is consumed, the less marginal utility it gives us. Perhaps this person would be better off not consuming an extreme amount of one good, but an average amount of both goods.

[3] If there was only one good, it wouldn't be an interesting analysis. If we use more than two, graphical explanations become increasingly difficult.

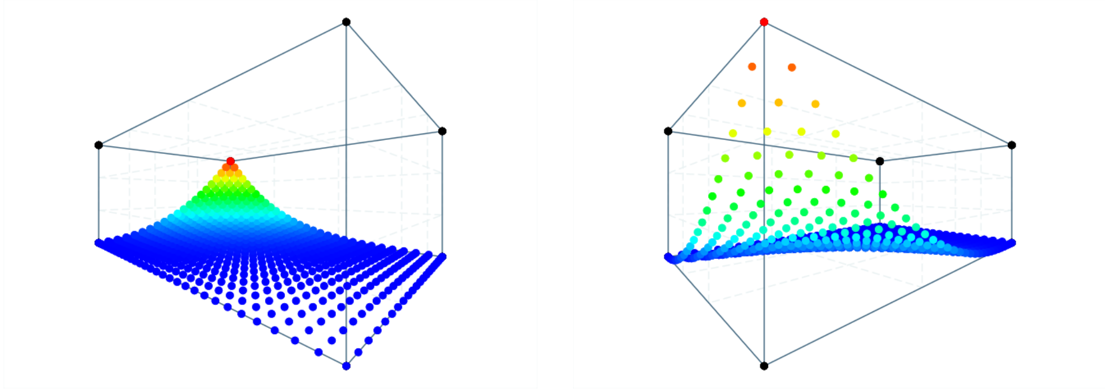

(a) (b)

Figure 1.1: Different Perspectives on the 3D Graph of a Utility Function with Two Variables. Utility is on the z-axis, Apples and Bananas consumed are on the x and y axis, respectively.

Does this figure resemble anything to you? I rather think it looks like the contour map of a hill, such as that in Figure 1.2. Notice that on our contour map, levels of identical elevation are marked with a line, e.g. 1500 feet, or closer to the top, 2100 feet. This is just like the lines shown on the utility graph in Figure 1.1. The lines there represent identical amounts of utility. Put another way, on those lines, a consumer would be indifferent between bundles along the line because they would all grant the same level of utility. Because of this indifference, these curves are called Indifference Curves.

Definition 1.3. Indifference Curve

A graphical representation of all bundles that a consumer is indifferent between. Indifference curves are derived from an individual's preferences. ♣

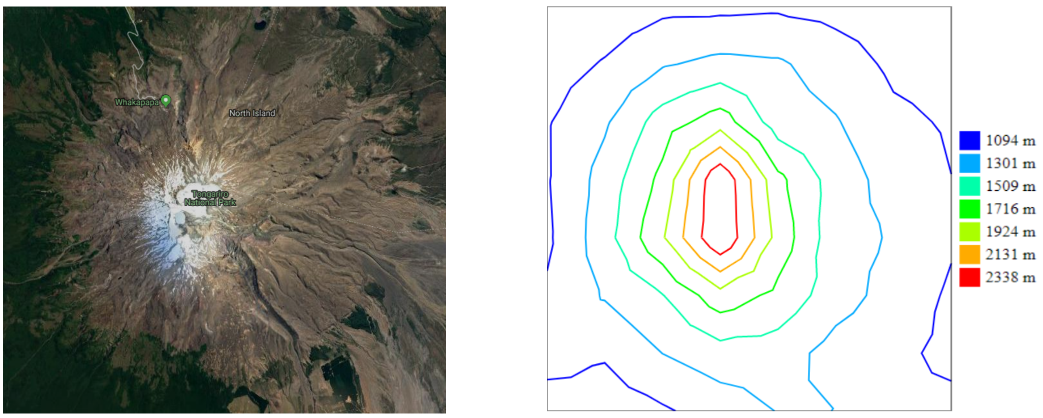

(a) Overhead Google Maps Image of Mount Ruapehu in New Zealand, known in popular culture as one of the volcanos used to film the Mount Doom portions of the Lord of the Rings films.

(b) The Contour Map of Mount Ruapehu.

Figure 1.2: Juxtaposition of Mount Ruapehu (Mount Doom) and Its Contour Map

Most often, for simplicity, we will draw these curves from an overhead view. In terms of comparing it to the contour map, it will be as if we were flying in a helicopter right above them. If we did this to Figure 1.1, it would look like Figure 1.3.

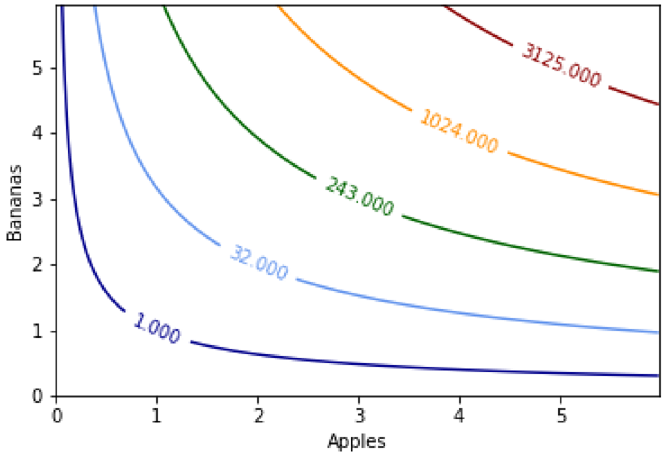

Figure 1.3: The indifference curve: the contour map of utility.

Technical Explanation 1.3. How To Graph An Indifference Curve

If an indifference curve is identified with a given level of utility, a constant, setting the utility function equal to that desired level of utility and solving for the y-axis variable would allow us to graph the curve.

Example1.1 *Suppose we are given the following preferences over two goods, x_1 and x_2:*

$$u(x_1, x_2) = x_1x_2$$

If we wanted to graph an indifference curve with a utility value of some constant, k, then what we want is $k = x_1x_2$. By convention, we will graph x_2 on the y-axis, and will thus solve for x_2. Doing so gives us $x_2 = \frac{k}{x_1}$, the equation we can use to graph the indifference curve.

1.3.1 Types of Indifference Curves

Cobb-Douglas

The most common type of utility function we will deal with is called a Cobb-Douglas utility function. It is of the form:

$$u(x_1, x_2) = Ax_1^{\alpha}x_2^{\beta} \tag{1.1}$$

where the coefficient A is some constant, and the greek letters α and β stand in for two other constants that often, but not always, sum to 1. We have seen the indifference curves for these types of preferences already, in Figure 1.3 as well as the three dimensional version in Figure 1.2.

Let us return to thinking of the indifference curves representing bundles of goods that an individual is indifferent between for a given level of utility. Now, using Figure 1.3 as a reference, let us ask ourselves what is implied by that at different levels of consumption for our goods. To start, what do we notice about the shape of the indifference curves at lower levels of consumption of the x-axis good? It appears to be relatively much more steep of a slope there. What does this imply? It means that this individual would be willing to give up large amounts of the y-axis good in order to consume some of the x-axis good. This effect diminishes as we move to the right on the x-axis and in fact, at extreme values in that direction, we see a similar phenomenon: the flatness of the curve on the x-axis at large values of x implies the individual would be willing to give up a lot of the x-axis good in order to get a relatively smaller amount of the y-axis good. It appears that this individual prefers averages to extremes. This is a common theme for Cobb-Douglas functions. The slope of the indifference curves is referred to as the Marginal Rate of Substitution.

Definition 1.4. Marginal Rate of Substitution

The Marginal Rate of Substitution (MRS) is the slope of the indifference curves. As such, it represents, at a given point, how much the consumer is willing to give up of the y-axis good in return for a unit of the x-axis good. ♣

Key Point: The MRS is equal to the ratio of marginal utilities, with a negative sign:

$$MRS_{x,y} = -\frac{MU_x}{MU_y} \tag{1.2}$$

Technical Explanation 1.4. Marginal Utility With Multiple Goods

Earlier, we discussed that the Marginal Utility of a good is the extra utility gained from consumption of one unit of another good and that we could take the derivative of the utility function with respect to the good in order to find the marginal utility. The same applies now that we have added a second good. The difference is that now we need to take the partial derivative with respect to the variable we want the marginal utility for.

$$MU_1 = \frac{\Delta U}{\Delta x_1} = \frac{\partial U}{\partial x_1}$$

As a reminder, when taking the partial derivative, we follow the same procedure as with a derivative with only one variable, except that we pretend the other variable is a constant. For a Cobb-Douglas utility function, the marginal utilities of both goods are:

$$MU_1 = \frac{\partial U}{\partial x_1} = A\alpha x_1^{\alpha-1} x_2^{\beta}$$

and

$$MU_2 = \frac{\partial U}{\partial x_2} = A\beta x_1^{\alpha} x_2^{\beta-1}$$

♡

Technical Explanation 1.5. MRS = the Ratio of Marginal Utilities

The mathematical definition of marginal utility for a two variable utility function is: $MU_x = \frac{\Delta U}{\Delta x}$ *and* $MU_y = \frac{\Delta U}{\Delta y}$. *Rearranging, we get:* $MU_x \Delta X = \Delta U$ *and* $MU_y \Delta y = \Delta U$. *If the combination of these changes leaves us indifferent (*$\Delta U = 0$*) as we would be on an indifference curve, then* $MU_x \Delta X + MU_y \Delta y = 0$.

Rearranging, we have $\frac{\Delta y}{\Delta x} = -\frac{MU_x}{MU_y}$. $\frac{\Delta y}{\Delta x}$ *is a way of expressing the rate of change of y given a change in x, which is another way of saying the slope (in this case, of an indifference curve), which we have named the MRS. Thus, the MRS equals the ratio of marginal utilities.* ♡

1.4 The Budget Constraint

If more is better, then why don't we just consume constantly to keep raising our utility? The answer is that since resources are limited, they cost money to get and even the most wealthy have limits to what they can afford. These limitations are described by a budget constraint.

Definition 1.5. Budget Constraint

With a budget constraint, a person's income (denoted M) must not be exceeded by the cost of the goods purchased. In a world where an individual is only concerned with purchasing two goods, x_1 *and* x_2 *at prices* p_1 *and* p_2 *respectively, the constraint is represented below in Equation 1.3 as:*

$$M \geq p_1x_1 + p_2x_2 \tag{1.3}$$

♣

In this course, we make a simplification on the budget constraint. We assume that individuals cannot save. Imagine that their money will expire after this time period. If individuals' income will not be good next period, then is there any point to not spending exactly what they have? With this in mind, we can imagine that this would lead to individuals spending all of their money on x_1 and x_2:

$$M = p_1x_1 + p_2x_2 \tag{1.4}$$

As we are wont to do in economics, we will make a graph. We make a point to graph because, as a tool, it can help us learn more about our subject. If we graph in two dimensional space (since we have two goods), we will need to solve for good we wish to be on the y-axis, just as we did with the indifference curves. This is solved for us in Equation 1.5 and the result is graphed in Figure 1.4

$$\begin{aligned} M &= p_1x_1 + p_2x_2 \\ p_2x_2 &= M - p_1x_1 \\ x_2 &= \frac{M}{p_2} - \frac{p_1}{p_2}x_1 \end{aligned} \tag{1.5}$$

Notice the following key facts about this equation:

- The y-intercept is simply the amount of good x_2 (with x_2 being the good represented by the y-axis) that the consumer can afford if they spent their entire budget on x_2, $\frac{M}{p_2}$.
- The x-intercept is analogous, with an intercept of $\frac{M}{p_1}$.
- The slope is constant and is the ratio of the prices of the two goods, $-\frac{p_1}{p_2}$.

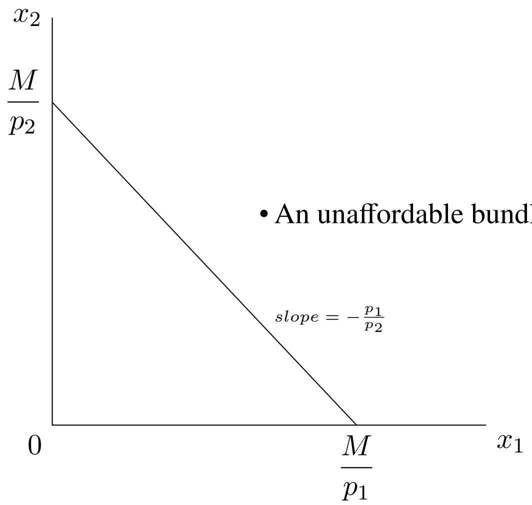

Figure 1.4: The Budget Constraint

As the income, M, is the numerator in both the x and y intercepts, an increase in income would result in a parallel shift of the budget constraint (see Figure 1.5a). This should make intuitive sense. If we have more money, we can spend it on more of everything. If, on the other hand, the price of one of the goods changes, it would result in a change of the intercept for that good only. For example, in Figure 1.5b, the price of good 1 has increased, decreasing the possible amount of good 1 we can afford. This is visually represented by the budget constraint pivoting around the y-intercept (the y-intercept does not change...and why should it? The change in the price of good one has not impacted the total number of good 2 it is possible to purchase.).

1.5 The Optimal Bundle

Superimposing the budget constraint and the indifference curves can help us solve the mystery of how to maximize utility under the constraint of our budget. Figure 1.6 illustrates the budget constraint and an indifference curve that is just tangent to it. It turns out that the point of tangency represents the optimal consumption of the x-axis good and the y-axis good. Let us prove this by considering cases where our consumption bundle is not (y^*,x^*). First, more is better right? But if we are at (y^*,x^*), we cannot increase our consumption as it would be greater than what we could afford. Look at Figure 1.6 again. Is there any point below or on the budget constraint that could get us on a higher utility curve? The answer is no. Any movement away from (y^*,x^*) in the set of affordable points would result in us being on a lower indifference curve. This is clearly not optimal. Thus, the optimal point is (y^*,x^*). For a more mathematical

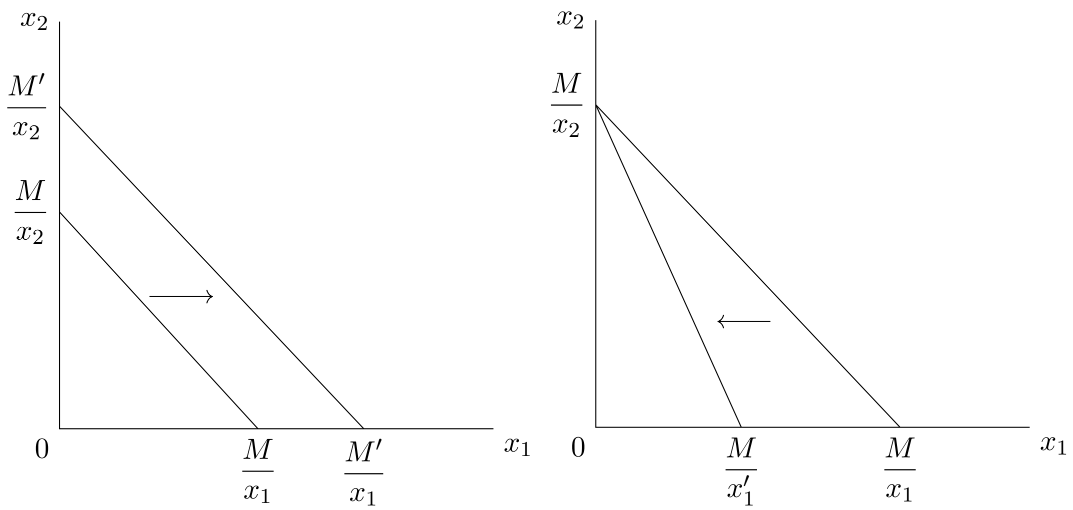

(a) The effect of an increase in income on the budget constraint

(b) The effect of an increase in the price of good 1 on the budget constraint

Figure 1.5: Possible Changes to a Budget Constraint

explanation of why this point is optimal, see Appendix A.2.4.

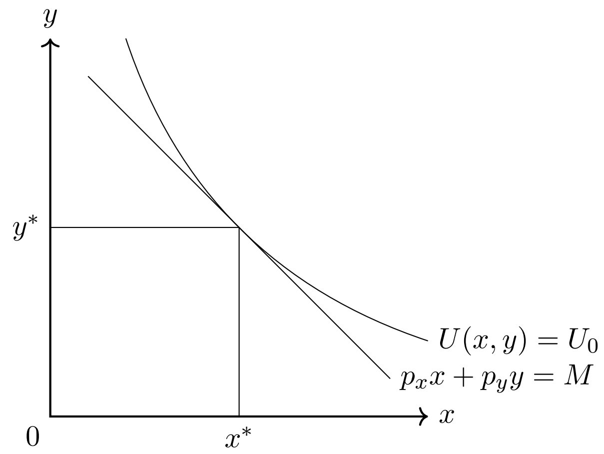

Figure 1.6: The Tangency of the Budget Constraint and the Indifference Curves. Optimal levels of consumption are indicated with an asterisk(*).

Our next step on the journey to calculate the optimal values of our goods lies in the mathematical fact that at the point of tangency, the slopes are equal. Luckily for us, we already know how to find the slopes of both of these curves. Equations 1.2 and 1.5 give us the slopes of the indifference curve and the budget constraint respectively. Thus, we have our tangency condition in Equation 1.6.[4]

$$\frac{MU_1}{MU_2} = \frac{p_1}{p_2} \tag{1.6}$$

We can also rearrange this formula as below in Equation 1.7. Each side of the equation represents the new utility you would get for consuming another unit of the good per cost of that good. In other words, the "bang for buck" for consuming that good. Think about what would

[4]Notice that there are no negative signs anymore. Both slopes are negative, so they cancelled out.

happen if this equation were not true. If your bang for buck was better for one of the two goods, would you not spend less on the good that did not provide as much new utility per dollar as the other good until this equality was achieved?

$$\frac{MU_1}{p_1} = \frac{MU_2}{p_2} \tag{1.7}$$

Example1.2 Doris's preferences over apples, A and bananas, B are represented by the Cobb-Douglas utility function; $u(x_A; x_B) = A^2B^3$. What is her optimal bundle if her income is \$500, and bananas are \$1 a pound and apples are \$2 a pound?

The following steps will guide us through the process.

1. Find the Marginal Utilities of Apples and Bananas.
   - First, let's find MU_A:
   $$MU_A = \frac{\partial U}{\partial A} = 2AB^3 \tag{1}$$
   - Next, let's find MU_B:
   $$MU_B = \frac{\partial U}{\partial B} = 3A^2B^2 \tag{2}$$
   - Now we have the MRS by taking the ratio of the results so that MRS=(1)/(2):
   $$MRS = \frac{2AB^3}{3A^2B^2} = \frac{2B}{3A} \tag{3}$$
   - Now that we have the MRS, let us set it equal to the price ratio in order to satisfy the tangency condition. Since we are satisfying the condition, if we solve for one of the variables, we would find the optimal consumption. In this case, I solve for B.
   $$\begin{aligned} \frac{2B}{3A} &= \frac{2}{1} \\ B &= 3A \end{aligned} \tag{4}$$
   - So we have found the optimal level of B, but it is a function of the other variable A. It would be nice to have an actual value of B that did not depend on A. Fortunately, we have another trick up our sleeve. By definition, our budget constraint is binding, which means it must be true at the optimal levels of A and B. If we substitute Result (4) into the budget constraint, we can remove a variable from the equation. Let me demonstrate:
   $$\begin{aligned} M &= P_AA + P_BB && \text{Start with budget constraint} \\ 500 &= 2A + B && \text{Sub in the income and prices} \\ 500 &= 2A + 3A && \text{Sub in Result (4)} \\ 500 &= 5A \\ A^* &= 100 \end{aligned} \tag{5}$$
   - We have successfully found the optimal level of apples for this person to consume! But what about bananas? Recall that we have already found a function that gives the

optimal level of bananas as a function of apples (see Result (4)). If we evaluate this function at the optimal level of bananas, it will give us the level of bananas we want at our optimal bundle!

$$\begin{aligned} B &= 3A && \text{Start with Result (4)} \\ B &= 3 * 100 && \text{Sub in } A^* \text{ from Result (5)} \\ B^* &= 300 \end{aligned} \tag{6}$$

Thus, our optimal bundle of Apples and Bananas is $(A^*, B^*) = (100, 300)$.

1.5.1 Demand

In the Section 1.5, we learned how to find the optimal level of consumption of a good for an individual given some parameters (income, price of the good, price of the other good). But what if the price for a good changes? How will the amount demanded by the individual change with it?

Exercise1.1 Repeat Example 1.2 with a different price for A and B: $p_A = 1$ and $p_B = 3$. How does the optimal amount this individual demands change as the price changes? Graph the budget constraints at $p_A = \{1, 2, 3\}$ and mark the optimal points for each budget constraint.

If we repeated Example 1.2 with a generic price indicator, p_A, the result would be $A = \frac{2M}{5p_A}$. In other words, we can see how the demand for apples can change as a function of the price. We call this function the Demand Function. Keeping income at \$500, the demand function becomes $A = \frac{200}{p_A}$. In order to visually examine this relationship, we must first note that in Economics, we always graph the relationship between price and quantity of goods with price on the y-axis and quantity on the x-axis. This means we will have to solve for p_A.

$$\begin{aligned} A &= \frac{200}{p_A} \\ p_A &= \frac{200}{A} \end{aligned}$$

This is graphed in Figure 1.7 below. Notice that there is an inverse relationship between price and quantity demanded. This should make intuitive sense. In general, if something is more expensive, I will not buy as much of it.[5]

Often, we are interested in the market demand, or demand of all individuals wishing to participate in the market for some good. We can find the market demand by aggregating the total quantities at a given price. If we assume individuals are identical, then we can simplify the process. Equation 1.8 applies this process to our running example and Figure 1.8 graphs the market demand for two identical individuals.

[5]There are exceptions to this rule. What if price for this good was an indicator of quality, for example? Of course, this would mean we would have differentiated goods. What about conspicuous consumption?

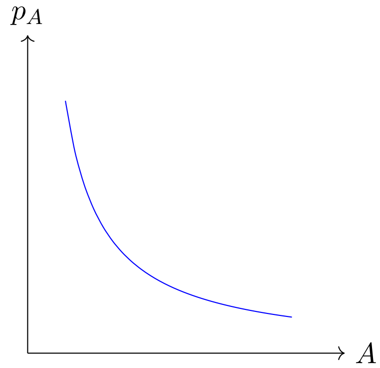

Figure 1.7: Inverse Demand Curve for Apples

$$\begin{aligned} \sum_{i=1}^{n} A &= \sum_{i=1}^{n} \frac{200}{p_A} \\ A(Market) &= \frac{200n}{p_A} \end{aligned} \tag{1.8}$$

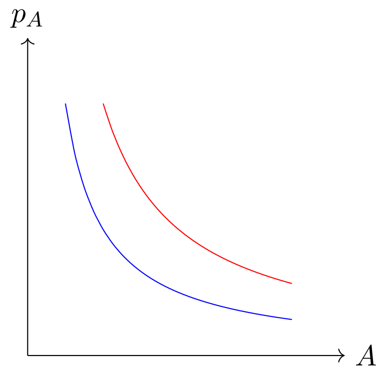

Figure 1.8: Market Inverse Demand Curve for Apples

So how can these curves change over time? Clearly anything changing that went into the optimization process in the first place can have an impact (income, the price of other goods, preferences), but now we can see that the number of people that participate in the market can affect the market demand curve as well.

✍ **Exercise1.2** We just mentioned a list of things that could affect the shape of the indifference curves. As an exercise, go through each item on the list and write down how each one would change the appearance of an indifference curve.

We can learn if we are dealing with a normal good or an inferior good based on how the consumption changes with a change in income. We do this by drawing a line through the optimal bundles as income increases. This line is called the Income Expansion Path. If the path shows that as income increases, more of the good is demanded, as in Figure 1.9, the good is normal. If the path shows that as income increases less of a good is demanded, as in Figure 1.10, the good is inferior.

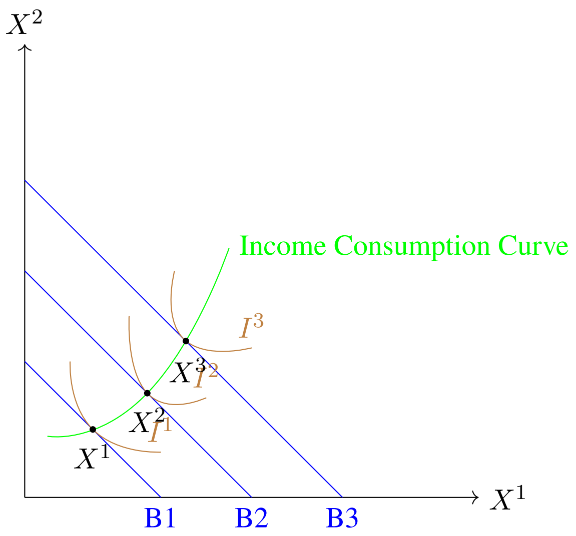

Figure 1.9: The Income Expansion Path of A Normal Good

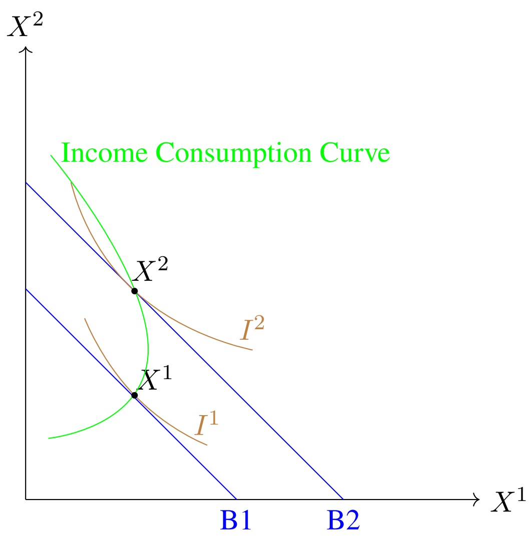

Figure 1.10: The Income Expansion Path of An Inferior Good

Exercise1.3 List 5 goods that are normal and 5 goods that are inferior. What makes them so?

1.5.2 Substitution and Income Effects

In Section 1.5.1, we learned how quantity demanded changes as price changes. But this change can be decomposed into two effects. The first effect, the income effect, is due to the change in real income that a price change causes. The second effect, the substitution effect, is due to the fact that even with an income adjustment to return to the original indifference curve, we would still choose a different bundle because of the change in relative prices. This is most easily shown graphically. Figure 1.11 illustrates the process. Assume we are originally optimal at point x^A. Then, the price of good A decreases causing the budget constraint to shift from A to C. Now we have a new optimal bundle at x^C. But let's assume, hypothetically, that we had our income reduced so that budget constraint would be just tangent to our original indifference curve. This would put us at x^B. The difference between the original level of consumption and this hypothetical level of consumption is the Substitution effect. The difference between the hypothetical level of consumption and the new realized level of consumption is the income effect.

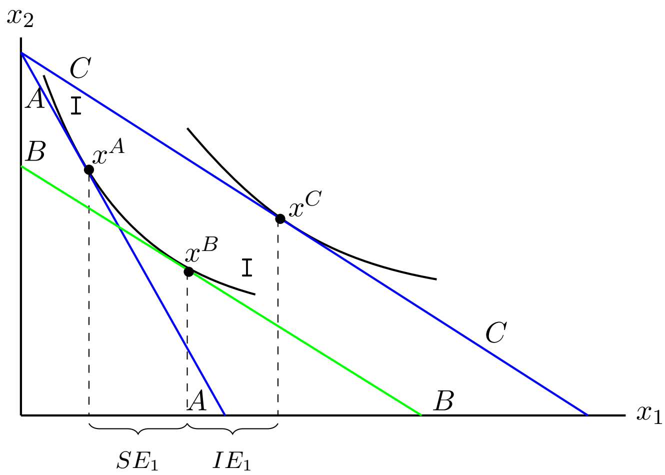

Figure 1.11: Hicks Decomposition

1.6 Other Utility Functions

There are any number of other types of utility functions we might see. For this course, we will only touch on a couple of others: Perfect Complements and Perfect Substitutes, which are illustrated in Figure 1.12. These preferences are two opposing extremes. Details on each are below.

1.6.1 Perfect Substitutes

In a world of perfect substitutes, a consumer is perfectly indifferent between any linear combination of two goods. In the pictured example, a person is perfectly indifferent between 1 dime or two nickels, 2 dimes and 4 nickels, and so on. So how does a consumer make an optimal choice? First, note that these preferences are linear and are in the form:

$$u(x_1, x_2) = ax_1 + bx_2$$

Now, let's try to use the tangency condition, as before:

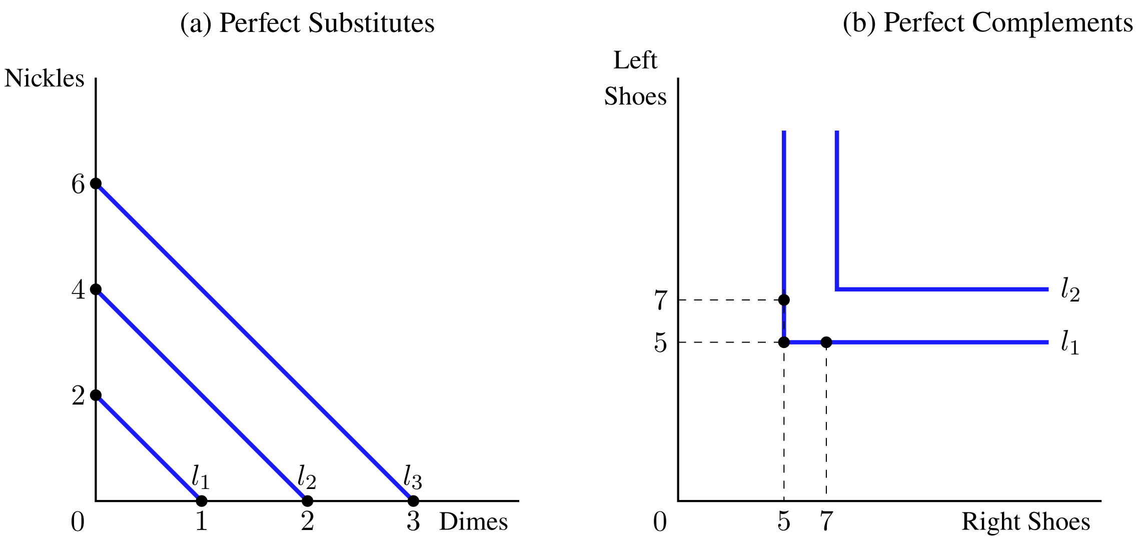

Figure 1.12: Additional Indifference Curves

$$MRS = -\frac{MU_1}{MU_2} = -\frac{a}{b}$$

Problem: $\frac{a}{b}$ is a number. We cannot set this equal to the price ratio. Solution? Well, if

$$\frac{MU_1}{MU_2} = \frac{a}{b} > \frac{p_1}{p_2}$$

or

$$\frac{MU_1}{p_1} > \frac{MU_2}{p_2}$$

then our Marginal Utility per dollar is better with good 1, so we should spend all of our money on good one, and no money on good 2, i.e. $x_1^* = (p_1, p_2, M) = \frac{M}{p_1}$, $x_2^* = 0$. This is known as a "corner solution" because they only consume one good. If instead, $\frac{MU_1}{p_1} < \frac{MU_2}{p_2}$, then $x_2^* = (p_1, p_2, M) = \frac{M}{p_2}$, $x_1^* = 0$.

Example1.3 Colley's utility function over apples, x_A, and bananas, x_B, is given by

$$u(A, B) = 2x_A + 2x_B$$

. If the price of apples is \$1 and the price of bananas is \$3, what is Colley's optimal consumption bundle when his income is \$100?

The first step is to find the marginal utilities for apples and bananas:

- $MU_A = 2$
- $MU_B = 2$

Now we need to ask ourselves which good gives us more "bang for the buck" ($\frac{MU}{P}$).

- Apples "bang for buck": $\frac{2}{1}$
- Bananas "bang for buck": $\frac{2}{3}$

Since $\frac{2}{1} > \frac{2}{3}$, apples provide the most "bang for buck" and should all funds should be spent on apples and none on bananas. Therefore, $x_A^* = \frac{M}{p_A} = \frac{100}{1} = 100$ and $x_B^* = 0$.

Exercise1.4 In Example 1.3, if the price of apples were to increase to \$2, what would happen to his banana consumption?

1.6.2 Perfect Complements

Perfect Complement preferences, also known as "Leontief" preferences or "Constant-Proportion" preferences have "L" shaped indifference curves. This is a result of the form the underlying utility function takes, which is that of a min function: $u(x, y) = \min\{ax_1, bx_2\}$. Optimiality is achieved at the "kink" of the function.[6][7] At this kink, $ax_1 = bx_2$.[8] Thus, we have a new optimiality condition, which we can substitute into our budget constraint, just as in the Cobb-Douglas case. Here, though, we skip some of the steps of the Cobb-Douglas function, and skip directly to the optimality condition.

Example1.4 Jeanelle's preferences over x and y are given by: $u(x, y) = \min\{2x, y\}$. Derive her demand for x and y.

Since Jeanelle has "perfect complement" preferences, as we can tell from the "min" function, we know that at the optimal point,

$$2x = y \tag{1}$$

As with Cobb-Douglas preferences, we can exploit the fact that the budget constraint is binding and substitute the above expression for y into it.

$$\begin{aligned} M &= p_x x + p_y y \\ M &= p_x x + p_y 2x && \text{Substitute Result (1) into constraint} \\ M &= x(p_x + 2p_y) && \text{Collect like terms} \\ x^* &= \frac{M}{p_x + 2p_y} \end{aligned} \tag{2}$$

And from Result (1), we know that $y = 2X$ at the optimal bundle; therefore,

$$y^* = 2(x^*) = \frac{2M}{p_x + 2p_y} \tag{3}$$

1.7 Extensions

Our earlier models were very restrictive. Here, we relax some assumptions to learn more about our agents under more realistic circumstances.

---

[6]Think about why this is the case. What would happen if we increased the consumption of only one of the two goods. Would the utility change? We would be better off spending proportionally on both goods, because then our utility would actually increase.

[7]The math adept among us may recognize that this kink means that this function is not differentiable. We will therefore need to use the following approach.

[8]You should prove this to yourself!

1.7.1 Inter-temporal Utility

Prior to this, we did not allow agents to save - they spent all of their income in the period in which they received it. Now we will gently relax this restriction by assuming that an agent can save their income for one period. Assume that your income is \$10 today and \$10 tomorrow. You could spend all of today's endowment now, and all of tomorrow's endowment tomorrow (spending all of your income in the period you receive it is called the **endowment point**)... or you could choose to save some (or all) of your earnings in the first period. Figure 1.13 illustrates the possible combinations. \$20 is the y-intercept because if we save all of our first period funds, we would spend \$20 in period 2 and \$0 in period 1. If we spend everything as we get it (no saving), we would spend \$10 in period 1 and \$10 in period 2 (the endowment point).

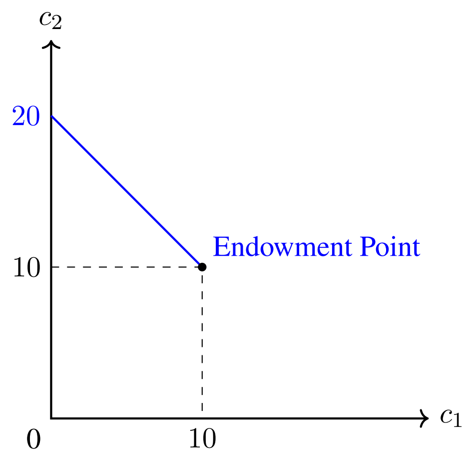

Figure 1.13: Budget Constraint If Agent Can Save For One Period

Now let us further assume that you can earn interest on this money. Using 10% as an example, we would be able to receive a maximum of $10 + 10(1 + .1) = 21$ in the second period, if we saved our entire income from period 1. This change is reflected in Figure 1.14.

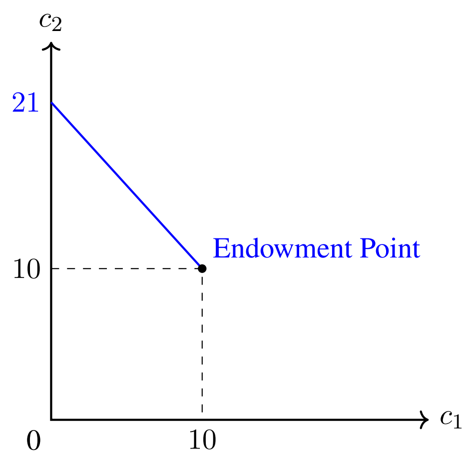

Figure 1.14: Budget Constraint If Agent Can Save For One Period with $r = 10\%$

Now we are feeling pretty good about relaxing our restrictions, so let's add another wrench into the gears. Now assume that we can borrow as well, so long as we pay it back with interest by the end of the next period. We will use the same rate for simplicity. If we borrowed against our upcoming income of $10, we could take out a loan of $\frac{10}{(1+.1)} = 9.09$. Think about why this must be so. If collecting on savings earns us interest of $P \times (1+r)$, then by doing the "opposite" (taking out a loan), we should divide instead of multiply by 1+r. Keeping in mind our income in period 1 as well, if we do this, we could possibly spend $10 + 9.09 = 19.09$ in the first period. The updated constraint is illustrated in Figure 1.15

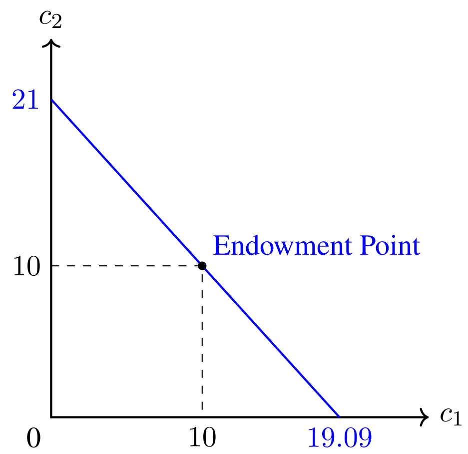

Figure 1.15: Budget Constraint If Agent Can Save For One Period with $r = 10\%$

Note that the slope of the budget constraint is constant. This is because we assumed that both borrowers (those that spend to the right of the endowment point in period 1) and savers (those that spend to the left of the endowment point in period 1) borrow and save at the same rate, r. Thus, the "price"(benefit) in terms of c_2 of borrowing(saving) is the constant -(1+r), the slope of the budget constraint. More generally,

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

where: c_1 and c_2 are consumption in periods 1 and 2, and m_1 and m_2 represent income in periods 1 and 2. Note that the term $(m_1 - c_1)$ represents the "savings" in period one, which could be negative if we are a borrower. Rewriting this to reflect slope-intercept form, we have Equation 1.9 below which is graphed in Figure 1.16.

$$c_2 = \underbrace{m_1(1+r) + m_2}_{\text{intercept (max } c_2 \text{ consumption)}} \overbrace{-(1+r)}^{\text{slope}} c_1 \tag{1.9}$$

When the interest rate changes, the budget constraint will shift about the endowment point. Think about why this is: at the endowment point, we neither save nor borrow, so would a change in the interest rate directly affect us? But if we are a borrower and the rate increases, we would not be able to borrow as much. In fact, we may even switch to being a saver, depending on our preferences. If we are a saver, on the other hand, we would have even more to spend in period 2 if the rate increases. This type of shift is shown in Figure 1.17.

We will continue relaxing our restrictions by now assuming that price levels may change between periods. In other words, we consider inflation. We define inflation as the rate of change

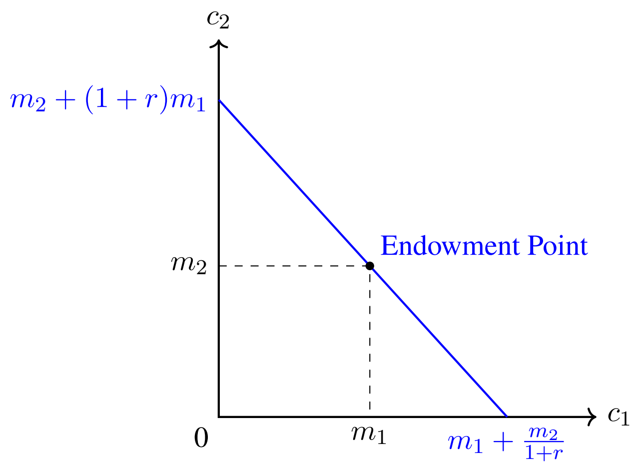

Figure 1.16: General Budget Constraint If Agent Can Save and Borrow

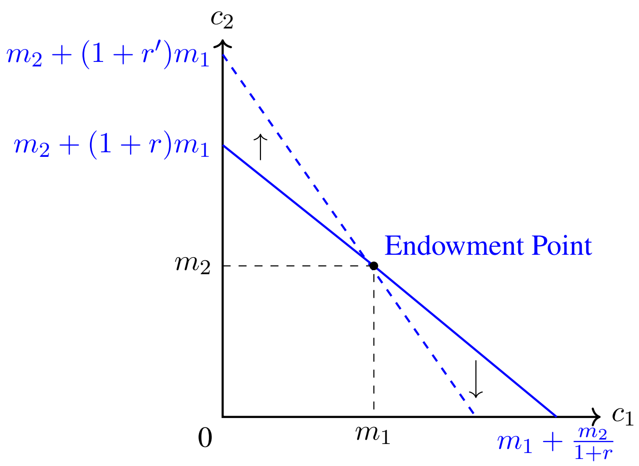

Figure 1.17: Change In Budget Constraint Due To Increase In Interest Rate

of price levels, $\pi = \dfrac{P_2 - P_1}{P_1}$. When we solve the budget constraint for c_2, what we are finding is the purchasing power of c_2 in terms of c_1. In the presence of inflation, this purchasing power must be discounted by the inflation to reflect the reduction in the ability to purchase in real terms, resulting in the following:

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

Rearranging this to match slope-intercept form, we have:

$$c_2 = \underbrace{m_1 \frac{1+r}{1+\pi} + \frac{m_2}{1+\pi}}_{\text{y-intercept}} - \overbrace{\frac{1+r}{1+\pi}}^{Slope} c_1 \tag{1.10}$$

Notice that Equation 1.9 is a special case of Equation 1.10 when $\pi = 0$.

Example1.5 The process for solving for the optimal bundle is identical to the methods discussed in prior sections with a few exceptions: instead of choosing between bundles of two different goods, the consumer is choosing between consumption in different time periods of c_i worth of money, and the budget constraint's shape is more complicated, as shown above. But these facts do not overly complicate our calculations. As an example, let us assume the following information: our individual of interest has preferences over consuming in period 1 and 2 of $U(c_1, c_2) = 400ln(c_1) + c_2$. This person also has an income of $240 in period 1 and $320 in period 2. The interest rate is $\frac{1}{3}$. There is no inflation. Our procedure is the same as with the Cobb-Douglas example in Example 1.2. We will set Marginal Rate of Substitution equal to the slope of the budget constraint. First, we find the MRS:

$$MU_1 = \frac{400}{c_1}$$

and

$$MU_2 = 1$$

Thus, the $MRS = -\frac{MU_1}{MU_2} = -\frac{400}{c_1}$. Now we need to set this equal to the slope of the budget constraint. Inflation is zero, so our budget constraint simplifies to that in Equation 1.9, giving us a slope of $-(1 + r)$. Setting this equal to the MRS, we have:

$$\frac{400}{c_1} = (1 + r)$$

or

$$c_1^* = \frac{400}{1 + r} = \frac{400}{1 + \frac{1}{3}} = 300$$

So this individual will consume $300 worth of goods in the first period. This amount is greater than their income in period 1, so this person must be a borrower. As before, we can substitute this value into another equation in order to get the value of consumption of the other good, in this case-consumption in period 2. Subbing 300 in for c_1 into the budget constraint gives us:

$$c_2^* = (\frac{4}{3})240 + 320 - 300(\frac{4}{3}) = 240$$

1.7.2 Utility Under Uncertainty

So far, we have dealt with our income being given. Now we will consider uncertainty in terms of outcomes, including income, and seek to learn about the risk profiles of consumers. To illustrate this, consider Figure 1.18. This figure graphs consumption, c, on the x-axis, and the resulting utility, u(c), on the y-axis. Let us assume that there are two states of the world. In state 1, the consumer will get an income of x. In state 2, the consumer will receive an income of y. If the probability of state 1 occurring was 100%, then the consumer would receive x with certainty and would gain utility of u(x). If the probability of being in state 2 was 100%, then the consumer would receive y with certainty and would gain utility of u(y).

But, what if we were unsure of the probabilities of the two events occurring? First, let's review two key aspects of probability:

1. The probabilities all sum to 1: $\sum p_i = 1$
2. Each probability must be between zero and one: $0 \leq p_i \leq 1, \forall_i$.

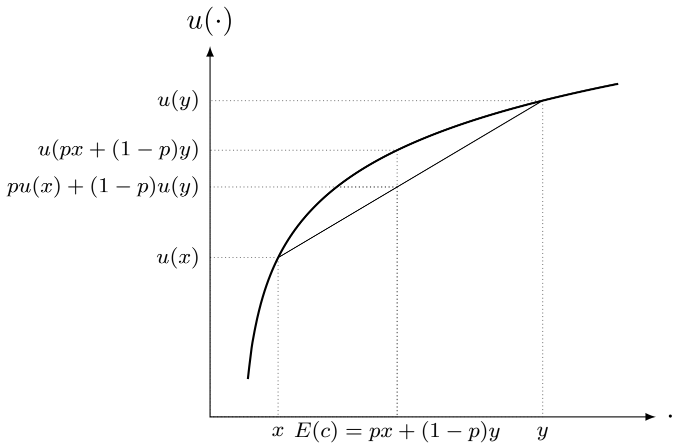

Figure 1.18: Risk Aversion

To expand more on our notation, since all probabilities must sum to 1, if we have only two states of the world occurring, we can denote the probability of state 1 occurring as p and the probability of state 2 occurring as (1-p). Prove to yourself that this must be the case. In this situation, the expected level of utility can be expressed as:

$$E(u(c)) = p^{*}u(x) + (1-p)^{*}u(y) \tag{1.11}$$

If $p = 100\%$, then we get x with certainty, as above. If $p = 0\%$, then $(1-p) = 100\%$ and so we get y with certainty, as above. But for any other value that p can take, $0 < p_i < 1$, the expected utility is some amount in-between the two. This is represented graphically as the chord between (x,u(x)) and (y,u(y)) in Figure 1.18.

We can also determine the expected amount of consumption, given these probabilities. This is expressed as Equation 1.12 below and is illustrated in Figure 1.18 as being between x and y on the x-axis.

$$E[c] = p^{*}x + (1-p)^{*}y \tag{1.12}$$

Upon inspection of Figure 1.18, we notice that the expected utility at this level of consumption is below that of the utility curve, the utility we would get for sure if we were given the expected value of consumption instead of having to face the uncertainty. Because of this, we would say this individual is **risk averse**.

Definition 1.6. Risk Aversion

An individual is risk averse if they would prefer to receive a given amount for sure, as opposed to, on average, receiving the same amount through a gamble/lottery/uncertainty. As these functions are ***concave****, we will primarily represent these types of functions as either* $u(c) = \sqrt{c}$ *or* $u(c) = ln(c)$*.* ♣

Some additional questions that we may want answered are: what level of consumption, given for certain, would an agent be willing to get in order to avoid the gamble? Related, how much would an individual pay to avoid risk? It is easiest to answer these questions through the example below.

Example1.6 Conan is a warrior who enjoys seeing his enemies slain before him. Currently, he terminates 2 foes per day. Suppose that he buys a sword of unknown quality. If it is a good sword, he can increase the number of enemies that fall by his hand by 1 per day. If it is a bad sword, the number of families that lament the loss of their breadwinner reduces by 1. His preferences of consumption is represented by $v(c) = \ln(c)$.

- First, write the function for the number of opponents slain, assuming that the probability of getting either type of sword is 50%:

$$E(c) = \frac{1}{2}(1) + \frac{1}{2}(3) = 2$$

- What is the expected utility given the above parameters?

$$E(u(c)) = \frac{1}{2}(ln(1)) + \frac{1}{2}(ln(3)) = .5493$$

- What is the utility if the expected number of opponents slain occurred with certainty?

$$u(2) = ln(2) = .69$$

- How many certain kills would a sword have to be able to provide for Conan to be willing to accept in lieu of facing the uncertainty provided by buying the other sword? This number of kills would be known as the **Certainty Equivalent**.

Definition 1.7. Certainty Equivalent

The Certainty Equivalent of a gamble (uncertain situation) is the amount of consumption an individual is indifferent between receiving for certain and facing the gamble. ♣

Since individuals make decisions based on what would maximize their utility, we find the Certainty Equivalent (CE) by looking for a level of utility given for certain that would match the expected utility at the expected level of consumption:

$$ln(CE) = E(u(E(g)))$$

Compare the expected utility on the gamble line to that given by some amount for certain, CE, on u(c)

$$ln(CE) = \frac{1}{2}(ln(1)) + \frac{1}{2}(ln(3)) = .5493$$

$$e^{ln(CE)} = e^{.5493}$$

$$CE = 1.73 \tag{1.13}$$

So, Conan would get the same amount of utility, .5493, if he for sure got 1.73 kills as opposed to, on average, getting 2 kills with uncertainty. Another way of saying this is that he is willing to "pay" $2 - 1.73 = 0.27$ to avoid the gamble. This number is called the **Risk Premium**.

Definition 1.8. Risk Premium

The risk premium is the amount of consumption/wealth that an individual is willing to part with in order to avoid the gamble. ♣

Other preferences over risk may be apparent depending on the person and situation. Figure 1.19 represents an individual that is Risk Neutral, or they are indifferent between a fair gamble and receiving the expected level of wealth for sure.

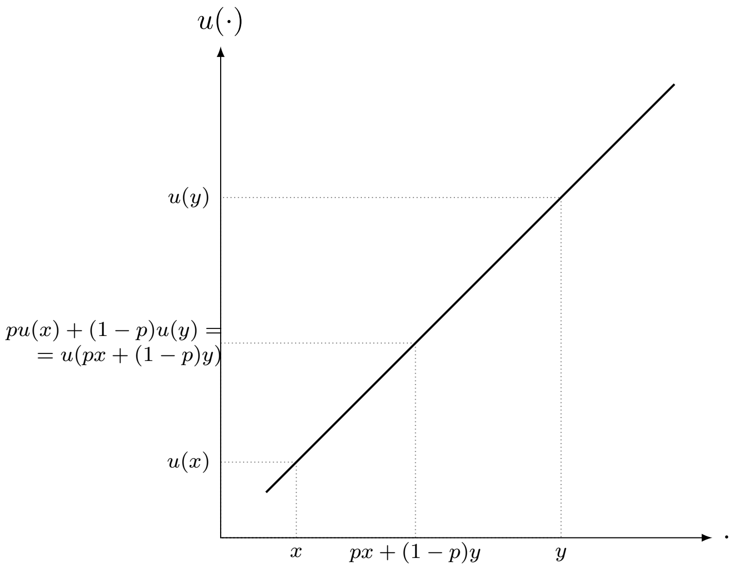

Figure 1.19: Risk Neutrality

✍ **Exercise1.5** How would you imagine a utility curve looking if an individual loved the risk of a gamble? How might the solutions in Example 1.6 change if Conan loved risk?

Chapter Monopoly

2.1 Putting It All Together

Consider a monopoly where we are given two bits of information:

$$\text{Demand: } Q = 80 - \frac{1}{5}p \tag{A}$$

and

$$\text{Total Cost: } TC = 100 + 20Q^2 \tag{B}$$

From just these two equations, we can learn much. To start, let us derive some more equations from these two:

$$\begin{aligned}
\text{Inverse Demand: } P &= 400 - 5Q && \text{Solve A for P} && \text{(C)}\\
\text{Marginal Cost: } MC &= 40Q && \text{Derivative of B} && \text{(D)}\\
\text{Average Total Cost: } ATC &= \frac{100}{Q} + 20Q && \frac{B}{Q} && \text{(E)}\\
\text{Revenue: } Rev &= P * Q && && \\
&= (400 - 5Q)Q && \text{Sub in P from C} && \\
&= 400Q - 5Q^2 && && \text{(F)}\\
\text{Marginal Revenue: } MR &= 400 - 10Q && \frac{dF}{dQ} && \text{(G)}\\
\text{Profit: } \Pi &= 400Q - 25Q^2 - 100 && \text{F} - \text{B} && \text{(H)}
\end{aligned}$$

In Figure 2.1, these figures are graphed and the process of finding the optimal profit using the graphs as a reference is undertaken. Following this are Technical Explanations that use a more mathematical approach to find the same solutions.

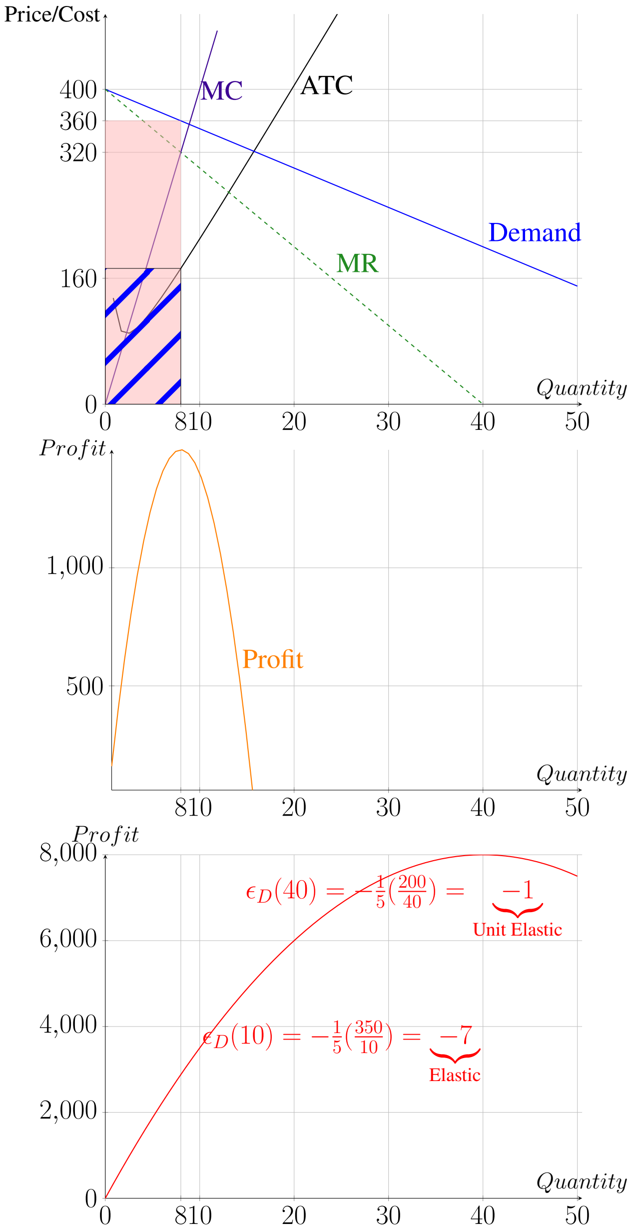

Figure Notes

To calculate the optimal production for a monopoly, set $MC = MR$:

$$40Q = 400 - 10Q$$
$$50Q = 400$$
$$Q_m = 8 \tag{2.1}$$

The price would therefore be:

$$P(8) = 400 - 5(8)$$
$$= 360 \tag{2.2}$$

The revenue ($P \times Q$) at this price, quantity is then:

$$Revenue = 360 * 8$$
$$= 2880 \tag{2.3}$$

This amount is represented geometrically by the pink rectangle.[a] To find the total cost, we can use the TC formula:

$$TC = 100 + 20Q^2$$
$$TC = 1380 \tag{2.4}$$

Note that this amount is represented geometrically by the striped blue lines. [b] Now we can find profit (Π):

$$\Pi = Revenue - Cost$$
$$= 2880 - 1380$$
$$= 1500 \tag{2.5}$$

We can see in the middle graph that profit seems to maximize at this number when graphed.

[a] What are the dimensions and area of the pink rectangle?

[b] What are the dimensions and area of THIS rectangle?

Figure 2.1: Graphs of many functions for a monopolist and explanations of solving associated problems

Technical Explanation 2.1. Maximize Profit w/ Calculus

The monopolists maximization problem is as follows:

$$\Pi = \max_{Q}\{400Q - 25Q^2 - 100\} \tag{2.6}$$

Since the slope of the profit function will equal zero at a maximum, we can take the first derivative and set it equal to zero in order to find the point where profit is maximized:[a]

$$\begin{aligned}
\frac{d\Pi}{dQ} &= 0 \\
400 - 50Q &= 0 \\
\underbrace{400 - 10Q}_{MR} &= \underbrace{40Q}_{MC} && \textit{Add 40Q to both sides} \quad (2.7) \\
400 &= 50Q && \textit{Add 10Q to both sides} \\
Q_M &= 8 && \textit{Divide both sides by 50} \quad (2.8)
\end{aligned}$$

Notice that in Equation 2.7, mathematical reasoning is given for why MR must equal MC for the monopolist. Also, we can plug 8, the optimal production quantity (Q), into the profit function to find the maximum profit:

$$\begin{aligned}
\Pi(Q) &= 400Q - 25Q^2 - 100 \\
\Pi(8) &= 400(8) - 25(8)^2 - 100 \\
\Pi(8) &= 1500 && (2.9)
\end{aligned}$$

[a] Note that the slope could equal zero at the minimum of the function as well. You should consider taking the second derivative to be sure, but this is not necessary here.

♡

Technical Explanation 2.2. Maximize Revenue w/ Calculus

It should be clear from Figure 2.1 why maximizing revenue should not be the goal; however, drawing the figure manually would be made easier by figuring out where the maximization occurs. We will follow the same procedure as maximizing profit, with the maximization problem as follows:

$$\Pi = \max_{Q}\{400Q - 5Q^2\}$$

The derivation of the maximum is as follows:

$$\frac{dRevenue}{dQ} = 0$$

$$\underbrace{400 - 10Q}_{MR} = 0$$

$$400 = 10Q$$

$$Q_R = 40 \tag{2.10}$$

with the maximum revenue being:

$$Revenue(Q) = 400(Q) - 5Q^2$$

$$Revenue(40) = 400(40) - 5(40)^2$$

$$= 8000 \tag{2.11}$$

Chapter Game Theory

Learning Objectives

- *Duopoly Models*
- *Nash Equilibria*
- *Sustained Cooperation*
- *Mixed Equilibria*
- *Sequential Games*
- *Evolutionary Game Theory*
- *Other Toy Games*

In order to model various types of strategic interactions between two or more agents, Economists use **Game Theory**. Solving problems using Game Theory involves finding the *best response* to other agents' decisions, where other agents are referred to as "players" or even "opponents." In order to utilize Game Theory, there must be several conditions:

- There must be more than one player
- There must be more than one strategy for each player
- Players must be rational

This chapter describes some of the more common ways that Game Theory is utilized. In general, Game Theory follows the preceding themes, but can take many forms. In other words, while the goal is to determine the best choices given opponents' choices, many things can vary, creating a different environment. For example, agents can choose simultaneously or sequentially; the number of players can vary from 2 players to many players; agents may care about relative payoffs more than level payoffs. These differences occur because there are many ways in which strategic interactions can occur: competition between firms, war strategies; and even how to behave on a date.

3.1 Duopoly Models

So far we have seen markets where there is a large number of suppliers (perfect competition) and markets where there is but one supplier (monopoly) and we have discussed the differences between the two. But what lies in-between these extremes? The natural starting place is to consider the case of two firms, or a **Duopoly**.

Definition 3.1. Duopoly

A Duopoly is a market where there exists exactly two firms for a given good. ♣

Duopoly firms can engage in strategic competition by competing on price or quantity. In this book, we will focus on competing on quantity. Additionally, the firms could choose the quantity they produce simultaneously, or one firm could choose before the other. Whether the assumptions of the model fit a particular real life situation well is a case-by-case situation. Let us consider each model.

3.1.1 Stackelberg Competition

The first model of duopoly we consider is **Stackelberg competition**. In this model, the competitors compete on the quantity they will each produce and do so one after the other. This model is named after German economist Heinrich Freiherr von Stackelberg who formulated the model. In describing this market, we would like to be informed about the market price of the good in question, the amount the firms each produce, their profit, and the welfare of consumers.

To start, we will consider a case where the producers face a linear demand curve: $p = a - b(y)$ where y is the total market production of the good ($y = y_1 + y_2$), where y_1 and y_2 represent the production of Firm 1 and Firm 2, respectively. We will establish the convention that Firm 1 will move first and Firm 2 will move second. We next simplify the profit function by *assuming* that there are *no costs* to produce their good. This means that the profit function is equal to the revenue function. Note that this is not necessary to perform this analysis, but it does make the initial process much easier.

To formulate our strategy in solving this problem, we consider this: Firm 1, the "leader", knows the demand curve that the two firms face and if they further also know the costs for Firm 2 (in this case, both firms face no costs), then Firm 1 could figure out Firm 2's optimal choice given Firm 1's choice. This information can inform Firm 1's optimal choice, giving it an advantage to moving first. This process of starting at the final choice and working backward is called **Backward Induction**.

Definition 3.2. Backward Induction

Backward Induction is the strategic process of considering optimal choices that will be made at the end of a sequence of events and using those choices to inform earlier choices. ♣

Given this, we first consider Firm 2's (the follower) position first. Given that our simple example has no costs, Firm 2's profit function is simply revenue:

$\Pi_2(y_1, y_2) = p(y_1, y_2)y_2$

$\Pi_2(y_1, y_2) = (a - b(y_1 + y_2))y_2$ Substitute inverse demand curve given above

Thus, Firm 2's maximization problem is:

$$\max_{\{y_2\}}[a - b(y_1 + y_2)]y_2 \tag{3.1}$$

In order to find the maximum, we set the derivative of the function equal to zero. In this case, since we have no costs, the derivative is simply marginal revenue. Setting this equal to zero, we have:

$$a - by_1 - 2by_2 = 0$$
$$2by_2 = a - by_1$$
$$y_2^* = \frac{a - by_1}{2b} \tag{3.2}$$

Thus, we have the optimal amount that Firm 2 should make as a function of how much Firm 1 makes. We could call this Firm 2's **Best Response** to Firm 1. Since the leader, Firm 1, knows this, they can make an informed decision on how much to produce by plugging this information into their maximization problem:

$$\Pi_1 = \max_{\{y_1\}}[a - b(y_1 + y_2)]y_1 \qquad \text{s.t. } y_2 = \frac{a - by_1}{2b}$$
$$\Pi_1 = \max_{\{y_1\}}[a - b(y_1 + \frac{a - by_1}{2b})]y_1 \qquad \text{Plug in constraint}$$
$$\Pi_1 = \frac{a}{2}y_1 - \frac{b}{2}y_1^2 \qquad \text{Simplifying} \tag{3.3}$$

Now, we execute the maximization problem by setting the derivative, with respect to y_1, equal to zero:[1]

$$\frac{a}{2} - by_1 = 0$$
$$y_1^* = \frac{a}{2b} \tag{3.4}$$

Notice that this is not a function of y_2, this is the actual amount of optimal production for Firm 1. This is because we eliminated y_2 from the equation when we substituted Firm 2's best response function in Equation 3.3. We can use this to our advantage in finding the numerical value for y_2 by plugging Equation 3.4 into Equation 3.2:

$$y_2^*(y_1^*) = \frac{a - b(y_1^*)}{2b}$$
$$y_2^* = \frac{a - b(\frac{a}{2b})}{2b} \qquad \text{Substitute } y_1^* = \frac{a}{2b}$$
$$y_2^* = \frac{a}{4b} \qquad \text{Simplify} \tag{3.5}$$

3.1.2 Cournot

Now we consider the case where each firm chooses the optimal level of production simultaneously. Assuming the same conditions as before, i.e. $p = a - b(y_1 + y_2)$ and no costs for production, each firm has the same maximization problem.

[1] This is equivalent to setting Marginal Revenue equal to Marginal Cost. Marginal cost, in our simple example, is zero.

Don't Forget!

Remember that in these simple examples, there is no cost. It is quite possible that you may get questions where there is a cost and this cost could differ for each firm. This does not change the process. This just adds a new term, which means more algebra.

The process for each firm would be to substitute the inverse demand function into their profit function and maximize. We actually already did this in the prior problem, for Firm 2 in Equation 3.2. Since we are assuming the same profit structure for Firm 1 in this problem (both have no cost), we should see a symmetric best response for Firm 1:

$$y_1^*(y_2^*) = \frac{a - by_2}{2b} \tag{3.6}$$

We can eliminate one variable by substituting one equation into the other. Let us solve for y_1^* by subbing in $y_2(y_1)$ from Equation 3.2 into 3.6:

$$y_1 = \frac{a - b[\frac{a-by_1}{2b}]}{2b}$$

$$y_1^* = \frac{a}{3b} \qquad \text{Simplifying} \tag{3.7}$$

And since both firms are identical, we have a symmetric solution for y_2: $y_2^* = \frac{a}{3b}$. Let's compare these results to that of Stackelberg in Table 3.1. For consumers, Cournot is worse than Stackelberg because Stackelberg generates more production and a lower price. But what about welfare for producers? This is presented in Table 3.2. For Stackelberg, the Leader (Firm 1) ends up with the best possible profit from all possibilities, while the Follower (Firm 2) ends up with the worst. This implies a distinct *First Mover Advantage*.

		Production	Price
Stackelberg	Firm 1	$y_1^* = \frac{a}{2b}$	
	Firm 2	$y_2^* = \frac{a}{4b}$	
	Total	$y = \frac{3}{4}\frac{a}{b}$	$P = a - b(\frac{3}{4}\frac{a}{b}) = \frac{1}{4}a$
Cournot	Firm 1	$y_1^* = \frac{a}{3b}$	
	Firm 2	$y_2^* = \frac{a}{3b}$	
	Total	$y^* = \frac{2a}{3b}$	$P = a - b(\frac{2}{3}\frac{a}{b}) = \frac{1}{3}a$

Table 3.1: Comparison of Price and Production Between Stackelberg and Cournot Competitors.

		Profit
Stackelberg	Firm 1	$\Pi_1 = \frac{a}{2b}\frac{1}{4}a = \frac{1}{8}\frac{a^2}{b}$
	Firm 2	$\Pi_2 = \frac{a}{4b}\frac{1}{4}a = \frac{1}{16}\frac{a^2}{b}$
	Total	$\Pi = \frac{3}{16}\frac{a^2}{b}$
Cournot	Firm 1	$\Pi_1 = \frac{a}{3b}\frac{1}{3}a = \frac{1}{9}\frac{a^2}{b}$
	Firm 2	$\Pi_2 = \frac{a}{3b}\frac{1}{3}a = \frac{1}{9}\frac{a^2}{b}$
	Total	$\Pi = \frac{2}{9}\frac{a^2}{b}$

Table 3.2: Comparison of Profit Between Stackelberg and Cournot Competitors.

3.2 Normal Form Games

We will now discuss a class of games where players move simultaneously. These simultaneous form, or Normal Form, games are expressed through cells that represent payoffs. Figure 3.1 is an example of such a game.

		Player II	
		Cooperate	Defect
Player I	Cooperate	3,3	1,4
	Defect	4,1	2,2 *

Figure 3.1: An Example Normal Form Game

By convention, the rows represent strategies available to the first player while the columns represent strategies available to the second player. The individual cells in the table represent the combinations of the two players' moves. Within these cells are numbers. These numbers represent the payoffs for each player at a particular combination of strategies. The first number is the payoff to Player I (the row player) and the second number is the payoff to Player II (column).Note that we need not only include two players, but two is all that is necessary to get the general idea of how normal form games work.

3.2.1 Nash Equilibrium

The goal of working with these games will be to "solve" them. For this, we must have a solution concept. The first solution concept will be a **Nash Equilibrium**. The Nash Equilibria is so named for John Nash, winner of the 1994 Nobel Memorial Prize in Economic Sciences for his work in Game Theory.

Definition 3.3. Nash Equilibrium

*A Nash Equilibrium is a combination of strategies where no player has an incentive to **unilaterally** deviate from.* ♣

There are actually multiple ways we could find Nash Equilibria. The most straightforward way are to find mutual best responses for each of the players:

1. **Put Yourself in a Player's Shoes**
    - For this step, pretend you are one of the players. It does not matter which one as we will do this process for each player.
2. **Hold An Opponent's Moves Constant**
    - Assume that your opponent picks a particular strategy. This means their other choices are not relevant at this time. You may choose to cover up other choices the opponent could have made with your hand or a notecard to demonstrate that fact.
3. **Determine your Best Response (BR) Given Your Opponent's Action**
    - If your opponent has chosen to move as in Step 2., which strategy yields the best payoff for you? Circle (or underline) the payoff for such a strategy in the payoff matrix to denote the best response. Note: It is possible that the best payoffs could be achieved from multiple strategies. In this case, each of the strategies is a best response.
4. **Repeat Steps 2 & 3 For All Opponent's Strategies**
5. **Repeat Steps 1 through 4 For All Remaining Players**
6. **Mutual Best Responses Are Nash Equilibria**
    - Any entry with all numbers circled (or underlined if you prefer) is NE. Typically, a NE is denoted with an asterisk (*) in the given cell.

As an example, consider Figure 3.2. First let us put ourselves in the shoes of Player I. As Player I, we have two possible moves: Cooperate or Defect. If Player II chooses Cooperate, then I could also choose Cooperate and get a payoff of three... or we could choose Defect and get a payoff of four. Four is the better option, so we circle the four, effectively notating that Defect is Player I's best response to Player II playing Cooperate. If Player II chooses Defect, we could choose Cooperate and get a payoff of one... or we could choose Defect and get a payoff of two. Two is better than one, so we circle two, indicating the best response for Player II playing Defect is for Player I to also play Defect. So, in this case, no matter what the opponent does, the player chooses a particular strategy. This is known as a **dominant strategy**. Note that this particular game is symmetric, meaning that both players have the same strategies and same payoffs given the other player's move. We don't have to repeat the process for Player II in this case, because we know that defecting is a dominant strategy. Drawing the circles to reflect that, we see that there is one cell that is a mutual best response (has all numbers circled): {Defect, Defect}, our Nash Equilibria.

Prisoners' Dilemma

Here we will discuss a particular type of normal form game called a Prisoners' Dilemma. The narrative behind this game is that the FBI have arrested two gangsters and have separated them, disallowing contact between them. The prosecutors tell each of them, individually, that if they cooperate with their investigation and confess, they will get a lesser punishment (a higher payoff), especially if their partner does not confess. On the other hand, they know the reverse is true. If the other player confesses and you do not, you will go to jail for quite some time (you will have a worse payoff). The ranking of these outcomes have been reflected in the payoffs presented in Figure 3.2, which we solved in the previous section. Note that the best outcome, socially, is for the two of them to cooperate with each other and not confess; however, there is incentive for each player to deviate from this strategy, and defect (confess), thus the dilemma.

		Player II	
		Cooperate	Defect
Player I	Cooperate	3,3	1,④
	Defect	④,1	②,②*

Figure 3.2: Collective action problem in the Prisoners' Dilemma

This type of structure can be applied to many situations, making it the most popular example of a normal form game. For example, two firms may be considering an advertisement war with each other. . . the payoff structure in that case, may be similar to that described above. Figure 3.3 illustrates a generic Prisoners' Dilemma game if and only if $C > A > D > B$.

		Player II	
		Cooperate	Defect
Player I	Cooperate	A,A	B,C
	Defect	C,B	D,D

Figure 3.3: This game represents a Prisoners' Dilemma if and only if $C > A > D > B$

3.2.2 Mixed Strategies

Strategies where a player must choose one of their options definitively (pure strategies) are not the only strategy profile available. A player could also choose to randomize between their strategies. This is called a **mixed strategy**.[2]

Definition 3.4. Mixed Strategy

A Mixed Strategy is a strategy where a player plays a set of strategies with probability weights assigned to them. ♣

In Figure 3.4, we have a normal form game where Player I plays Up with probability p and plays down with probability (1-p). Player II plays Left with Probability q and Right with probability (1-q). In order to find all equilibria, including mixed, we need to follow some

[2]Another interpretation could be that there are a large number of players and the population plays different moves in certain proportions.

		q	(1-q)
		Left	Right
p	Up	3,2	1,1
(1-p)	Down	0,0	2,3

Figure 3.4: A two-player game.

procedures:

1. Find the expected payoffs for each players' strategies if their opponent mixes (assume that each player is risk neutral). This is simple to do. We simply take the average of the payoffs for a pure strategy, weighted by the probability that the opponent puts on their strategies. Let us do this in detail for Player I:
   - If Player I chooses Up and Player II mixes, Player I's expected payoff is:

$$E[\Pi]_{Up} = \underbrace{3q}_{\substack{\text{She gets a payoff of 3 in \{Up, Left\}} \\ \text{and Player II plays Left with probability q}}} + \underbrace{1-q}_{\substack{\text{She gets a payoff of 1 in cell \{Up, Right\}} \\ \text{with probability (1-q).}}} = 1 + 2q \qquad (3.8)$$

   - If Player I chooses Down and Player II mixes, Player I's expected payoff is:

$$E[\Pi]_{Down} = \underbrace{0q}_{\substack{\text{She gets a payoff of 0 in \{Down, Left\}} \\ \text{and Player II plays Left with probability q}}} + \underbrace{2-2q}_{\substack{\text{She gets a payoff of 2 in cell \{Down, Right\}} \\ \text{with probability (1-q).}}} = 2 - 2q \qquad (3.9)$$

   Now that we have found the expected payoffs for Player I given mixing by Player II, let us update the original game to reflect the new strategy of "Mix" for Player II and the resulting expected payoff for Player I:

		q	(1-q)	
		Left	Right	Mix
p	Up	3,2	1,1	1+2q
(1-p)	Down	0,0	2,3	2-2q

Figure 3.5: A two-player game, updated with Player I's expected payoffs, given mixing by Player II.

   If we do the same procedure for Player II given Player I mixing, we get the following:

2. We can now use this information to define the best responses for each player given the other payers probability mix. Let us do this in detail for Player I.
   - If Player I's expected payoff for playing Up is greater than her expected payoff for playing Down, she will exploit this by playing Up for certain (she will set $p = 1$).

		q Left	(1-q) Right	Mix
p	Up	3,2	1,1	1+2q
(1-p)	Down	0,0	2,3	2-2q
	Mix	2p	3-2p	

Figure 3.6: A two-player game, updated with Player I's and Player II's expected payoffs, given mixing by the other player.

This happens where:

$$1 + 2q > 2 - 2q$$
$$4q > 1$$
$$q > \frac{1}{4} \tag{3.10}$$

- If Player I's expected payoff for playing Down is greater than her expected payoff for Playing Up, she will exploit this by playing Down for certain (she will set $p = 0$, so that $(1 - p) = 1$)). This happens where $q < \frac{1}{4}$.
- If Player I's expected payoff for her strategies are the same, she will be indifferent between Up and Down (any value of p works: $p \in [0, 1]$). This happens where $q = \frac{1}{4}$.

We can organize/summarize Player I's best response function (and Player II's function from doing the same procedure with her) based on this information as follows:

$$BR_1 = \begin{cases} \text{Up} & (p = 1) & \text{if} \quad q > \frac{1}{4} \\ \text{Down} & (p = 0) & \text{if} \quad q < \frac{1}{4} \\ \text{Up or Down} & (p \in [0, 1]) & \text{if} \quad q = \frac{1}{4} \end{cases} \tag{3.11}$$

$$BR_2 = \begin{cases} \text{Left} & (q = 1) & \text{if} \quad p > \frac{3}{4} \\ \text{Right} & (q = 0) & \text{if} \quad p < \frac{3}{4} \\ \text{Left or Right} & (q \in [0, 1]) & \text{if} \quad p = \frac{3}{4} \end{cases} \tag{3.12}$$

3. Now, we can illustrate the best responses and show equilibria graphically. To do this, we let the best response functions serve as a guide. I will discuss the process of graphing Player I's best response function (Equation 3.11) in detail. First set up an axis for q and for p, from zero to one. Next, consider each line in Equation 3.11 and interpret it graphically as below:
   (a). Draw a line everywhere at p=1 where q is greater than 1/4
   (b). Draw a line everywhere at p=0 where q is less than 1/4
   (c). Here, Player I is indifferent on Up or Down, so all values of p should be drawn in

where q=1/4

The leftmost graph in Figure 3.7 shows this illustration. The middle graph is for Player II and the rightmost graph combines the two. Notice that the two curves intersect at three places, two at corners and one in the middle. The intersections on the corners represent Pure Strategy Nash Equilibria. If we solve the normal form game as usual, we would find that there would be Nash Equilibria at {Down, Right} and {Up, Left}. In terms of probability, {Down, Right} is where p and q both equal zero and {Up, Left} is where p and q both equal 1. This matches up with the corner locations on our graph. Finally, we have a third equilibria, this one where Player I plays Up 3/4 of the time and Down 1/4 of the time, and where Player II plays Left 1/4 of the time and Right 3/4 of the time.

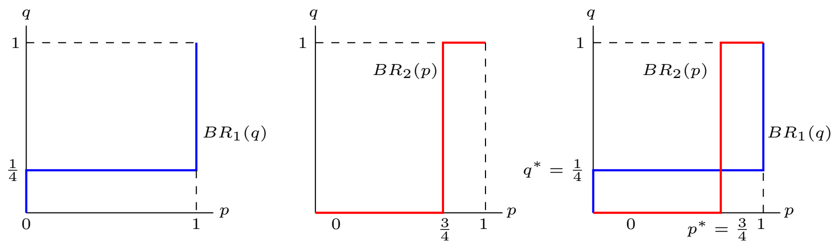

Figure 3.7: Best Response Curves

3.2.3 Sustained Cooperation

In some types of games, the Prisoners' Dilemma for example, we may wish to sustain cooperation, but we have the incentive not to. In fact, defecting strictly dominates cooperation. However, in society, we often observe some amount of cooperation. Why?

One difference between our normal form games and real life is that in real life, we often do not play a game just once. For example, a pair of firms trying to decide if they want to engage in an advertising campaign or not do not decide their choice one day, and then never consider it again. We may have it that these games are repeated. Could having to face our opponents repeatedly give us incentive to cooperate? Let's find out.

First, let us consider the case of repeating a Prisoners' Dilemma some finite number of times, T. In time T (the final period), do the firms have any incentive to cooperate? No. In fact, since the game will terminate after this round, the game is, at this point, exactly the same as our standard normal form game of a Prisoners' Dilemma. If the firms understand this, then in period T-1, they also have no incentive to cooperate, because they know there will be no cooperation in the following period. In period T-2, they will see that there will be no cooperation in period T-1 and will not have incentive to cooperate there, either. This analysis pattern will continue until we get back to the present; thus, repeating the game finitely does not help to generate cooperation.

However, if there IS no final period, we would not have this problem! To see if we can sustain cooperation for an infinitely repeated game, let us introduce some notation. First, consider that since we will be dealing with future income, we will need to discount payoff flows to the present in order to compare the payoffs from cooperating and defecting. To do this, we use a discount rate, $\delta \in [0, 1)$. δ itself can have many justifications for its value. For example, if we are dealing with money, we can consider it to be related to the interest rate, r: $\delta = \frac{1}{1+r}$. It could represent the chance that the game continues to the next round or even a term that measures an individuals' patience. As an example of how this might look, consider a three-round game with payoffs in each period, u_t, represented below:

$$\Pi = \sum_{t=1}^{3} \delta^{t-1} u_t = u_1 + \delta u_2 + \delta^2 u_3, \tag{3.13}$$

where an infinitely repeated game would then be represented as:

$$\Pi = \sum_{t=1}^{\infty} \delta^{t-1} u_t, \tag{3.14}$$

which is equal to:

$$\Pi = u + \delta u + \delta^2 u + \cdots + \delta^{n-1} u \tag{3.15}$$

This type of series is called a Geometric Series. To simplify this, we will use some tricks. First, we will exploit the fact that if you multiply each side of an equality by the same number, the equality is maintained. This in mind, we multiply Equation 3.15 by δ. This gives us:

$$\delta\Pi = \delta u + \delta^2 u + \cdots + \delta^{n-1} u + \delta^n u \tag{3.16}$$

The next trick we use will be to exploit the fact that you can add/subtract the same amount from each side of an equation and maintain the equality. We therefore subtract Equation 3.16 from Equation 3.15:

$$\Pi = u + \delta u + \delta^2 u + \cdots + \delta^{n-1} u$$

$$\delta\Pi = \delta u + \delta^2 u + \cdots + \delta^{n-1} u + \delta^n u$$

Notice the lines connecting identical values from the two equations. The only parts of the expression without a match are the left hand side, u, and $\delta^n u$. Thus, the result of subtracting the two is:

$$\begin{aligned} \Pi - \delta\Pi &= u - \delta^n u \\ \Pi(1-\delta) &= u(1-\delta^n) \\ \Pi &= u\frac{1-\delta^n}{1-\delta} \end{aligned} \tag{3.17}$$

We can further simplify this by considering that since we play this game infinity times, our equation can change as n gets arbitrarily large:

$$\lim_{x\to\infty} \Pi = \lim_{x\to\infty} u\frac{1-\delta^n}{1-\delta}$$
$$= \frac{u}{1-\delta} \tag{3.18}$$

Thus, we can also say that for a geometric series going to infinity where $u = 1$:

$$\Pi = 1 + \delta + \delta^2 + \delta^3 + \ldots = \frac{1}{1-\delta} \tag{3.19}$$

3.2.4 Grim Trigger

We will review two particular strategies for attempting to cooperate in an infinitely repeated Prisoners' Dilemma. The first will be a **trigger** strategy. In particular, we will look at a particularly harsh strategy, the **Grim Trigger** strategy.

Definition 3.5. Trigger Strategy

A trigger strategy is a strategy whereby players threaten other players with a punishment if they deviate from an agreed upon strategy. ♣

Definition 3.6. Trigger Strategy

The Grim Trigger strategy is a strategy whereby players threaten to punish players forever after a single deviation from cooperation. ♣

		Player II	
		Cooperate	Defect
Player I	Cooperate	3,3	1,④
	Defect	④,1	②,②*

Number left (right) of comma refers to A's (B's) preference ordering
(1 = worst outcome; 4 = best outcome). * indicates the equilibrium.

Figure 3.8: Collective action problem in the Prisoners' Dilemma

Consider the game in Figure 3.8. If a given player (this is a symmetric game, so in this case, it does not matter which player we are) were to cooperate forever, they would receive a payoff of three for every round. If we consider discounting, the present value of the payoffs would be:

$$\Pi_C = 3 + 3\delta + 3\delta^2 \ldots \tag{3.20}$$

And if one chooses to betray the other player by defecting, they would get a higher payoff during the period of their betrayal, but would get a reduced payoff for all time, because their opponent is playing Grim Trigger:

$$\Pi_D = 4 + 2\delta + 2\delta^2 + \ldots \tag{3.21}$$

So, if the discounted payoff for cooperation (Π_C, Equation 3.20) is greater than or equal to the discounted payoff for defection (Π_D, Equation 3.21), then we have an incentive to cooperate. In our example, we have the following:

$$\begin{aligned} \Pi_C &= 3 + 3\delta + 3\delta^2 + \dots \\ &= \frac{3}{1-\delta} && \text{Using Equation 3.19} \qquad (3.22) \end{aligned}$$

and

$$\begin{aligned} \Pi_D &= 4 + 2\delta + 2\delta^2 + \dots \\ &= 2 + (2 + 2\delta + 2\delta^2 + \dots) && \text{Try to make it look like a geometric series} \\ &= 2 + (\frac{2}{1-\delta}) && \text{Using Equation 3.19} \qquad (3.23) \end{aligned}$$

So we know we can sustain cooperation if δ is such that:

$$\begin{aligned} \Pi_C &\geq \Pi_D \\ \frac{3}{1-\delta} &\geq 2 + (\frac{2}{1-\delta}) \\ 3 &\geq 2 - 2\delta + 2 \\ 2\delta &\geq 1 \\ \delta &\geq \frac{1}{2} \qquad (3.24) \end{aligned}$$

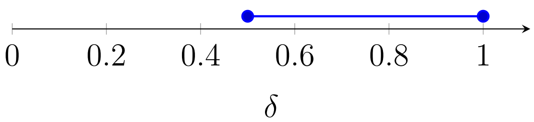

Figure 3.9: δ such that cooperation can be maintained when players play {Grim; Grim}

This range is expressed graphically in Figure 3.9. So, we seem to be able to sustain cooperation under certain conditions. . . but this is a very harsh strategy. Might we be able to maintain cooperation with a nicer strategy? We attempt a second strategy in Section 3.2.5

Example3.1

		Player II	
		Cooperate	Defect
Player I	Cooperate	0,0	-3,2
	Defect	2,-3	-2,-2

Figure 3.10: A Prisoners' Dilemma

Question: You are given a normal form game in Figure 3.10. Suppose two players play this game repeatedly and know that there is a probability $(1-\delta)$ that the game terminates.[3] Further suppose that both players use a "Grim Trigger" strategy. This means that defection is punished by the opponent playing "defect" forever afterward. Find all values of δ such that {Grim Trigger, Grim Trigger} is a Nash equilibrium of this game (there is no profitable unilateral deviation for either player).

To answer this question we need to know discounted payoffs for cooperation and for defection.

For cooperation:

$$\Pi_C = 0 + \delta 0 + \delta^2 0 + \delta^3 0 + \dots$$

For defection:

$$\Pi_D = 2 + (-2)\delta + (-2)\delta^2 + (-2)\delta^3 + \dots$$

In words: if we cooperate, we get payoffs of zero each period. If we defect, we get a nice payoff of two (better than the cooperate payoff for that period) for this period, but we get punished by our grim trigger playing opponent and we get payoffs of -2 forever. The future -2 payoffs are each discounted by δ^n such that in one time period in the future, the -2 is discounted by δ; in two periods it is discounted by δ^2; in three it is discounted by δ^3, etc.

While the payoff to cooperating (in THIS case) was not that complicated, the defection payoffs will require some work to analyze. In particular we notice a pattern. This is very similar to the geometric series we saw as an example in class. How can we make $2+(-2)\delta+(-2)\delta^2+(-2)\delta^3+\dots$ look like the standard geometric series?

In this case, it might be easier to pull out -2δ from the payoffs:

$$\Pi = 2 + -2\delta(1+\delta+\delta^2+\delta^3)$$

Now we should notice that inside the parenthesis is a geometric series. We can replace this with $\frac{1}{1-\delta}$, giving us:

$$\Pi_D = 2 + -2\delta(\frac{1}{1-\delta}) = 2 - \frac{2\delta}{1-\delta}$$

In order for cooperation to be possible, we must have it that the discounted payoffs for

[3] Or a probability, δ, that the game continues.

cooperation is greater than or equal to the discounted payoffs for defection:

$$\begin{aligned}
\Pi_C &\geq \Pi_D \\
0 &\geq 2 - \frac{2\delta}{1-\delta} \\
\frac{2\delta}{1-\delta} &\geq 2 \\
2\delta &\geq 2 - 2\delta \\
\delta &\geq 1 - \delta \\
2\delta &\geq 1 \\
\delta &\geq \frac{1}{2}
\end{aligned}$$

Here, our Nash Equilibrium is to cooperate so long that $\delta \in [\frac{1}{2}, 1]$ Note that if we interpret our problem such that we are discussing interest rates, we can substitute in for δ. If δ is the same as $1/(1+r)$ then δ^2 is the same as $1/(1+r)^2$ and so on. If this is the case, then what interest rate could sustain cooperation? Simply convert the discount rate to interest rate to find out!

$$\begin{aligned}
\frac{1}{1+r} &\geq \frac{1}{2} \\
1 &\geq r
\end{aligned}$$

Wow! So the interest rate must be at least 100% in order to sustain cooperation here.

3.2.5 Tit-for-tat

Since the Grim Trigger strategy may be a bit harsh, is there a less strong punishment we can use? Tit-for-tat is a strategy with which we play the same move our opponent did the turn preceding. Using the game from Example 3.1, if a player defects in round 1, we will have the following pattern emerge for the first nine rounds:

	1	2	3	4	5	6	7	8	9
Player i	C	D	C	D	C	D	C	D	C
Player -i	D	C	D	C	D	C	D	C	D

where C and D stand for "Cooperate" and "Defect" respectively. Notice that we are actually repeating the same two rounds over and over again. We can greatly simplify our problem by only considering the payoffs in these rounds. In this case we have the following:

$$\Pi_C = 0$$

$$\Pi_D = 2 + (-3)\delta$$

In order to sustain cooperation, we must have it that $\Pi_C \geq \Pi_D$:

$$0 \geq 2 + (-3)\delta$$

$$3\delta \geq 2$$

$$\delta \geq \frac{2}{3}$$

So, with Tit-for-Tat, we are able to sustain cooperation with a much higher discount rate (or lower interest rate) than with Grim Trigger. Here, our Nash Equilibrium is to cooperate so long as $\delta \in [\frac{2}{3}, 1]$. Figure 3.11 compares the results between Tit-for-Tat and the Grim Trigger.

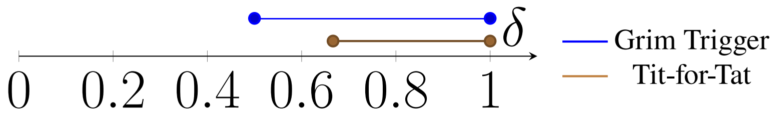

Figure 3.11: Range of δ where Cooperate is a Nash Equilibria

3.3 Sequential Games

3.3.1 Subgame-Perfect Nash Equilibria

Definition 3.7. Subgame-Perfect Nash Equilibria

A Nash equilibrium is said to be subgame perfect if and only if it is a Nash equilibrium in every subgame of the game. ♣

3.4 Introduction to Evolutionary Game Theory

"This preservation of favourable variations and the rejection of injurious variations, I call Natural Selection, or the Survival of the **Fittest**."

— Charles Darwin

Now consider a world with the following conditions:

1. Individuals no longer use strategies to make their choices. Instead, they are hardcoded to use a given strategy.
2. There are many individuals/firms in the world and they will be randomly matched with another individual/firm to play a game
3. To start the game, most individuals play a particular strategy. These individuals are the "incumbents."
4. An arbitrarily small percentage, $\epsilon\%$, of the population play a "mutant" or "invading" strategy. If There are only two strategies, the incumbents make up $(1 - \epsilon)\%$ of the population.
5. Strategies that are relatively more successful are more "fit" and the relative population of the group playing that strategy will grow.

Definition 3.8. Fitness

Fitness refers to the capacity of a variant type to invade and displace the incumbent population in competition for available resources. Fitness level is how we will define our payoffs. ♣

While the concept of Game Theory was borrowed and altered for their purposes by biologists, economists have, in turn, borrowed back these concepts. One of these concepts is a new solution concept, the **Evolutionary Stable Strategy**.

Definition 3.9. Evolutionary Stable Strategy

A strategy is an Evolutionary Stable Strategy (ESS) if its population rejects the invasion of any other mutant strategy. ♣

Biologists define an incumbent's strategy as being ESS mathematically when the following is true:

$$\underbrace{(1-\epsilon)u(\hat{s},\hat{s})+\epsilon u(\hat{s},s')}_{\text{Avg. payoff to incumbent}} > \overbrace{(1-\epsilon)u(s',\hat{s})+\epsilon u(s',s')}^{\text{Avg. payoff to mutant}},\ \forall \epsilon < \bar{\epsilon} \tag{3.25}$$

where:

- ϵ : is an initially small proportion of the invading mutant [gene,company,etc.] in the population.
- $\bar{\epsilon}$ is some small level of ϵ that we choose for all ϵ's to be smaller than to start out with.
- $\hat{s}$ is the incumbent population's hardwired strategy
- s' is the mutant's hardwired strategy

The left-hand side of this inequality represents the average payoff (fitness) of the incumbent. $(1-\epsilon)\%$ of the time, the incumbent will be randomly paired with another incumbent, in which case they will get the payoff associated with both players playing the incumbent strategy, $\hat{s}$. $\epsilon\%$ of the time, the incumbent will be randomly paired with a mutant and the incumbent will receive the payoff associated with playing $\hat{s}$ while the mutant plays the mutant strategy, s'. The resulting average payoff to incumbents, $(1-\epsilon)u(\hat{s},\hat{s})+\epsilon u(\hat{s},s')$, must be greater than the average payoff to the mutant in order to be stable.[4] The mutants' payoffs are calculated in a similar manner: mutants will meet incumbents $(1-\epsilon)\%$ of the time and they will receive a payoff associated with playing the mutant strategy, s', against the incumbent strategy, $\hat{s}$. $\epsilon\%$ of the time, they will be paired against another mutant. The resulting average fitness level is $(1-\epsilon)u(s',\hat{s})+\epsilon u(s',s')$.

There is also an economics definition of an ESS and as it turns out they are equivalent. To see this, first we need to rewrite Equation 3.25 to have both terms on the same side:

[4] If this was not the case, then the mutant would be more fit and would grow in size until it was the new incumbent.

$$\underbrace{(1-\epsilon)[u(\hat{s},\hat{s})-u(s',\hat{s})]}_{\text{Payoff of inc. \& mutant vs inc.}}+\overbrace{\epsilon[u(\hat{s},s')-u(s',s')]}^{\text{Payoff of inc. \& mutant vs mutant}}>0 \tag{3.26}$$

We have defined this equation to hold for ALL ϵ smaller than $\bar{\epsilon}$, so if ϵ is very small and the first term is negative, then the entire left-hand side is negative; however, in Equation 3.26 we show that the left-hand side must be greater than 0 and thus this cannot be the case for an ESS. So:

$$u(\hat{s},\hat{s})-u(s',\hat{s})\geq 0 \tag{3.27}$$

$$u(\hat{s},\hat{s})\geq u(s',\hat{s}) \tag{3.28}$$

ESS Requirement 1

ES strategies must be a symmetric Nash Equilibrium.

But there is more than that to our analysis. Our previous result was a weak inequality, $u(\hat{s},\hat{s})\geq u(s',\hat{s})$. Consider both cases: if $u(\hat{s},\hat{s})$ was greater than $u(s',\hat{s})$ and if they are equal:

$$u(\hat{s},\hat{s})>u(s',\hat{s}) \tag{3.29}$$

and

$$u(\hat{s},\hat{s})-u(s',\hat{s})=0 \tag{3.30}$$

In Equation 3.29, the in-economic-words description is that playing the incumbent strategy is a *strict* Nash Equilibrium. Because of this, the incumbent strategy clearly rewards better fitness for the incumbent and so the incumbent strategy is stable.

ESS Requirement 2a

If the strategy is a *strict* NE, it is ESS.

Now consider if the NE is not strict: $u(\hat{s},\hat{s})-u(s',\hat{s})=0$. If this is the case, the inequality can still be maintained, but now it is all up to the second term to do so. Now, for the incumbent strategy to be ESS, it must be the case that $u(\hat{s},s')>u(s',s')$.

ESS Requirement 2b

If the NE is not strict, in order for the incumbent to be ES, the mutant is allowed to get the same payoffs as incumbents when playing against them, provided they perform relatively poorly when playing against other mutants.

Example3.2 Using the Economics Method These concepts need some grounding in some examples. Consider Figure 3.12. This two-player normal form game has two strategies, appropriately labeled "Incumbent" and "Mutant." Is the incumbent strategy evolutionary stable? To answer this, we review the ESS Requirements. ESS Requirement 1 states that the strategy must be a symmetric Nash Equilibrium, which is the case. ESS Requirement 2a states that if the strategy is a *strict* NE, it is ESS. This is the case, therefore the incumbent has an ESS.

Player II

		Incumbent	Mutant
Player I	Incumbent	1,1*	0,0
	Mutant	0,0	0,0*

Figure 3.12: A Normal Form Game with Multiple NE

Example3.3 Using the Biology Method

Now, consider the Biology Method. Consider the Prisoners' Dilemma in Figure 3.13.

	Cooperate	Defect
Cooperate	2,2	0,3
Defect	3,0	1,1
	$1-\epsilon$	ϵ

Figure 3.13: A Prisoners' Dilemma

Is cooperation an ESS? To answer this, let us calculate the fitness for each type. First, the cooperators. Assuming cooperation is the incumbent strategy, there is a $(1-\epsilon)\%$ chance of meeting another cooperator and therefore getting a payoff of two. There would also be an $\epsilon\%$ chance of meeting a mutant and getting a payoff of zero. This means that the average fitness level of a cooperator is $(1-\epsilon)[2]+\epsilon[0]=2(1-\epsilon)$.

Now for the mutant defectors. As a defector, there is a $(1-\epsilon)\%$ that I am paired with a cooperator, in which case, I get a payoff of 3. There is an $\epsilon\%$ chance that I get paired with another mutant and would therefore get a payoff of 1. This means that the average fitness level of a mutant is $(1-\epsilon)[3]+\epsilon[1]=3(1-\epsilon)+\epsilon$. If we compare the two, we see that $2(1-\epsilon)<3(1-\epsilon)+\epsilon$. In other words, the fitness of the defector is greater on average than the cooperator and so defectors will grow in proportion and eventually take over the population. Therefore, cooperation is not ESS.

We have seen that equilibria exist where a player plays a mixed strategy wherein they play a particular strategy a proportion of the time and another strategy (or strategies) the remaining proportion of the time. When thinking about stable strategies here, we are dealing with hard-coded players that don't make choices, but rather reproduce asexually at rates depending on relative fitness; therefore, we would not have a case where a person can choose between choices. . . however, the evolutionary analogue is that there may exist some stable proportion of the population that play different strategies. In other words, our population is **polymorphic**. We are essentially asking, "what if a mutant (or incumbent) did not die out, but also did not completely take over?"

Definition 3.10. Polymorphic

To be polymorphic is to have more than one form. ♣

Example3.4 Polymorphic Equilibrium

Consider the game of chicken in Figure 3.14 with "Wimp" being the incumbent strategy.

	Wimp	Macho
Wimp	0,0	-1,1
Macho	1,-1	-2,-2
	$1-\epsilon$	ϵ

Figure 3.14: A Game of Chicken

Wimps, the incumbent, will meet other Wimps $(1-\epsilon)\%$ of the time and get a payoff of zero. They will meet Macho types $\epsilon\%$ of the time and get a payoff of -1. This results in an average fitness of a Wimp of $0 * (1-\epsilon) - 1 * \epsilon = -\epsilon$. Macho types will meet Wimps $(1-\epsilon)\%$ of the time and get a payoff of 1. They will meet other Macho types $\epsilon\%$ of the time and get a payoff of -2. This results in an average fitness of a Macho type of $1 * (1-\epsilon) - 2\epsilon = 1 - 3\epsilon$. So, which type is better? It is difficult to compare $-\epsilon$ with $1-3\epsilon$ because which one is better depends on the size of ϵ! The Macho type is fitter if its fitness exceeds that of the Wimp type:

$$1 - 3\epsilon > -\epsilon$$
$$2\epsilon < 1$$
$$\epsilon < \frac{1}{2}$$

So if $\epsilon < \frac{1}{2}$, the Macho type is more fit, logically if $\epsilon > \frac{1}{2}$, the Wimp type is more fit. Intuitively, the Macho type is more fit when there are less Macho guys out there because it is less likely that they will pair with other Macho types and crash their vehicles, but they do well when they have lots of Wimps to pick on. Wimps do well when there are lots of Macho types because Machos are all crashing into each other, but the Wimp never gets hurt. We therefore have a "Polymorphic" ESS of 50% Macho types and 50% Wimp types. This is further illustrated in Figure 3.15 which graphs the expected fitness of each type as ϵ increases. At low levels of ϵ, the Macho type's fitness is higher than the Wimp's. At $\epsilon = \frac{1}{2}$ the fitness is identical and at larger values of ϵ, the fitness of Wimps exceeds that of Macho types.

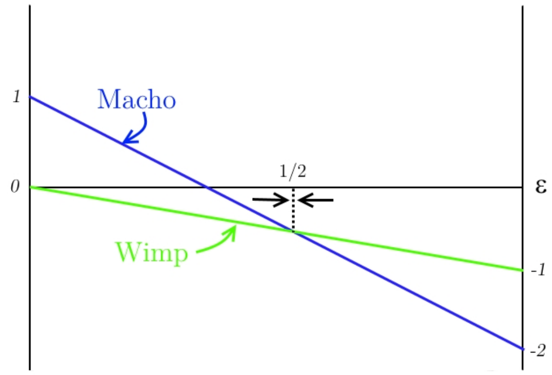

Figure 3.15: Fitness Graphs and Polymorphic Eq.

Appendix Mathematical Tools

This appendix covers some of the basic mathematics used in microeconomics.

A.1 Algebra

A.1.1 Algebra Tips

Here are some key Algebra tips that we will make frequent use of.

- It is often useful to rearrange fractions in the numerator and denominator to be easier to work with. We can rewrite the fraction below as follows:

$$\frac{\frac{A}{B}}{\frac{C}{D}} = \frac{A}{B} \times \frac{D}{C} = \frac{AD}{BC}$$

- In this course, we will often divide by the same variables but with different exponents. It is important to know how to deal with this. The key is to follow the following rule:

$$\frac{X^\alpha}{X^\beta} = X^{\alpha-\beta}$$

A.1.2 The Number e and the Natural Log

The Number e

Like π, e is a number. Specifically, $e \approx 2.72$. Also, like π, it is a special number.

Technical Explanation A.1. Where e Comes From

Let's put ourselves in the banking world and consider interest rates. Specifically, interest rates that compound continuously. The formula for the value of an investment given some interest rate, i, and the number of times the investment is compounded, n, is $P = (1 - \frac{i}{n})^n$. *What if we compounded an arbitrarily large number of times? This would lead to the following:*

$$P = \lim_{n\to\infty} (1 - \frac{i}{n})^n = e^i$$

♡

The function itself is graphed below in Figure A.1. Can you think of anything that follows a pattern like this?

The Natural Log

Log functions are the inverse of exponential functions, like e. They may appear to look something like this: $log_4 16 = ?$. To translate what this is in English, it is asking, "What power can I raise 4 to, such that I would get 16?" The answer to which would be 2, as $4^2 = 16$. The

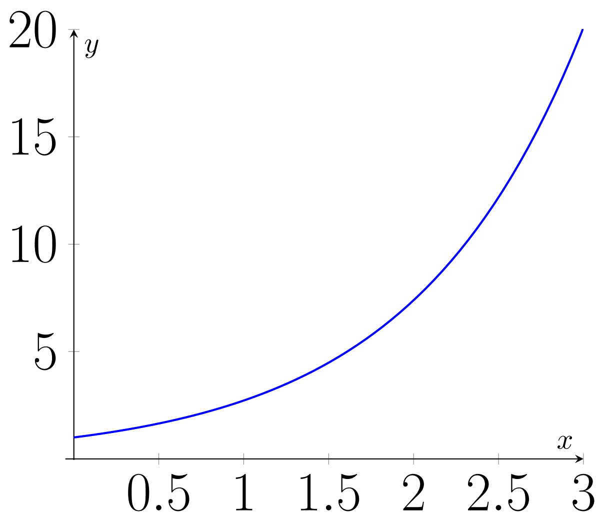

Figure A.1: $y = e^x$

Natural Log (or ln) is a log function has a specific base of e: $log_e(x) = ln(x)$, and thus the natural log of x is the inverse function of e^x. Figure A.2 graphs $ln(x)$.

Logs are a very useful tool in Economics. They have numerous applications:

- As exponential functions are important in Economics, the inverse naturally would be as well. One particular use is in making useful graphs. A popular graph in Economics is showing GDP over time, as in Figure A.3a. But this is a little hard to read. GDP is partially increasing due to inflation and the inputs of GDP. For example, population growth should increase output, right? It may be a better idea to remove this exponential growth from our graph to get a better idea of the actual health of our economy. Figure A.3b removes this exponential growth by taking the natural log of GDP. If the slope of the log adjusted graph is constant, it means the rate of increase of GDP is constant. [1]

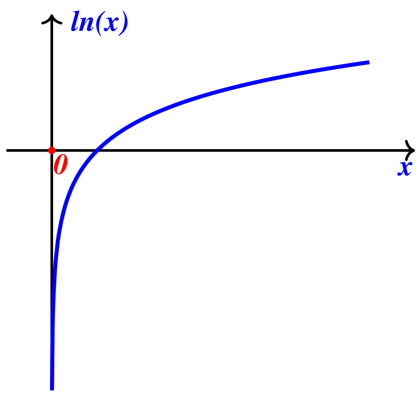

Figure A.2: $ln(x)$

[1]NOTE: You might see graphs that are adjusted in a similar way, but the y-values are not adjusted, simply the distance between the ticks is increasing exponentially. For example, the y-axis of a variable that has been transformed from x to $log_{10}(x)$ might have evenly spaced ticks such as 1, 10, 100, 1000, 10000, etc. This helps for interpretation because the reader does not have to think about how big $log_{10}(x)$ of something is.

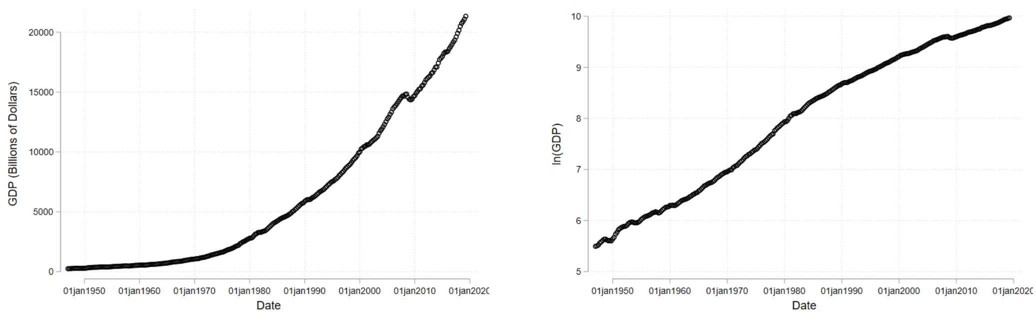

(a) GDP Over Time (b) ln(GDP) Over Time

Figure A.3: GDP Over Time on a Level and Natural Log Scale

- It can be useful in Econometrics. What if the relationship between variables is actually something along the lines of $y = \beta_0 + \beta_1 ln(x)$? Then perhaps the transformation of the x variable would more accurately represent the real relationship between the variables. Additionally, it can allow for the interpretation of coefficients to resemble that of elasticity, e.g. $ln(y) = \beta_0 + \beta_1 ln(x)$ would allow us to interpret the results as: "On average, a 1% change in x results in a $[\beta_1]$% change in y." See your Econometrics text for more.
- They have some useful Algebraic properties. Namely the following:
  - Product Rule: $ln(x \cdot y) = ln(x) + ln(y)$
  - Quotient Rule: $ln(x/y) = ln(x) - ln(y)$
  - Power Rule: $ln(x^y) = y \cdot ln(x)$

  This is especially useful because the natural log is a positive monotonic transformation of x. This means that it preserves the ranking of all the possible numbers it is transforming, i.e. if $x_1 > x_0$, then $ln(x_1) > ln(x_0), \forall x$. Since the units of utility are arbitrary, then if using the properties above could make utility maximization problems more simple, then it is fair game, since its positive monotonic nature preserves the ranking of bundles. Using the problem in Example 1.2 as a reference, if I transformed the utility function, A^2B^3, using natural log, I could use the log properties to write it as $2ln(A) + 3ln(B)$. If we are choosing to do the calculus instead of just memorizing the formula for the MRS, this is an easier derivative. We have:

$$\frac{\partial[2ln(A) + 3ln(B)]}{\partial A} = \frac{2}{A}$$

  and

$$\frac{\partial[2ln(A) + 3ln(B)]}{\partial B} = \frac{3}{B}$$

  Taking the ratio, we have $\frac{\frac{2}{A}}{\frac{3}{B}} = \frac{2B}{3A}$, which is the MRS we found earlier. This is completely optional and is JUST A TOOL. . . But some of you may find this useful.

A.2 Calculus

A.2.1 Basic Derivations

Calculus, in this course, is just a tool. In particular, we will use derivatives and partial derivatives in order to show the rate of change (or slope) of a function. The derivative of a typical function, $f(x) = Ax^3 + 4$, where A is a constant, is expressed below:

$$\frac{df(x)}{dx} = 3Ax^2$$

This came about through the following steps:

1. First, notice the addition sign in the original equation. We can actually separate the derivative process by addition and subtraction operators. In other words, we can take the

derivative of the first part, Ax^3, and the second part, 4, separately.

2. Let's take the derivative of the first part first. A is a constant that is attached to the variable of interest, x. As such, we leave it alone.
3. Next, we look at the variable of interest itself, x, and notice it has an exponent of 3. The process here is to multiply this part of the expression by this number, and then subtract 1 from the exponent. This gives us $3Ax^2$.This is the derivative of the first expression.
4. Next, let's look at the second component, 4. This 4 is a constant and is not attached in any way to a variable, so its derivative is equal to zero. This should make intuitive sense. Ask yourself what the rate of change of a constant is.
5. Therefore, the derivative of the whole function is $3Ax^2$.

A.2.2 Popular Derivatives

- $\frac{dC}{dx} = 0$, where C is a constant
- $\frac{dCx}{dx} = C$
- $\frac{dCx^\alpha}{dx} = \alpha C x^{\alpha-1}$
- $\frac{dln(x)}{dx} = \frac{1}{x}$
- $\frac{de^x}{dx} = e^x$

In order to perform a partial derivative, or a derivative of an expression that contains more than one variable, we perform the same operations as before, just operating as if the other variable is a constant. For example:

$$\frac{\partial 3x^2y^3 + 3x + 2y}{\partial x} = 6xy^3 + 3$$

A.2.3 The Chain Rule

Sometimes our function may be a function of another function, or a "compound function." The procedure here is to take the derivative of the "outside" function and multiply it by the derivative of the "inside" function. For example, suppose we have the following expression that we want a derivative for: $ln(2x^3)$. The first step is to take the derivative of the "outside" function, $ln(\cdot)$. Referring to the prior section, we know that the solution to this would be $\frac{1}{2x^3}$. But we also have the "inner" function to deal with. This function is: $2x^3$, the derivative of which would be $6x^2$. Multiplying the two, as per the chain rule, we get $\frac{6x^2}{2x^3} = \frac{3}{x}$

A.2.4 Constrained Optimization - The Lagrangian

In order to maximize a function under a constraint, we will use a mathematical trick. This trick is to rewrite our objective function with an additional element: add the constraint to the objective function, but have it set to zero. For example, a budget constraint for two goods set to zero would be: $M - p_1x_1 - p_2x_2$. Since, in our examples, $M = p_1x_1 + p_2x_2$, the preceding expression equals zero. Additionally, we add a coefficient, λ to the constraint expression. Then, maximize the entire new function by taking the First Order Conditions (taking the partial derivatives of all variables separately and setting them equal to zero). These conditions describe the optimal values of the inputs. For an example, let us use the same setup as that in Example 1.2.

First, write the Lagrangian function:

$$\mathcal{L}(x_1, x_2, \lambda) = max_{x_1,x_2,\lambda}\{A^2B^3 + \lambda[m - p_A A - p_B B]\}$$

Now, take the First Order Conditions (FOCs):

$$\begin{aligned} 2AB^3 &= \lambda p_A && \text{FOC-1} \\ 3A^2B^2 &= \lambda P_B && \text{FOC-2} \\ M &= p_A A + p_B B && \text{FOC-3} \end{aligned}$$

Since these are equalities, we can divide each side of any of the equations by another one of the equations and it will not change the result. Dividing FOC-1 by FOC-2 yields:

$$\frac{2B}{3A} = \frac{p_A}{p_B}$$

Notice that the left hand side is the MRS (the slope of the indifference curves) and the right hand side is the price ratio (the slope of the budget constraint). Thus, it is optimal when the two curves are equal. This occurs at the two curves' tangency. From here, we can solve the problem as was done in Example 1.2.[2]

[2]Though, in this version, I did not include the given prices.

A.3 Formulas

- Budget Constraint:

$$M = p_x x + p_y y$$

or

$$y = \underbrace{\frac{M}{p_y}}_{\text{y-intercept}} \overbrace{-\frac{p_x}{p_y}}^{Slope} x$$

- Intertemporal Budget Constraint:

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

or

$$c_2 = \underbrace{m_1 \frac{1+r}{1+\pi} + \frac{m_2}{1+\pi}}_{\text{y-intercept}} - \overbrace{\frac{1+r}{1+\pi}}^{Slope} c_1$$

- Marginal Rate of Substitution:

$$MRS = -\frac{MU_x}{MU_y}$$

- **Cobb Douglas**
  - Functional Form:

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

  - First derivative with respect to x_1 (MU_1):

$$\frac{\partial U(x_1, x_2)}{\partial x_1} = A\alpha x_1^{\alpha-1} x_2^{\beta}$$

  - First derivative with respect to x_2 (MU_2):

$$\frac{\partial U(x_1, x_2)}{\partial x_2} = A\beta x_1^{\alpha} x_2^{\beta-1}$$

- **Linear**
  - Functional Form

$$U(x_1, x_2) = A x_1 + B x_2$$

  - First derivative with respect to x_1 (MU_1):

$$\frac{\partial U(x_1, x_2)}{\partial x_1} = A$$

  - First derivative with respect to x_2 (MU_2):

$$\frac{\partial U(x_1, x_2)}{\partial x_2} = B$$