Chapter 1
Microeconomics of Consumer Theory

The two broad categories of decision-makers in an economy are consumers and firms. Each individual in each of these groups makes its decisions in order to achieve some goal – a consumer seeks to maximize some measure of satisfaction from his consumption decisions while a firm seeks to maximize its profits. We first consider the microeconomics of consumer theory and will later turn to a consideration of firms. The two theoretical tools of consumer theory are utility functions and budget constraints. Out of the interaction of a utility function and a budget constraint emerge the choices that a consumer makes.

Utility Theory

A utility function describes the level of "satisfaction" or "happiness" that a consumer obtains from consuming various goods. A utility function can have any number of arguments, each of which affects the consumer's overall satisfaction level. But it is only when we consider more than one argument can we consider the *trade-offs* that a consumer faces when making consumption decisions. The nature of these trade-offs can be illustrated with a utility function of two arguments, but is completely generalizable to the case of any arbitrary number of arguments.[2]

Figure 2 illustrates in three dimensions the square-root utility function $u(c_1, c_2) = \sqrt{c_1} + \sqrt{c_2}$, where c_1 and c_2 are two different goods. This utility function displays **diminishing marginal utility** in *each* of the two goods, which means that, holding consumption of one good constant, increases in consumption of the other good increase total utility at ever-decreasing rates. Graphically, diminishing marginal utility means that the slope of the utility function with respect to each of its arguments in isolation is always decreasing.

The notion of diminishing marginal utility seems to describe consumers' preferences so well that most economic analysis takes it as a fundamental starting point. We will consider diminishing marginal utility a fundamental building block of all our subsequent ideas.

---

[2] An advantage of considering the case of just two goods is that we can analyze it graphically because, recall, graphing a function of two arguments requires three dimensions, graphing a function of three arguments requires four dimensions, and, in general, graphing a function of n arguments requires n+1 dimensions. Obviously, we cannot visualize anything more than three dimensions.
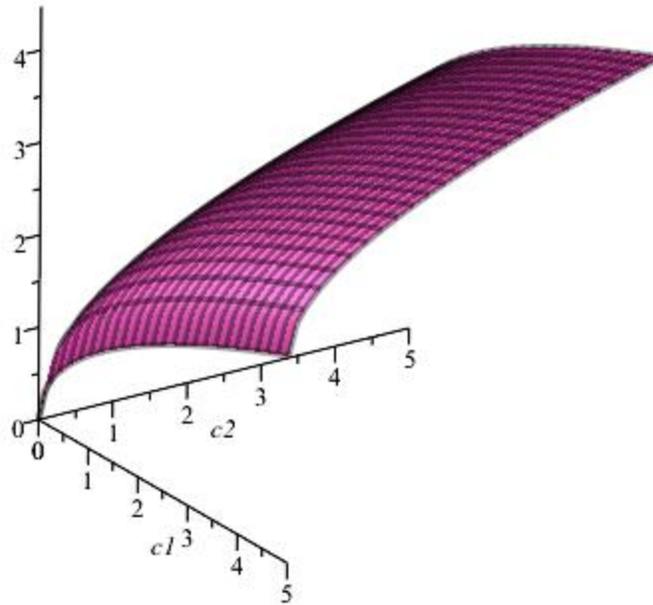
Figure 2. The utility surface as a function of two goods, c_1 and c_2. The specific utility function here is the square-root utility function, $u(c_1, c_2) = \sqrt{c_1} + \sqrt{c_2}$. The three axes are the c_1 axis, the c_2 axis, and the utility axis.

The first row of Figure 3 displays the same information as in Figure 2 except as a pair of two-dimensional diagrams. Each diagram is a rotation of the three-dimensional diagram in Figure 2, which allows for complete loss of depth perspective of either c_2 (the upper left panel) or of c_1 (the upper right panel).

The bottom row of Figure 3 contains the diminishing marginal utility functions with respect to c_1 (c_2), holding constant c_2 (c_1).
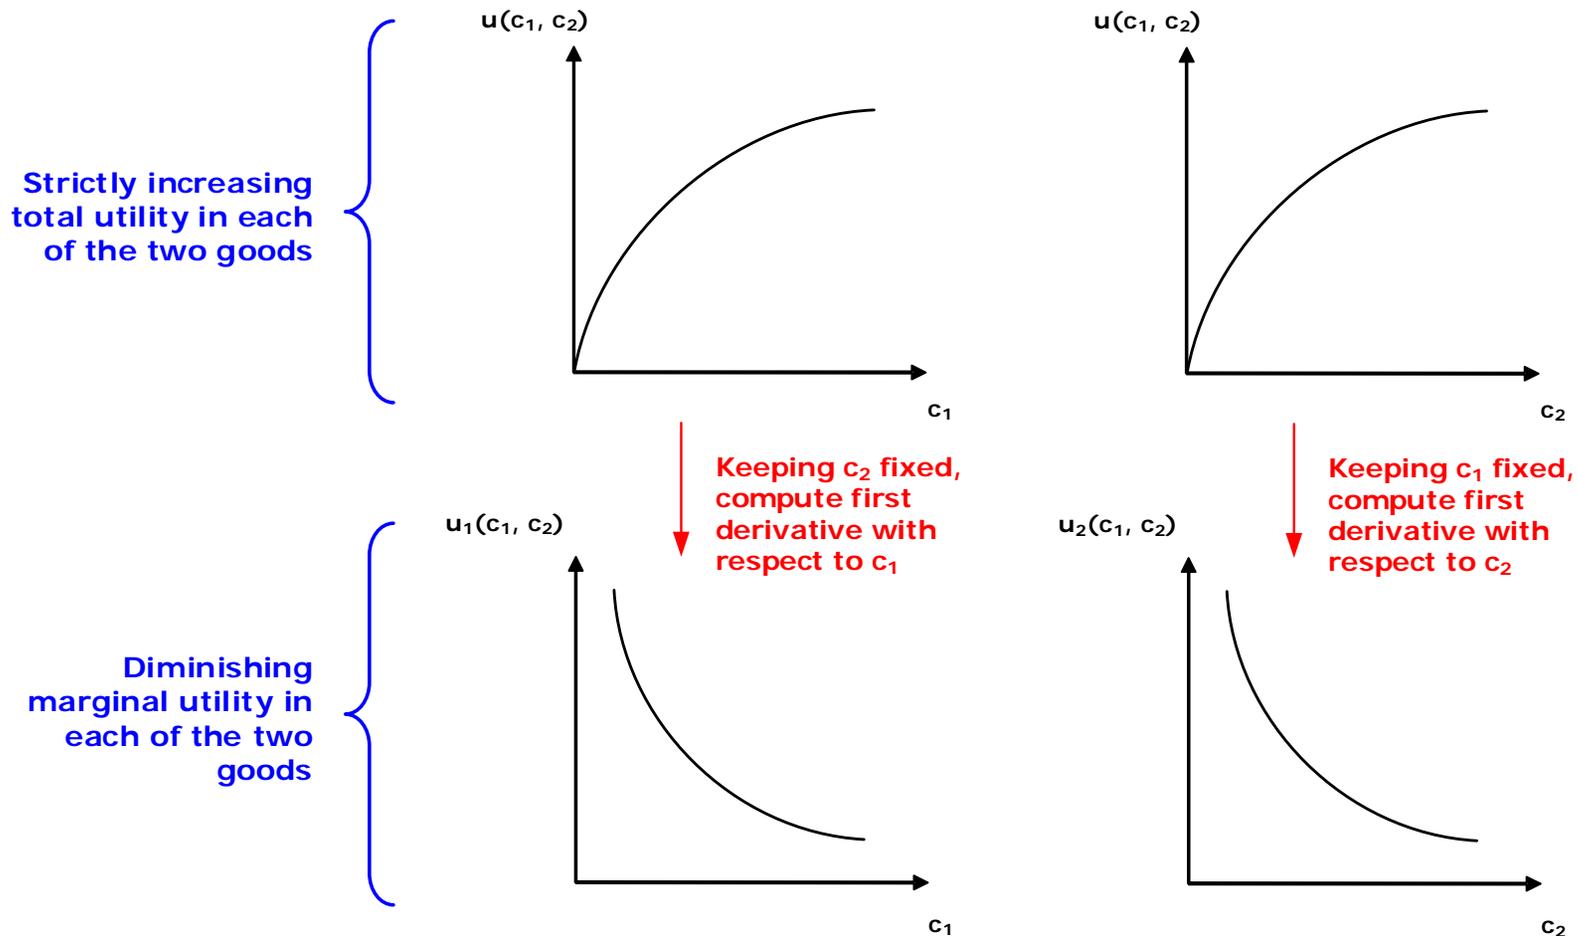
Figure 3. Top left panel: total utility as a function of c_1, holding fixed c_2. Top right panel: total utility as a function of c_2, holding fixed c_1. Lower left panel: (diminishing) marginal product function of c_1, holding fixed c_2. Lower right panel: (diminishing) marginal product function of c_2, holding fixed c_1. For the utility function $u(c_1, c_2) = \sqrt{c_1} + \sqrt{c_2}$, the marginal utility functions are $u_1(c_1, c_2) = (1/2) \cdot \left(1/\sqrt{c_1}\right)$ (lower left panel) and $u_2(c_1, c_2) = (1/2) \cdot \left(1/\sqrt{c_2}\right)$ (lower right panel).

Indifference Curves

Figure 4 returns to the three-dimensional diagram using the same utility function, with a different emphasis. Each of the solid curves in **Figure 4** corresponds to a particular level of utility. This three-dimensional view shows that a given level of utility corresponds to a given height of the function $u(c_1, c_2)$ above the $c_1 - c_2$ plane.[3]
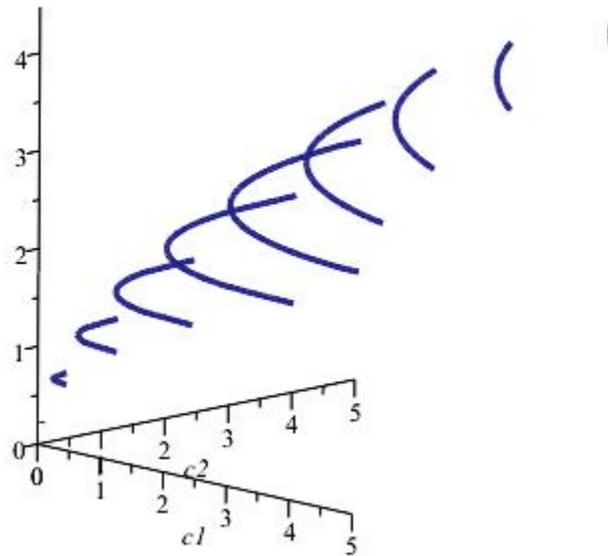

Figure 4. An indifference map of the utility function $u(c_1, c_2) = \sqrt{c_1} + \sqrt{c_2}$, where each solid curve represents a given height above the c_1-c_2 plane and hence a particular level of utility. The three axes are the c_1 axis, the c_2 axis, and the utility axis.

If we were to observe **Figure 4** from directly overhead, so that the utility axis were coming directly at us out of the $c_1 - c_2$ plane, we would observe **Figure 5**. **Figure 5** displays the contours of the utility function. In general, a **contour** is the set of all combinations of function arguments that yield some pre-specified function value. In our application here to utility theory, each contour is the set of all combinations of the two goods c_1 and c_2 that deliver a given level of utility. The contours of a utility function are called **indifference curves,** so named because each indifference curve shows all combinations

---

[3] Be sure you understand this last point very well.

(sometimes called "bundles") of goods between which a consumer is *indifferent* – that is, deliver a given amount of satisfaction. For example, suppose a consumer has chosen 4 units of c_1 and 9 units of c_2. The square-root utility function then tells us that his level of utility is $u(4,9) = \sqrt{4} + \sqrt{9} = 5$ (utils, which is the fictional measure of utility). There are an infinite number of combinations of c_1 and c_2, however, which deliver this same level of utility. For example, had the consumer instead been given 9 units of c_1 and 4 units of c_2, he would have obtained the same level of utility. That is, from the point of view of his overall level of satisfaction, the consumer is indifferent between having 4 units of good 1 in combination with 9 units of good 2 and having 9 units of good 1 in combination with 4 units of good 2. Thus these two points in the $c_1 - c_2$ plane lie on the same indifference curve.
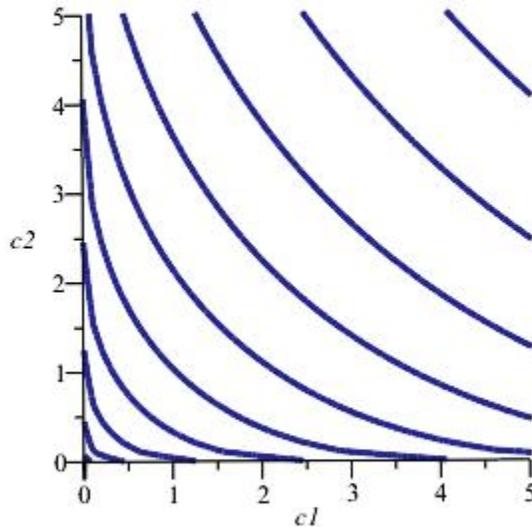

Figure 5. The contours of the utility function $u(c_1, c_2) = \sqrt{c_1} + \sqrt{c_2}$ viewed in the two-dimensional $c_1 - c_2$ plane. The utility axis is coming perpendicularly out of the page at you. Each contour of a utility function is called an indifference curve. Indifference curves further to the northeast are associated with higher levels of utility.

A crucial point to understand in comparing **Figure 4** and **Figure 5** is that indifference curves which lie further to the northeast in the latter correspond to higher values of the utility function in the former. That is, although we cannot actually "see" the height of the utility function in **Figure 5**, by comparing it to **Figure 4** we can conclude that indifference curves which lie further to the northeast provide higher levels of utility. Intuitively, this means

that if a consumer is given more of *both* goods (which is what moving to the northeast in the $c_1 - c_2$ plane means), his satisfaction is unambiguously higher.[4]

Once we understand that **Figure 4** and **Figure 5** are conveying the same information, it is clearly much easier to use the latter because drawing (variations of) **Figure 4** over and over again would be very time-consuming! As such, much of our study of consumer analysis will involve **indifference maps** such as that illustrated in **Figure 5**.

Marginal Rate of Substitution

Each indifference curve in **Figure 5** has a negative slope throughout. This captures the idea that, starting from any consumption bundle (that is, any point in the $c_1 - c_2$ plane), if a consumer gives up some of one good, *in order to maintain his level of utility* he must be given an additional amount of the other good. The crucial idea is that the consumer is willing to **substitute** one good for another, even though the two goods are not the same. Some reflection should convince you that this is a good description of most people's preferences. For example, a person who consumes two pizzas and five sandwiches in a month may be just as well off (in terms of total utility) had he consumed one pizza and seven sandwiches.[5]

The slope of an indifference curve tells us the **maximum number of units of one good the consumer is willing to substitute to get one unit of the other good.** This is an extremely important economic way of understanding what an indifference curve represents. The slope of an indifference curve varies depending on exactly which consumption bundle is under consideration. For example, consider the bundle ($c_1 = 3, c_2 = 2$), which yields approximately 3.15 utils using the square-root utility function above. If the consumer were asked how many units of c_2 he would be willing to give up in order to get one more unit of c_1, he would first consider the utility level (3.15 utils) he currently enjoys. Any final bundle that left him with less total utility would be rejected. He would be indifferent between his current bundle and a bundle with 4 units of c_1 that also gave him 3.15 total utils. Simply solving from the utility function, we have that $\sqrt{4} + \sqrt{c_2} = 3.15$, which yields (approximately) $c_2 = 1.32$. Thus, from the initial consumption bundle ($c_1 = 3, c_2 = 2$), the consumer is willing to trade at most 0.68 units of c_2 to obtain one more unit of c_1.

---

[4] You can probably readily think of examples where consuming more does not always leave a person better off. For example, after consuming a certain number of pizza slices and sodas, you will have probably had enough, to the point where consuming more pizza and soda would decrease your total utility (because it would make you sick, say). While this may be an important feature of preferences (the technical name for this phenomenon is "satiation"), for the most part we will be concerned with those regions of the utility function where utility is increasing. A way to justify this view is to suppose that the goods that we speak of are very broad categories of good, not very narrowly-defined ones such as pizza or soda.

[5] The key phrase here is "just as well off." Given our assumption above of increasing utility, he would *prefer* to have more pizzas *and* more sandwiches.

What if we repeated this thought experiment starting from the new bundle? That is, with $(c_1 = 4, c_2 = 1.32)$, what if we again asked the consumer how many units of c_2 he would be willing to give up to obtain yet another unit of c_1? Proceeding just as above, we learn that he would be willing to give up at most 0.48 units of c_2, giving him the bundle $(c_1 = 5, c_2 = 0.84)$, which yields total utility of 3.15.[6]

The preceding example shows that the more units of c_1 the consumer has, the fewer units of c_2 the consumer is willing to give up to get *yet another* unit of c_1. The economic idea here is that consumers have preferences for balanced consumption bundles – they do not like "extreme" bundles that feature very many units of one good and very few of another. Some reflection may also convince you that this feature of preferences is a good description of reality.[7] In more mathematical language, this feature of preferences leads to indifference curves that are **convex to the origin.**

Thus, the slope of the indifference curve has very important economic meaning. It represents the **marginal rate of substitution** between the two goods – the maximum quantity of one good that the consumer is willing to trade for one more unit of the other. Formally, the marginal rate of substitution at a particular consumption bundle is the negative of the slope of the indifference curve passing through that consumption bundle.

Budget Constraint

The cost side of a consumer's decisions involves the price(s) he must pay to obtain consumption. Again maintaining the assumption that there are only two types of consumption goods, c_1 and c_2, let P_1 and P_2 denote their prices, respectively, in terms of money. We will assume for simplicity for the moment that each consumer spends all of his income, denoted by Y, (more generally, all of his resources, which may also include wealth) on purchasing c_1 and c_2.[8] We further assume (for now) that he has no control over his income – he simply takes it as given.[9] The **budget constraint** the consumer must respect as he makes his choice about how much c_1 and c_2 to purchase is therefore

$$P_1 c_1 + P_2 c_2 = Y .$$

---

[6] Make sure you understand how we arrived at this.

[7] When we later consider how consumers make choices across time (as opposed to a specific point in time), we will call this particular feature of preferences the "consumption-smoothing" motive.

[8] Assuming this greatly simplifies the analysis yet does not alter any of the basic lessons to be learned. Indeed, if we allowed the consumer to "save for the future" so that he didn't spend of all of his current income on consumption, the additional choice introduced (consumption versus savings) would also be analyzed in using exactly the same procedure. We will turn to such "intertemporal choice" models of consumer theory shortly.

[9] Also very shortly, using the same tools of utility functions and budget constraints, we will study how an individual decides what his optimal level of income is.

The term $P_1 c_1$ is total expenditure on good 1 and the term $P_2 c_2$ is total expenditure on good 2, the sum of which is equal to (by our assumption above) income. If we solve this budget constraint for c2, we get

$$c_2 = -\frac{P_1}{P_2} c_1 + \frac{Y}{P_2},$$

which, when plotted in the $c_1 - c_2$ plane, gives the straight line in **Figure 6**. In this figure, for illustrative purposes, the prices are chosen to both equal one (that is, $P_1 = P_2 = 1$) so that the slope of the budget line is negative one, and income is arbitrarily chosen to be $Y = 5$. Obviously, when graphing a budget constraint, the particular values of prices and income will determine its exact location.
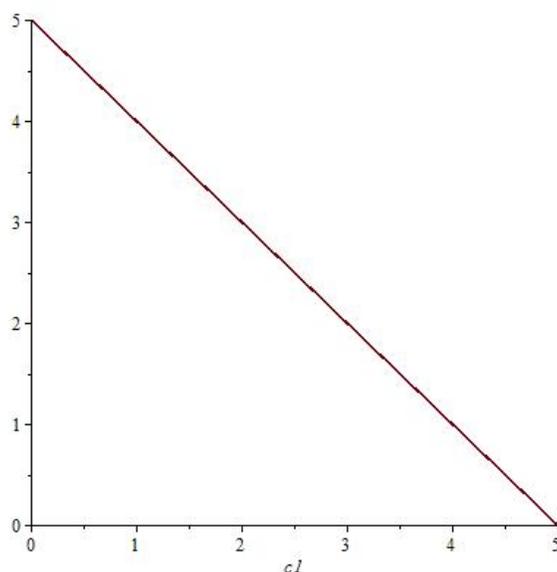

Figure 6. The budget constraint, plotted with c$_2$ as a function of c$_1$. For this example, the chosen prices are P$_1$ = P$_2$ = 1, and the chosen income is Y=5.

We discussed in our study of utility functions the idea that we needed three dimensions – the c_1 dimension, the c_2 dimension, and the utility dimension – to properly visualize utility. We see here that utility plays no role in the budget constraint, as it should not because the budget constraint only describes expenditures, not the benefits (i.e., utility) a consumer obtains from those expenditures. That is, the budget constraint is a concept completely independent of the concept of a utility function – this is a key point. We **could** graph the budget constraint in the same three-dimensional space as our utility function – it simply would be independent of utility. The graph of the budget constraint (which we call a budget plane when we construct it in three-dimensional space) in our $c_1 - c_2 - u$ space is shown in **Figure 7**.
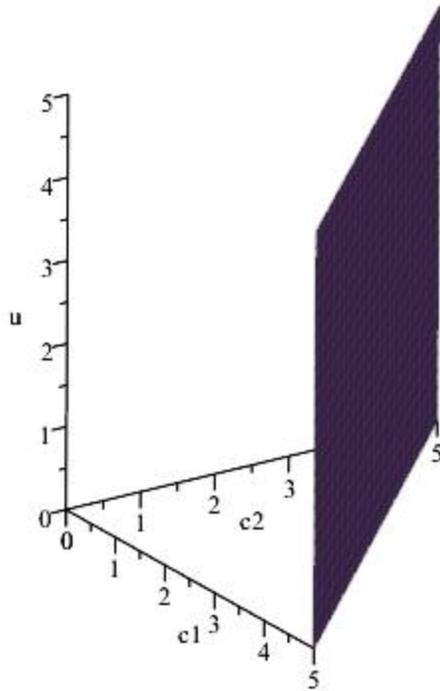
Figure 7. The budget constraint drawn in the three-dimensional c_1-c_2-u space. The budget constraint is a plane here because it is independent of utility.

Optimal Choice

We are now ready to consider how consumers make choices. The benefits of consumption are described by the utility function, and the costs of consumption are described by the budget constraint. Graphically, the decision the consumer faces is to choose that bundle (c_1, c_2) that yields the highest utility (i.e., lies on the highest indifference curve) that also satisfies his budget constraint (i.e., lies in the relevant budget plane).
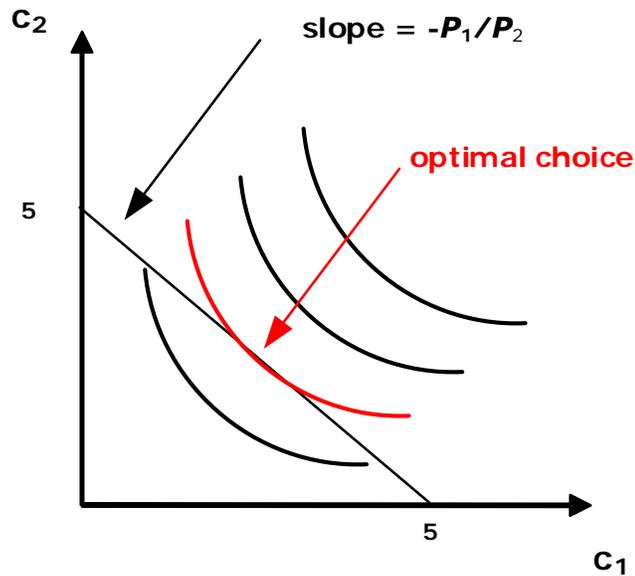
Figure 8. The optimal consumption choice features a tangency between the budget line and an indifference curve. The optimal choice must lie on the budget line and attain the highest possible utility for the consumer.

Imagine that both the budget constraint and the utility function were plotted in the three dimensions of **Figure 7** – and then imagine that we are observing that figure from directly overhead, so that the utility axis were coming straight out of the $c_1 - c_2$ plane at us, so that we lose perspective of the utility axis. What we would see are an indifference map and a budget line. **Figure 8** shows that the optimal decision (the one that yields the highest attainable utility) features a tangency between the budget constraint and an indifference curve. Consider what would happen if the optimal choice did not feature such a tangency. In this case, it must be that the indifference curve through which the chosen bundle passes also crosses the budget line at another point. Given that indifference curves are convex to the origin, this must mean that there is another consumption bundle that is both affordable and also yields strictly higher utility, so a rational consumer would choose it.[10]

At the point of tangency that describes the consumer's optimal choice, the slope of the budget line must equal the slope of the indifference curve. The slope of the budget line, as we saw above, is simply the price ratio $-P_1/P_2$. And recall from our discussion of utility functions that the (negative of the) slope of an indifference curve is the marginal rate of substitution – the maximum amount of one good that the consumer is willing to give up in order to obtain one more unit of the other good *such that* his total utility remains the same. These two points lead us to a very important description of a

---

[10] The assumption of a "rational" consumer must also be augmented with other strong assumptions, some of which are that there is no income uncertainty, prices are fixed and the consumer has no bargaining power, and there is no uncertainty about the quality or nature of the products. We will discuss some of these strong assumptions later.

consumer's optimal choice, one that we will refer to as the **consumer's optimality condition:**

$$\frac{P_1}{P_2} = MRS \, .$$

When markets are functioning well (and we have yet to discuss what "functioning well" means), this optimality condition is what guides the decisions of consumers. When markets are not functioning well, policy discussions at both the microeconomic level and the macroeconomic level can use this optimality condition as a benchmark to strive to achieve when considering intervening in markets.[11]

The economic logic of the optimality condition is as follows. Without regard to prices, the MRS describes the consumer's internal (i.e., utility-based) willingness to trade one good for another. The price ratio describes the market trade-off between the two goods. To understand this last point, suppose that $P_1 = \$3$ per unit of good 1, and $P_2 = \$2$ per unit of good 2.[12] The price ratio, therefore, keeping explicit track of units, is

$$\frac{P_1}{P_2} = \frac{\$3/\text{unit of good 1}}{\$2/\text{unit of good 2}} = \frac{3 \text{ units of good 2}}{2 \text{ units of good 1}}.$$

Notice the units here – it is units of good 2 per unit of good 1, which is exactly what it must be in our two-dimensional graph with c_1 on the horizontal axis and c_2 on the vertical axis. This demonstrates that the price ratio does indeed describe the market trade-off between the two goods.

Suppose the consumer had chosen a bundle at which his MRS was higher than the market price ratio. This means that he is willing to give up more units of c_2 for a little more c_1 than the market requires him to – so he should use the markets to trade some of his c_2 for c_1 because he would be made unambiguously better off! Now suppose he has traded himself in this way all the way to the point where his MRS equals the price ratio. Should he trade yet more units of c_2 to obtain a little more c_1? The answer is no, because doing so would now mean having to give up more units of c_2 than he willing to for a little more c_1. Thus, once he has traded his way to the bundle at which his MRS equals the price ratio, he can do no better – he has arrived at his optimal consumption choice, the one that maximizes $u(c_1, c_2)$ subject to his budget constraint.

---

[11] We will have much more to say later in the course about the role of government intervention in markets.
[12] It is very easy to lose sight of the fact that prices have units. That is, when a price tag on a T-shirt says "$10," the implicit units attached to this are "10 *per T-shirt*" – because obviously if you want to buy 2 T-shirts you will have to pay $10x2 = $20. Unit analysis is often helpful in thinking about how economic variables relate to each other.

Lagrange Characterization

Let's now study the optimality condition using our Lagrange tools. To cast the problem we are studying here into the mathematical form we encountered in our general introduction to the Lagrange method: the objective function (i.e., the function that the consumer seeks to maximize) is the utility function $u(c_1, c_2)$; the variables to be chosen are c_1 and c_2; and the maximization of utility is subject to the budget constraint $P_1c_1 + P_2c_2 = Y$. To cast the budget constraint into the form $g(.) = 0$, let's write $g(c_1, c_2) = Y - P_1c_1 - P_2c_2 = 0$.[13] The Lagrange function is thus

$$L(c_1, c_2, \lambda) = u(c_1, c_2) + \lambda[Y - P_1c_1 - P_2c_2],$$

where λ is the Lagrange multiplier. The first-order conditions with respect to c_1, c_2, and λ are, respectively,

$$\frac{\partial u}{\partial c_1} - \lambda P_1 = 0,$$

$$\frac{\partial u}{\partial c_2} - \lambda P_2 = 0,$$

and

$$Y - P_1c_1 - P_2c_2 = 0.$$

From the first two of these expressions, we can obtain the optimality condition we obtained qualitatively/graphically earlier. The first expression immediately above can be solved for the multiplier to give us $\lambda = \dfrac{\partial u / \partial c_1}{P_1}$. Next, insert this in the second expression immediately above, giving us

$$\frac{\partial u}{\partial c_2} = \frac{P_2 \cdot (\partial u / \partial c_1)}{P_1}.$$

Rearranging this in one more step gives us

$$\frac{P_1}{P_2} = \frac{\partial u / \partial c_1}{\partial u / \partial c_2},$$

which states that when the consumer is making the optimal choice between consumption of the two types of goods, his ratio of marginal utilities (the right-hand-side of this last

---

[13] Alternatively, we could equivalently construct $g(.)$ as $g(c_1, c_2) = P_1c_1 + P_2c_2 - Y = 0$, and we would obtain exactly the same result we are about to obtain; it would be a good exercise for you to try the subsequent manipulations for yourself using this alternate definition of the function $g(.)$..

expression) is equal to the ratio of prices of the two goods (the left-hand-side of this last expression). Compare this expression to the expression earlier that we named the optimality condition: inspecting the two reveals that they must be the same and furthermore that **the MRS between two goods is equal to the ratio of marginal utilities.** This latter very important result (that the MRS between any two goods is equal to the ratio of marginal utilities between those two goods) can be derived more rigorously mathematically, but we defer this. Instead, the important idea to understand here is how to apply the Lagrange method to the basic consumer optimization problem and how it yields the same intuitive result we qualitatively obtained earlier.

To link the result here back to our introduction to the Lagrange method, note that if we compute the partial derivatives of the constraint function (the budget constraint in this case) with respect to c_1 and c_2, we have $\partial g / \partial c_1 = -P_1$ and $\partial g / \partial c_2 = -P_2$, which means that the ratio of partial derivatives of the constraint function is $\dfrac{\partial g / \partial c_1}{\partial g / \partial c_2} = \dfrac{P_1}{P_2}$. But this is obviously just the left-hand-side of the optimality condition above. Recall in our introduction to Lagrange theory we noted that a central result was that **at optimal choices,** the ratio of partials of the objective function (here, the utility function) would be equal to the ratio of partials of the constraint function (here, the budget constraint): here is our first specific instance of this important result.