MODERN MACROECONOMICS

Sanjay K. Chugh

Modern Macroeconomics

Modern Macroeconomics

Sanjay K. Chugh

The MIT Press
Cambridge, Massachusetts
London, England

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu

This book was set in Times New Roman by Toppan Best-set Premedia Limited. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Chugh, Sanjay K.
Modern macroeconomics / Sanjay K. Chugh.
pages cm
Includes bibliographical references and index.
ISBN 978-0-262-02937-7 (hardcover : alk. paper) 1. Macroeconomics. 2. Keynesian economics.
3. Comparative economics.
I. Title.
HB172.5.C48 2015
339—dc23

2015009283

10 9 8 7 6 5 4 3 2 1

Contents

Acknowledgments

This textbook began as a collection of notes that I prepared to distribute to undergraduate students more than ten years ago while I was a graduate teaching assistant at the University of Pennsylvania. From the start, I organized the notes in "chapter" form because that made the notes appear coherent in the flow of thoughts and information. Writing these notes also helped me learn what I wanted to discuss with students, and allowed easy communication within the classroom. Or, rather, I think, at least easier than if I were just repeating phrasings and approaches of other textbooks.

Over the years, inevitably, the collection of notes grew, and many students (I hesitate today to even call them "students" because I learned a lot from them) have read various chapters and versions of the text. In reverse chronological order, these students were in classes I taught at Boston College, Boston University, the University of Maryland, Johns Hopkins University, Georgetown University, and the University of Pennsylvania. I thank all the students whose discussions contributed to the early chapters and, occasionally, brand-new drafts of chapters written on the fly.

I also thank all the department chairs at these institutions that permitted me, knowingly or not, to use my "notes" in the classroom rather than a "formal" textbook. Among them all, I owe Frank Weiss at Johns Hopkins an enormous debt of gratitude for his patience and encouragement in developing my notes over the past decade.

I further owe an enormous debt of gratitude to Allan Drazen at the University of Maryland for suggesting my name to Jane Macdonald at the MIT Press in 2011. Without his reference, this textbook would not have come to fruition.

When Jane Macdonald approached me and asked if I would be interested in developing my notes into a textbook, that was the moment I felt that I actually had in hand the makings of a textbook. I can't thank the MIT Press enough, and in particular Jane and Emily Taber for their support and encouragement in my completing the manuscript, for there is no other way my collection of notes could have turned into a textbook I also thank Dana Andrus at MIT Press for her spectacular editing skills, advice, and suggestions.

I have had many teaching assistants over the years who helped students get through various parts and early drafts of notes. There are again way too many to thank, but one

really stood out, Dominique Brabant at Boston College. Dominique helped tremendously while I was starting to resume work on the textbook. She went through the draft in early 2014 with a fine-toothed comb, offering lots of suggestions, comments, advice, and ways to maintain consistency across the chapters. It was Dominique's input that finally pushed me to begin rewriting the chapters I wrote at the very start of my teaching career.

Additionally I have had feedback from faculty members who used parts of my notes in their own classes at different universities; I hope this was productive for their students. And then there are the bits and pieces scattered throughout the text based on discussions I have had with fellow researchers, coauthors, friends, colleagues, and members of the economics profession. I thank them all for enriching throughout my thinking experience.

Sanjay K. Chugh
May 18, 2015

Introduction to Modern Macroeconomics

Modern macroeconomics is built explicitly on microeconomic foundations. That is, the modern study and analysis of macroeconomics begins by considering how the microeconomic units, namely consumers and firms, in an economy make their decisions and then considers how the choices of these great many individuals interact with each other to yield economy-wide outcomes. This approach sounds quite reasonable because, after all, it is individuals in a society that ultimately make decisions. However, it may surprise you that macroeconomics was not always studied this way. Indeed much of the evolution of macroeconomic theory occurred without any reference to its microfoundations. We, however, will consider the microeconomic foundations of macroeconomics—as such, our consideration of macroeconomics will mostly be a "modern" one.

The two most fundamental microeconomic units in any economy are consumers and firms. In introductory microeconomics, you studied how these individual units make their decisions. Under economists' usual assumption of rational behavior, the posited goal of consumers is to maximize their utility, and the posited goal of firms is to maximize their economic (as opposed to accounting) profits. Concepts such as marginal utility, marginal revenue, and marginal cost should be familiar to you from your introduction to microeconomics, and they will provide the foundation of our consideration of macroeconomics.

In modern industrialized economies, consumption activity (i.e., purchases of goods and services by individuals) constitutes the largest share of all macroeconomic activity. For example, in the United States, consumption accounts for roughly 70 percent of all economic activity. Understanding how consumers make decisions and the factors, especially government policies, that affect these decisions will be of prime importance in our study of macroeconomics. We thus begin our study of macroeconomics by reviewing the microeconomics of consumer theory in chapter 1. The tools introduced there will be used repeatedly, so it is important to grasp these ideas fully. Following this review of consumer theory, we will develop the macroeconomic theory of consumption, including the impact of various government policies on consumption behavior. After this, we will introduce firms into our theoretical model of the economy, again considering the impact of various government policies on firms' decisions.

We are potentially faced with one daunting task, however. It is obvious that each consumer is different from every other consumer in his preferences for goods and services, and it is equally obvious that firms are very different from one another, both in the goods and services they produce as well as the technologies that they use in producing those goods and services. In short, there is a great deal of heterogeneity in the economy. This poses a potentially intractable theoretical problem because it should strike you as impossible to model theoretically the choices of *every single individual* and *every single firm* in the economy. Quite apart from the fact that there is no way we could know the exact choices of every single microeconomic unit, the point of any theoretical model is to be a simplified description of some complicated phenomenon—if we had to try to determine the choices of every single microeconomic unit, we would not achieve any simplification at all!

One approach, then, is to categorize the individual microeconomic units into broad groups: for example, categorize consumers into "upper class," "middle class," and "lower class" and categorize firms into "goods-producing firms" and "service-producing firms." We could then consider how individuals in these different groups make their decisions, and subsequently "sum up" their choices to yield macroeconomic outcomes. This seems an appealing way of proceeding—it turns out, however, that even doing this becomes quite cumbersome theoretically. The details of the theoretical problems associated with this approach are left to more advanced courses in macroeconomics, but, briefly, the main problems have to do with defining the appropriate broad categories and then determining an appropriate way of "summing up" the individuals' choices.

We will instead adopt what is known as the **representative agent** paradigm. In the representative agent approach, we suppose that there are a great many consumers in the economy *each of whom is identical to all other consumers in every way* and that there are a great many firms in the economy *each of which is identical to all other firms in every way.* This is obviously a gross simplification of reality. However, adopting this approach has the virtue that it becomes much simpler to theoretically model macroeconomic outcomes. Of particular interest for our purposes is that it still allows us to consider the general effects of macroeconomic policies, although we will not be able to say which groups are hurt versus which groups benefit from any given policy (because, by construction, there are no distinct "groups" at all).

A simple example may help illustrate how we will use the representative agent approach. Suppose that there are five different consumers in an economy: in a given year, person A spends $50 on consumption, person B spends $75 on consumption, person C spends $100 on consumption, person D spends $125 on consumption, and person E spends $150 on consumption. The total dollar value of consumption in this economy in this year is thus $500. If we wanted to model every microeconomic unit, we would have to describe how each of persons A, B, C, D, and E made his decisions. However, if our main focus is on studying the total consumption of $500, we could equivalently suppose that there are five individuals in the economy *each of whom spent $100 on consumption.* That is, we could

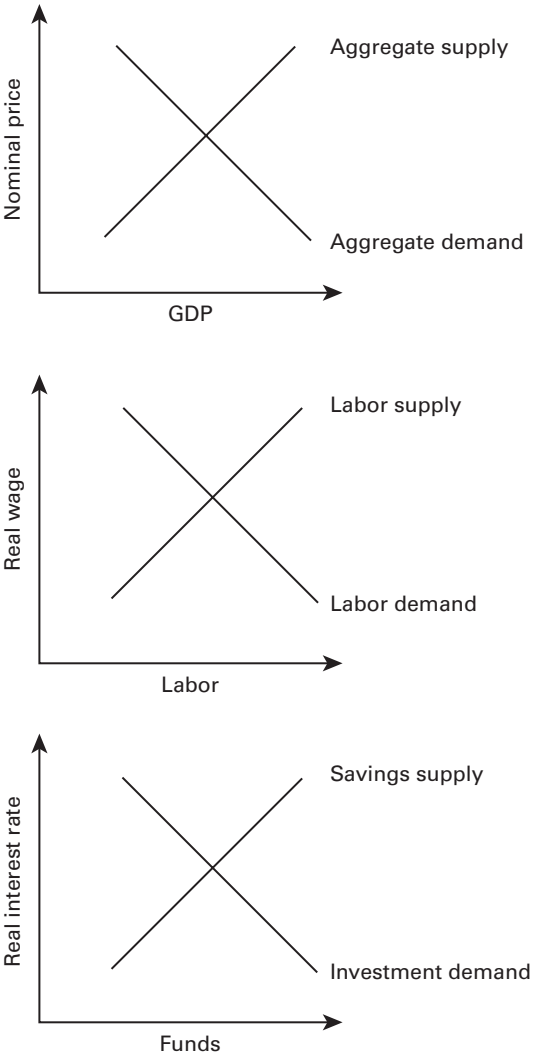

Three macro markets: goods and services markets, labor markets, and financial markets

suppose that each individual simply spent the economy-wide average on consumption. Then our task, at the microeconomic level, is to model just one individual, this "average consumer," because as soon as we know how he made his decisions we know the economy-wide outcome. This average consumer is exactly who the representative agent is. While seemingly a gross simplification of reality (as it is!), we will see that by modeling only this **representative consumer** in the economy we will be able to describe quite well many

macroeconomic outcomes and will also be able to consider the effects of macroeconomic policies.

Similarly we will also suppose that there is an "average firm" in the economy—the **representative firm.** This representative firm produces the average level of goods and services in the economy, guided by the usual principle of profit maximization familiar from introductory microeconomics. Once again, the way in which we model this representative firm will allow us to consider how firms respond to various macroeconomic policies.

In all to come, keep the following in mind: our goal is essentially to build a small theoretical model (using the representative agent paradigm) of the entire economy, one that includes consumers, firms, and the government. Putting these components together will allow us to see how they all interact with one another to yield macroeconomic outcomes and allow fairly rich consideration of the effects of macroeconomic policy, both fiscal policy (tax and spending initiatives of Congress) and monetary policy (control of interest rates and the money supply by the Federal Reserve). Throughout, we will be informed by basic microeconomic principles.

Our analysis will be concerned with demand, supply, and equilibrium in the "three macro markets," which are the aggregate goods and services market, the aggregate labor market, and the aggregate financial market depicted in the figure above. All of the demand and supply relationships are sketched as linear only for illustrative purposes.

Exogenous Variables versus Endogenous Variables

Before we begin, a crucial distinction to keep in mind throughout our study is that between **exogenous variables** and **endogenous variables.** In every particular framework and macro market we discuss, the exogenous variables are the *inputs* into the analysis. Exogenous variables are the ones that "are taken as given," as economic language so often puts it. In contrast, the endogenous variables are the *outputs* from the analysis conducted within the particular framework or market we are studying. Stated more mathematically, the endogenous variables are the ones that "need to be solved for," whether we're describing the consumer side of the economy or the firm side of the economy (or, for that matter, the government's role in the macroeconomy).

In each of the three macro markets as depicted in the figure, *prices are endogenously determined at the point at which economy-wide quantities demanded and economy-wide quantities supplied equate.* Of course, "distortions" arise in these perfect markets, and we will discuss many departures from perfect competition, but this diagram provides an important starting point.

Another important starting point is displayed in the next figure. The endogenous prices that arise in this figure are *exogenous ("taken as given") from the point of view of atomistic individuals actively participating in the markets,* be they individual consumers or individual businesses. Keep both figures in mind as we begin to construct our macroeconomic frameworks.

Each atomistic firm and each atomistic individual takes as given prices in markets. Prices are determined in equilibrium, hence are exogenous to atomistic firms and atomistic individuals.

Before we get into the foundations of modern macroeconomics, in chapter 1 we briefly review the microeconomics of consumer theory. Part I next takes us through the various building blocks of modern macro, not just on the consumer side but also with respect of firms and the government.

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 this case is completely generalizable to the case of any arbitrary number of arguments.[1]

Figure 1.1 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.

1. An advantage of considering the case of just two goods is that we can analyze it graphically. 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.

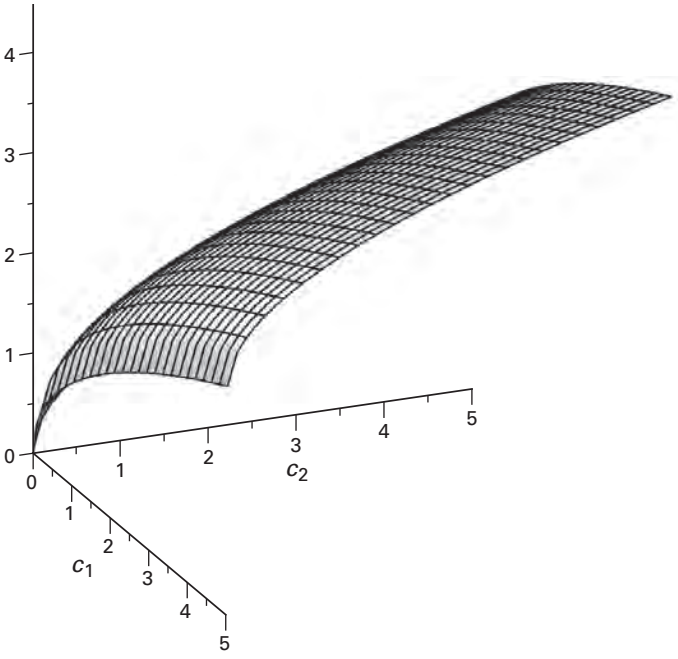

Figure 1.1
Utility surface as a function of two goods, c_1 and c_2. The specific utility function here is the square-root utility function, $u(c_1, c_2) = \sqrt{c_1} + \sqrt{c_2}$. The three axes are the c_1 axis, the c_2 axis, and the utility axis.

The 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.

The first row of figure 1.2 displays the same information as in figure 1.1 except as a pair of two-dimensional diagrams. Each diagram is a rotation of the three-dimensional diagram in figure 1.1, 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 1.2 contains the diminishing marginal utility functions with respect to c_1 (c_2), holding constant c_2 (c_1).

Indifference Curves

Figure 1.3 returns to the three-dimensional diagram using the same utility function, with a different emphasis. Each of the solid curves in figure 1.3 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.[2]

2. Be sure you understand this last point very well.

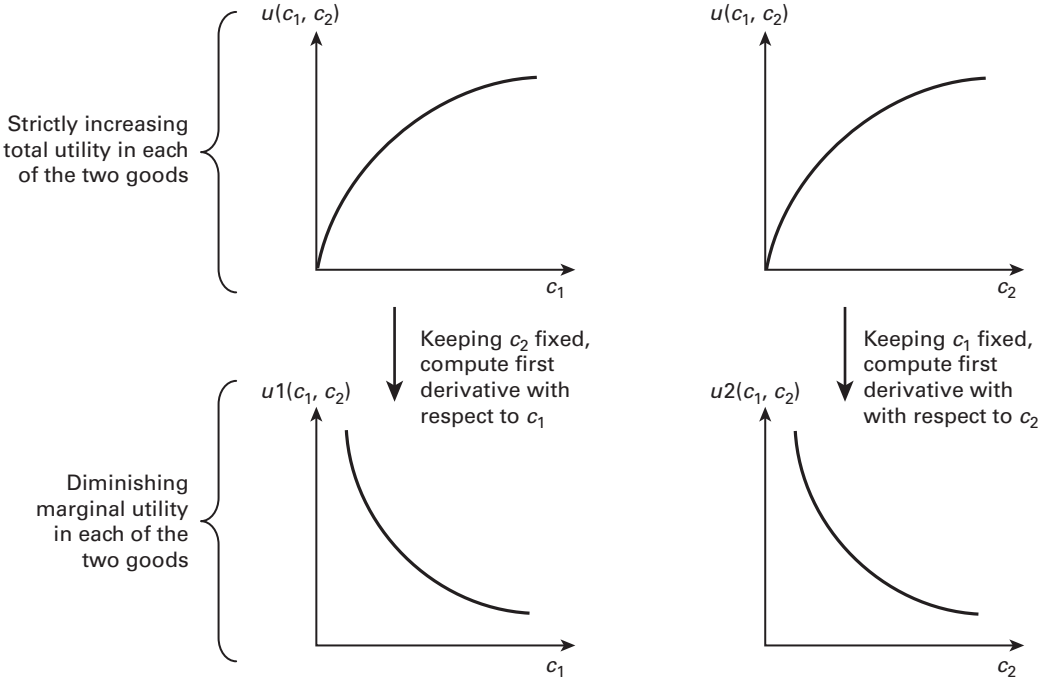

Figure 1.2
Top left: Total utility as a function of c_1, holding fixed c_2. *Top right*: Total utility as a function of c_2, holding fixed c_1. *Bottom left*: (Diminishing) marginal product function of c_1, holding fixed c_2. *Bottom right*: (Diminishing) marginal product function of c_2, holding fixed c_1. For the utility function $u(c_1,\ c_2) = \sqrt{c_1} + \sqrt{c_2}$, the marginal utility functions are $u_1(c_1,\ c_2) = (1/2)\cdot\left(1/\sqrt{c_1}\right)$ (*bottom left panel*) and $u_2(c_1,\ c_2) = (1/2)\cdot\left(1/\sqrt{c_2}\right)$ (*bottom right panel*).

If we were to observe figure 1.3 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 1.4. Figure 1.4 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. Here in our application 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 (sometimes called "bundles") of goods between which a consumer is *indifferent*—that is, deliver a given amount of satisfaction. For example, suppose that 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, that deliver this 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,

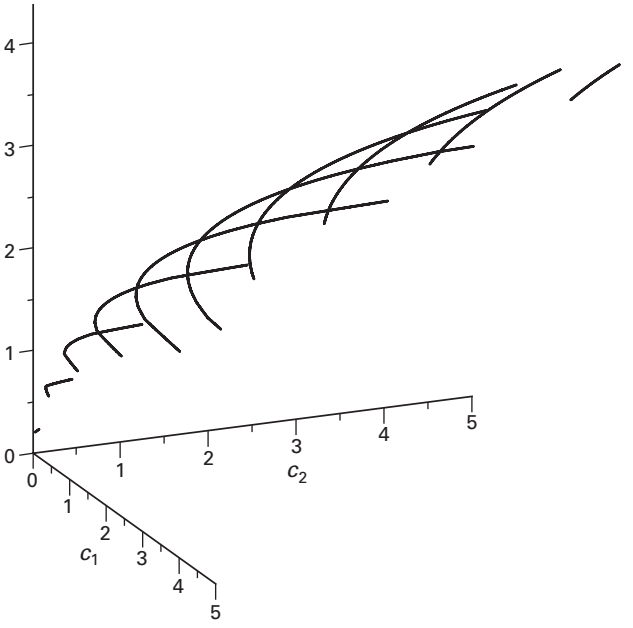

Figure 1.3
Indifference map of the utility function $u(c_1, c_2) = \sqrt{c_1} + \sqrt{c_2}$, where each solid curve represents a given height above the c_1–c_2 plane and hence a particular level of utility. The three axes are the c_1 axis, the c_2 axis, and the utility axis.

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.

A crucial point to understand in comparing figure 1.3 and figure 1.4 is that indifference curves that 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 1.4, by comparing it to figure 1.3, we can conclude that indifference curves that 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), then his satisfaction is unambiguously higher.[3]

3. You may 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 likely had enough, to the point where consuming more pizza and soda would decrease your total utility (i.e., it would make you sick). 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 goods, not very narrowly defined ones such as pizza or soda.

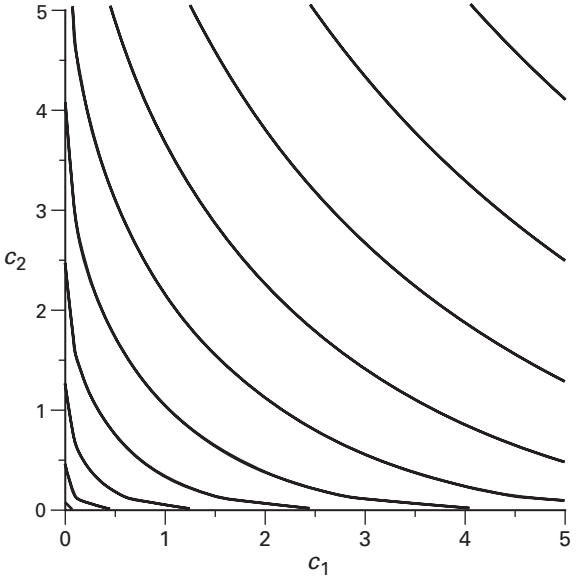

Figure 1.4
Contours of the utility function $u(c_1, c_2) = \sqrt{c_1} + \sqrt{c_2}$ viewed in the two-dimensional c_1–c_2 plane. The utility axis is coming perpendicularly out of the page at you. Each contour of a utility function is called an indifference curve. Indifference curves further to the northeast are associated with higher levels of utility.

Once we understand that figure 1.3 and figure 1.4 are conveying the same information, it is much easier to use the latter diagram because drawing (variations of) figure 1.3 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 1.4.

Marginal Rate of Substitution

Each indifference curve in figure 1.4 has a negative slope throughout. This captures the idea that starting from any consumption bundle (i.e., any point in the c_1–c_2 plane), when 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.[4]

4. 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.

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.

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 that 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.[5]

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.[6] 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. For sim-

5. Make sure you understand how we arrived at this.

6. 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.

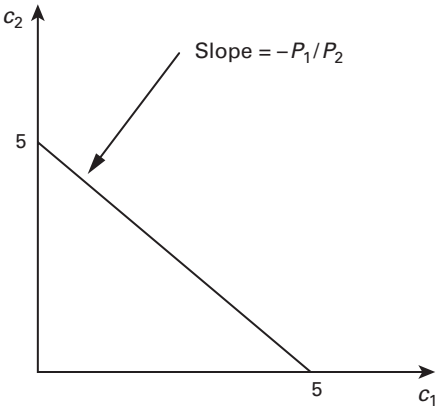

Figure 1.5
Budget constraint, plotted with c_2 as a function of c_1. For this example, the chosen prices are $P_1 = P_2 = 1$, and the chosen income is $Y = 5$.

plicity, we will assume 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.[7] We further assume (for now) that he has no control over his income—he simply takes it as given.[8] The **budget constraint** the consumer must respect as he makes his choice about how much c_1 and c_2 to purchase is therefore

$$P_1c_1 + P_2c_2 = Y.$$

The term P_1c_1 is total expenditure on good 1 and the term P_2c_2 is total expenditure on good 2, the sum of which is equal to income (by our assumption above). If we solve this budget constraint for c_2, 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 1.5. In this figure, for illustrative purposes, the prices are chosen to both equal one (i.e., $P_1 = P_2 = 1$) so that the slope of the budget line is a negative one, and income is arbitrarily chosen to be $Y = 5$.

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

8. 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.

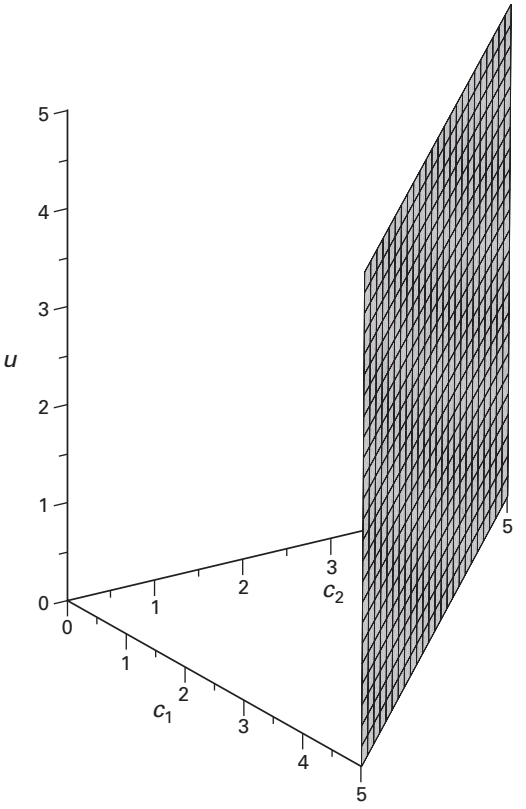

Figure 1.6
Budget constraint drawn in the three-dimensional c_1–c_2–u space. The budget constraint is a plane here because it is independent of utility.

Obviously, when graphing a budget constraint, the particular values of prices and income will determine its exact location.

We discussed in our study of utility functions the idea that we need 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 1.6.

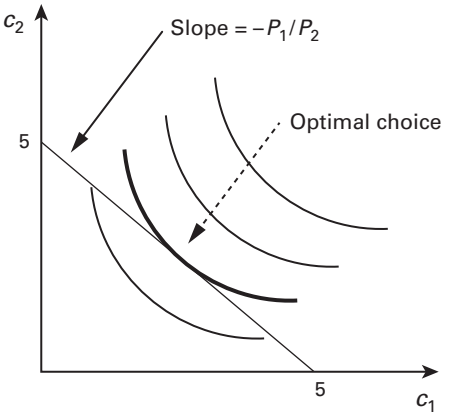

Figure 1.7
Optimal consumption choice displayed as a tangency between the budget line and an indifference curve. The optimal choice must lie on the budget line and attain the highest possible utility for the consumer.

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).

Imagine that both the budget constraint and the utility function were plotted in the three dimensions of figure 1.6—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 1.7 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 yields strictly higher utility, so a rational consumer would choose it.[9]

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

9. The assumption of a "rational" consumer must further be augmented by other strong assumptions, some of these being that there is no income uncertainty, prices are fixed, the consumer has no bargaining power, and no uncertainty exists as to the quality of the products. We will discuss some of these strong assumptions later.