

Notes on Microeconomic Theory

Contents

1 The Economic Approach	1
2 Consumer Theory Basics	5
2.1 Commodities and Budget Sets	5
2.2 Demand Functions	8
2.3 Three Restrictions on Consumer Choices	9
2.4 A First Analysis of Consumer Choices	10
2.4.1 Comparative Statics	11
2.5 Requirement 1 Revisited: Walras' Law	11
2.5.1 What's the Funny Equals Sign All About?	12
2.5.2 Back to Walras' Law: Choice Response to a Change in Wealth	13
2.5.3 Testable Implications	14
2.5.4 Summary: How Did We Get Where We Are?	15
2.5.5 Walras' Law: Choice Response to a Change in Price	15
2.5.6 Comparative Statics in Terms of Elasticities	16
2.5.7 Why Bother?	17
2.5.8 Walras' Law and Changes in Wealth: Elasticity Form	18
2.6 Requirement 2 Revisited: Demand is Homogeneous of Degree Zero	18
2.6.1 Comparative Statics of Homogeneity of Degree Zero	19
2.6.2 A Mathematical Aside ...	21
2.7 Requirement 3 Revisited: The Weak Axiom of Revealed Preference	22
2.7.1 Compensated Changes and the Slutsky Equation	23
2.7.2 Other Properties of the Substitution Matrix	27

3	The Traditional Approach to Consumer Theory	29
3.1	Basics of Preference Relations	30
3.2	From Preferences to Utility	35
3.2.1	Utility is an Ordinal Concept	36
3.2.2	Some Basic Properties of Utility Functions	37
3.3	The Utility Maximization Problem (UMP)	42
3.3.1	Walrasian Demand Functions and Their Properties	47
3.3.2	The Lagrange Multiplier	49
3.3.3	The Indirect Utility Function and Its Properties	50
3.3.4	Roy's Identity	52
3.3.5	The Indirect Utility Function and Welfare Evaluation	54
3.4	The Expenditure Minimization Problem (EMP)	55
3.4.1	A First Note on Duality	57
3.4.2	Properties of the Hicksian Demand Functions and Expenditure Function	59
3.4.3	The Relationship Between the Expenditure Function and Hicksian Demand	62
3.4.4	The Slutsky Equation	65
3.4.5	Graphical Relationship of the Walrasian and Hicksian Demand Functions	67
3.4.6	The EMP and the UMP: Summary of Relationships	71
3.4.7	Welfare Evaluation	72
3.4.8	Bringing It All Together	78
3.4.9	Welfare Evaluation for an Arbitrary Price Change	83
4	Topics in Consumer Theory	87
4.1	Homothetic and Quasilinear Utility Functions	87
4.2	Aggregation	89
4.2.1	The Gorman Form	90
4.2.2	Aggregate Demand and Aggregate Wealth	92
4.2.3	Does individual WARP imply aggregate WARP?	94
4.2.4	Representative Consumers	96
4.3	The Composite Commodity Theorem	98
4.4	So Were They Just Lying to Me When I Studied Intermediate Micro?	99
4.5	Consumption With Endowments	101

4.5.1	The Labor-Leisure Choice	104
4.5.2	Consumption with Endowments: A Simple Separation Theorem	106
4.6	Consumption Over Time	110
4.6.1	Discounting and Present Value	111
4.6.2	The Two-Period Model	112
4.6.3	The Many-Period Model and Time Preference	115
4.6.4	The Fisher Separation Theorem	118
5	Producer Theory	121
5.1	Production Sets	122
5.1.1	Properties of Production Sets	123
5.1.2	Profit Maximization with Production Sets	127
5.1.3	Properties of the Net Supply and Profit Functions	129
5.1.4	A Note on Recoverability	132
5.2	Production with a Single Output	133
5.2.1	Profit Maximization with a Single Output	134
5.3	Cost Minimization	136
5.3.1	Properties of the Conditional Factor Demand and Cost Functions	138
5.3.2	Return to Recoverability	139
5.4	Why Do You Keep Doing This to Me?	140
5.5	The Geometry of Cost Functions	141
5.6	Aggregation of Supply	142
5.7	A First Crack at the Welfare Theorems	143
5.8	Constant Returns Technologies	144
5.9	Household Production Models	148
5.9.1	Agricultural Household Models with Complete Markets	149
6	Choice Under Uncertainty	158
6.1	Lotteries	160
6.1.1	Preferences Over Lotteries	161
6.1.2	The Expected Utility Theorem	164
6.1.3	Constructing a vNM utility function	167
6.2	Utility for Money and Risk Aversion	168

6.2.1	Measuring Risk Aversion: Coefficients of Absolute and Relative Risk Aversion	173
6.2.2	Comparing Risk Aversion	174
6.2.3	A Note on Comparing Distributions: Stochastic Dominance	176
6.3	Some Applications	178
6.3.1	Insurance	178
6.3.2	Investing in a Risky Asset: The Portfolio Problem	180
6.4	Ex Ante vs. Ex Post Risk Management	182
7	Competitive Markets and Partial Equilibrium Analysis	185
7.1	Competitive Equilibrium	186
7.1.1	Allocations and Pareto Optimality	186
7.1.2	Competitive Equilibria	188
7.2	Partial Equilibrium Analysis	191
7.2.1	Set-Up of the Quasilinear Model	191
7.2.2	Analysis of the Quasilinear Model	192
7.2.3	A Bit on Social Cost and Benefit	195
7.2.4	Comparative Statics	195
7.3	The Fundamental Welfare Theorems	198
7.3.1	Welfare Analysis and Partial Equilibrium	201
7.4	Entry and Long-Run Competitive Equilibrium	205
7.4.1	Long-Run Competitive Equilibrium	206
7.4.2	Final Comments on Partial Equilibrium	209
8	Externalities and Public Goods	211
8.1	What is an Externality?	211
8.2	Bilateral Externalities	213
8.2.1	Traditional Solutions to the Externality Problem	217
8.2.2	Bargaining and Enforceable Property Rights: Coase's Theorem	219
8.2.3	Externalities and Missing Markets	221
8.3	Public Goods and Pure Public Goods	222
8.3.1	Pure Public Goods	223
8.3.2	Remedies for the Free-Rider Problem	227
8.3.3	Club Goods	229

8.4 Common-Pool Resources	229
9 Monopoly	233
9.1 Simple Monopoly Pricing	233
9.2 Non-Simple Pricing	236
9.2.1 Non-Linear Pricing	237
9.2.2 Two-Part Tariffs	239
9.3 Price Discrimination	241
9.3.1 First-Degree Price Discrimination	241
9.3.2 Second-Degree Price Discrimination	243
9.3.3 Third-Degree Price Discrimination	247
9.4 Natural Monopoly and Ramsey Pricing	249
9.4.1 Regulation and Incentives	251
9.5 Further Topics in Monopoly Pricing	252
9.5.1 Multi-Product Monopoly	252
9.5.2 Intertemporal Pricing	255
9.5.3 Durable Goods Monopoly	256

These notes are intended for use in courses in microeconomic theory taught at Harvard University. Consequently, much of the structure is inherited from the required text for the course, which is currently Mas-Colell, Whinston, and Green's Microeconomic Theory (referred to as MWG in these notes). They also draw on material contained in Silberberg's The Structure of Economics, as well as additional sources. They are not intended to stand alone or in any way replace the texts.

In the early drafts of this document, there will undoubtedly be mistakes. I welcome comments from students regarding typographical errors, just-plain errors, or other comments on how these notes can be made more helpful.

I thank Chris Avery, Lori Snyder, and Ben Sommers for helping clarify these notes and finding many errors.

Chapter 1

The Economic Approach

Economics is a social science.¹ Social sciences are concerned with the study of human behavior. If you asked the next person you meet while walking down the street what defines the difference between economics and other social sciences, such as political science or sociology, that person would most likely say that economics studies money, interest rates, prices, profits, and the like, while political science considers politicians, elections, etc., and sociology studies the behavior of groups of people. However, while there is certainly some truth to this statement, the things that can be fairly called economics are not so much defined by a subject matter as they are united by a common approach to problems. In fact, economists have written on topics spanning human behavior, from traditional studies of firm and consumer behavior, interest rates, inflation and unemployment to less traditional topics such as social choice, voting, marriage, and family.

The feature that unites these studies is a common approach to problems, which has become known as the “marginalist” or “neoclassical” approach. In a nutshell, the marginalist approach consists of four principles:

1. Economic actors have preferences over allocations of the world’s resources. These preferences remain stable, at least over the period of time under study.²
2. There are constraints placed on the allocations that a person can achieve by such things as wealth, physical availability, and social/political institutions.
3. Given the limits in (2), people choose the allocation that they most prefer.

¹See Silberberg’s Structure of Economics for a more extended discussion along these lines.

²Often preferences that change can be captured by adding another attribute to the description of an allocation.

More on this later.

4. Changes in the allocations people choose are due to changes in the limits on available resources in (2).

The marginalist approach to problems allows the economist to derive predictions about behavior which can then, in principle, be tested against real world data using statistical (econometric) techniques. For example, consider the problem of what I should buy when I go to the grocery store. The grocery store is filled with different types of food, some of which I like more and some of which I like less. Principle (1) says that the trade-offs I am willing to make among the various items in the store are well-defined and stable, at least over the course of a few months. An allocation is all of the stuff that I decide to buy. The constraints (2) put on the allocations I can buy include the stock of the items in the store (I can't buy more bananas than they have) and the money in my pocket (I can't buy bananas I can't afford). Principle (3) says that given my preferences, the amount of money in my pocket and the stock of items in the store, I choose the shopping cart full of stuff that I most prefer. That is, when I walk out of the store, there is no other shopping cart full of stuff that I could have purchased that I would have preferred to the one I did purchase. Principle (4) says that if next week I buy a different cart full of groceries, it is because either I have less money, something I bought last week wasn't available this week, or something I bought this week wasn't available last week.

There are two natural objections to the story I told in the last paragraph, both of which point toward why doing economics isn't a trivial exercise. First, it is not necessarily the case that my preferences remain stable. In particular, it is reasonable to think that my preferences this week will depend on what I purchased last week. For example, if I purchased chocolate chip cookies last week, this may make me less likely to purchase them this week and more likely to purchase some other sort of cookie. Thus, preferences may not be stable over the time period we are studying. Economists deal with this problem in two ways. The first is by ignoring it. Although widely applied, this is not the best way to address the problem. However, there are circumstances where it is reasonable. Many times changes in preferences will not be important relative to the phenomenon we are studying. In this case it may be more trouble than it's worth to address these problems. The second way to address the problem is to build into our model of preferences the idea that what I consumed last week may affect my preferences over what I consume this week. In other words, the way to deal with the cookie problem is to define an allocation as "everything I bought this week and everything I bought last week." Seen in this light, as long as the effect of having chocolate chip cookies last week on my preferences this week stay stable over time, my

preferences stay stable, whether or not I actually had chocolate chip cookies last week. Hence if we define the notion of preferences over a rich enough set of allocations, we can usually find preferences that are stable.

The second problem with the four-step marginalist approach outlined above is more troublesome: Based on these four steps, you really can't say anything about what is going to happen in the world. Merely knowing that I optimize with respect to stable preferences over the groceries I buy, and that any changes in what I buy are due to changes in the constraints I face does not tell me anything about what I will buy, what I should buy, or whether what I buy is consistent with this type of behavior.

The solution to this problem is to impose structure on my preferences. For example, two common assumptions are to assume that I prefer more of an item to less³ (monotonicity) and that I spend my entire grocery budget in the store (Walras' Law). Another common assumption is that only real opportunities matter. If I were to double all of the prices in the store and my grocery budget, this would not affect what I can buy, so it shouldn't affect what I do buy.

Once I have added this structure to my preferences, I am able to start to make predictions about how I will behave in response to changes in the environment. For example, if my grocery budget were to increase, I would buy more of at least one item (since I spend all of my money and there is always some good that I would like to add to my grocery cart).⁴ This is known as a testable implication of the theory. It is an implication because if the theory is true, I should react to an increase in my budget by buying more of some good. It is testable because it is based on things which are, at least in principle, observable. For example, if you knew that I had walked into the grocery store with more money than last week and that the prices of the items in the store had not changed, and yet I left the store with less of every item than I did last week, something must be wrong with the theory.

The final step in economic analysis is to evaluate the tests of the theories, and, if necessary, change them. We assume that people follow steps 1 - 4 above, and we impose restrictions that we believe are reasonable on their preferences. Based on this, we derive (usually using math) predictions about how they should behave and formulate testable hypotheses (or refutable propositions) about how they should behave if the theory is true. Then we observe what they really do. If

³Or, we could make the weaker assumption that no matter what I have in my cart already, there is something in the store that I would like to add to my cart if I could.

⁴The process of deriving what happens to people's choices (the stuff in the cart) in response to changes in things they do not choose (the money available to spend in the store) is known as comparative statics.

their behavior accords with our predictions, we rejoice because the real world has supported (but not proven!) our theory. If their behavior does not accord with our predictions, we go back to the drawing board. Why didn't their behavior accord with our predictions? Was it because their preferences weren't like we thought they were? Was it because they weren't optimizing? Was it because there was an additional constraint that we didn't understand? Was it because we did not account for a change in the environment that had an important effect on people's behavior?

Thus economics can be summarized as follows: It is the social science that attempts to account for human behavior as arising from consistent (often maximizing — more on that later) behavior subject to one or more constraints. Changes in behavior are attributed to changes in the constraints, and the test of these theories is to compare the changes in behavior predicted by the theory with the changes that actually occur.

Chapter 2

Consumer Theory Basics

Recall that the goal of economic theory is to account for behavior based on the assumption that actors have stable preferences, attempt to do as well as possible given those preferences and the constraints placed on their resources, and that changes in behavior are due to changes in these constraints. In this section, we use this approach to develop a theory of consumer behavior based on the simplest assumptions possible. Along the way, we develop the tool of comparative statics analysis, which attempts to characterize how economic agents (i.e. consumers, firms, governments, etc.) react to changes in the constraints they face.

2.1 Commodities and Budget Sets

To begin, we need a description of the goods and services that a consumer may consume. We call any such good or service a commodity. We number the commodities in the world 1 through L (assuming there is a finite number of them). We will refer to a “generic” commodity as l (that is, l can stand for any of the L commodities) and denote the quantity of good l by x_l . A commodity bundle (i.e. a description of the quantity of each commodity) in this economy is therefore a vector

$x = (x_1, x_2, \dots, x_L)$. Thus if the consumer is given bundle $x = (x_1, x_2, \dots, x_L)$, she is given x_1 units of good 1, x_2 units of good 2, and so on.¹ We will refer to the set of all possible allocations as the

commodity space, and it will contain all possible combinations of the L possible commodities.²

Notice that the commodity space includes some bundles that don’t really make sense, at least

¹For simplicity of terminology - but not because consumers are more or less likely to be female than male - we will call our consumer “she,” rather than “he/she.”

²That is, the commodity space is the L -dimensional real space \mathbb{R}^L .

economically. For example, the commodity space includes bundles with negative components. And, it includes bundles with components that are extremely large (i.e., so large that there simply aren't enough units of the relevant commodities for a consumer to actually consume that bundle). Because of this, it is useful to have a (slightly) more limited concept than the commodity space that captures the set of all realistic consumption bundles. We call the set of all reasonable bundles the consumption set, denoted by X . What exactly goes into the consumption set depends on the exact situation under consideration. In most cases, it is important that we eliminate the possibility of consumption bundles containing negative components. But, because consumers usually have limited resources with which to purchase commodity bundles, we don't have to worry as much about very large bundles. Consequently, we will, for the most part, take the consumption set to be the L dimensional non-negative real orthant, denoted R_+^L . That is, the possible bundles available for the consumer to choose from include all vectors of the L commodities such that every component is non-negative.

The consumption set eliminates the bundles that are "unreasonable" in all circumstances. We are also interested in considering the set of bundles that are available to a consumer at a particular time. In many cases, this corresponds to the set of bundles the consumer can afford given her wealth and the prices of the various commodities. We call such sets (Walrasian) budget sets.³ Let w stand for the consumer's wealth and p_l stand for the price of commodity l . Without any exceptions that I can think of, we assume that $p_l \geq 0$ for all l and that $w \geq 0$. That is, prices and wealth are either positive or zero, but not negative.⁴ We will let $p = (p_1, \dots, p_L)$ stand for the vector of prices of each of the goods. Hence if the consumer purchases consumption bundle x and the price vector is p , the consumer will spend

$$p \cdot x = \sum_{l=1}^L p_l x_l$$

on commodities.⁵ Since the consumer's total income is w , the consumer's Walrasian budget set is

³We call the budget set Walrasian after economist Leon Walras (1834-1910), one of the founders of this type of analysis.

⁴What do you imagine would happen if there were goods with negative prices?

⁵A few words about notation: In the above equation, x and p are both vectors, but they lack the usual notation \vec{x} and \vec{p} . Since economists almost never use the formal vector notation, you will need to use the context to judge whether an "x" is a single variable or actually a vector. Frequently we'll write something with subscript l to denote a particular commodity. Then, when we want to talk about all commodities, we put them together into a vector, which has no subscript. For example, p_l is the price of commodity l , and $p = (p_1, \dots, p_L)$ is the vector containing the prices of all commodities.




Figure 2.1: The Budget Set

defined as all bundles x such that $p \cdot x \leq w$ - in other words, all affordable bundles given prices and wealth. More formally, we can write the budget set as:

$$B_{p,w} = \{ x \in \mathbb{R}^L : p \cdot x \leq w \}$$

The term Walrasian is appended to the budget set to remind us that we are implicitly speaking of an environment where people can buy as much as they want of any commodity at the same price. In particular, this rules out the situations where there are limits on the amount of a good that a person can buy (rationing) or where the price of a good depends on how much you buy. Thus the Walrasian budget set corresponds to the opportunities available to an individual consumer whose consumption is small relative to the size of the total market for each good. This is just the standard “price taking” assumption that is made in models of competitive markets.

In order to understand budget sets, it is useful to assume that there are two commodities. In this case, the budget set can be written as

$$B_{p,w} = \{ x \in \mathbb{R}^2 : p_1x_1 + p_2x_2 \leq w \}$$

Or, if you plot x_2 on the vertical axis of a graph and x_1 on the horizontal axis, $B_{p,w}$ is defined by the set of points below the line $x_2 = \frac{-p_1x_1}{p_2} + \frac{w}{p_2}$. See Figure 2.1.

How does the budget set change as the prices or income change? If income increases, budget line AB shifts outward, since the consumer can purchase more units of the goods when she has more wealth. If the price of good 1 increases, when the consumer purchases only good 1 she can afford fewer units. Hence if p_1 increases, point B moves in toward the origin. Similarly, if p_2 increases, point A moves in toward the origin.

Exercise 1 Here is a task to show that you understand budget sets: Show that the effect on a budget set of doubling p_1 and p_2 is the same as the effect of cutting w in half. This is an illustration of the key economic concept that only relative prices matter to a consumer, which we will see over and over again.⁶

Now that we have defined the set of consumption bundles that the consumer can afford, the next step is to try to figure out which point the consumer will choose from the budget set. In order to determine which point from the budget set the consumer will choose, we need to know something about the consumer's preferences over the commodities. For example, if x_1 is onions and x_2 is chocolate, the consumer may prefer points with relatively high values of x_2 and low values of x_1 (unless, of course, p_2 is very large relative to p_1). If we knew exactly the trade-offs that the consumer is willing to make between the commodities, their prices, and the consumer's income, we would be able to say exactly which consumption bundle the consumer prefers. However, at this point we do not want to put this much structure on preferences.

2.2 Demand Functions

Now we need to develop a notation for the consumption bundle that a consumer chooses from a particular budget set. Let $p = (p_1, \dots, p_L)$ be the vector of prices of the L commodities. We will assume that all prices are non-negative. When prices are p and wealth is w , the set of bundles that the consumer can afford is given by the Walrasian budget set $B_{p,w}$. Assume that for any price vector and wealth (p, w) there is a single bundle in the budget set that the consumer chooses. Let $x_i(p, w)$ denote the quantity of commodity i that the consumer chooses at these prices and wealth. Let $x(p, w) = (x_1(p, w), \dots, x_L(p, w)) \in B_{p,w}$ denote the bundle (vector of commodities) that the consumer chooses when prices are p and income is w . That is, it gives the optimal consumption bundle as a function of the price vector and wealth. To make things easier, we will assume that $x_i(p, w)$ is single-valued (i.e. a function) and differentiable in each of its arguments.

Exercise 2 How many arguments does $x_i(p, w)$ have? Answer: $L + 1$: L prices and wealth.

Functions $x_i(p, w)$ represent the consumer's choice of commodity bundle at a particular price and wealth. Because of this, they are often called choice functions. They are also called demand

⁶The idea that only relative prices matter goes by the mathematical name "homogeneity of degree zero", but we'll return to that later.

functions, although sometimes that name is reserved for choice functions that are derived from the utility-maximization framework we'll look at later. Generally, I use the terms interchangeably, except when I want to emphasize that we are not talking about utility maximization, in which case I'll use the term "choice function."

At this point, we should introduce an important distinction, the distinction between endogenous and exogenous variables. An endogenous variable in an economic problem is a variable that takes its value as a result of the behavior of one of the economic agents within the model. So, the consumption bundle the consumer chooses $x(p, w)$ is endogenous. An exogenous variable takes its value from outside the model. Exogenous variables determine the constraints on the consumer's behavior. Thus in the consumer's problem, the exogenous variables are prices and wealth. The consumer cannot choose prices or wealth. But, prices and wealth determine the budget set, and from the budget set the consumer chooses a consumption bundle. Hence the consumption bundle is endogenous, and prices and wealth are exogenous. The consumer's demand function $x(p, w)$ therefore gives the consumer's choice as a function of the exogenous variables.

One of the main activities that economists do is try to figure out how endogenous variables depend on exogenous variables, i.e., how consumers' behavior depends on the constraints placed on them (see principles 1-4 above).

2.3 Three Restrictions on Consumer Choices

So, let's begin with the following question: What are the bare minimum requirements we can put on behavior in order for them to be considered "reasonable," and what can we say about consumers' choices based on this? It turns out that relatively weak assumptions about consumer behavior can generate strong requirements for how consumers should behave.⁷ We will start by enumerating three requirements.

- Requirement 1: The consumer always spends her entire budget (Walras' Law).

Requirement 1 is reasonable only if we are willing to make the assumption that "more is better." That is, for any commodity bundle x , the consumer would rather have a bundle with at least as much of all commodities and strictly more of at least one commodity. Actually, we can get away

⁷A "weak assumption" imposes less restriction on the behavior of an economic agent than a "strong assumption" does, so when designing a model, we prefer to use weaker assumptions if possible.

with a weaker assumption: Given any bundle x , there is always a bundle that has more of at least one commodity that the consumer strictly prefers to x . We'll return to this later. For now, just remember that the consumer spends all of her budget.

- Requirement 2: Only real opportunities matter (demand is homogeneous of degree zero).

The essence of requirement 2 is that consumers care about wealth and prices only inasmuch as they affect the set of allocations in the budget set. Or, to put it another way, changes in the environment that do not affect the budget set should not affect the consumer's choices. So, for example, if you double each price and wealth, the budget set is unchanged. Hence the consumer can afford the same commodity bundles as before and should choose the same bundle as before.

- Requirement 3: Choices reveal information about (stable) preferences.

So, suppose I offer you a choice between an apple and a banana, and you choose an apple. Then if tomorrow I see you eating a banana, I can infer that you weren't offered an apple (remember we assume that your preferences stay constant). Requirement 3 is known as the Weak Axiom of Revealed Preference (WARP). The essence is this. Suppose that on occasion 1 you chose bundle x when you could have chosen y . If I observe that on occasion 2 you choose bundle y , it must be because bundle x was not available. Put slightly more mathematically, suppose two bundles x and y are in the budget set $B_{p,w}$ and the consumer chooses bundle x . Then if at some other prices and wealth (p^0, w^0) the consumer chooses y , it must be that x was not in the budget set B_{p^0, w^0} . We'll return to WARP later, but you can think of it in this way. If the consumer's preferences remain constant over time, then if x is preferred to y once, it should always be preferred to y . Thus if you observe the consumer choose y , you can infer from this choice that x must not have been available. Or, to put it another way, if you observe the consumer choosing x when x and y were available on one day and y when x and y were available on the next day, then your model had better have something in it to account for why this is so (i.e., a reason why the two days were different).

2.4 A First Analysis of Consumer Choices

In the rest of this chapter, we'll develop formal notation for talking about consumer choices, show how the three requirements on consumer behavior can be represented using this notation, and determine what imposing these restrictions on consumer choices implies about the way consumers

should behave when prices or wealth change. Thus it is our first pass at the four-step process of economics: Assume consumers make choices that satisfy certain properties (the three requirements), subject to some constraints (the budget set); assume further that any changes in choices are due to changes in the constraints; and then derive testable predictions about consumer's behavior.

2.4.1 Comparative Statics

The analytic method we will use to develop testable predictions is what economists call comparative statics. A comparative statics analysis consists of coming up with a relationship between the exogenous variables and the endogenous variables in a problem and then using calculus to determine how the endogenous variables (i.e., the consumer's choices) respond to changes in the exogenous variables. Then, hopefully, we can tell if this response is positive, negative, or zero.⁸ We'll see comparative statics analysis used over and over again. The important thing to remember for now is that even though "comparative statics" as a phrase doesn't mean anything, it refers to figuring out how the endogenous variables depend on the exogenous variables.⁹

2.5 Requirement 1 Revisited: Walras' Law

Requirement 1 for consumer choices is that consumers spend all of their wealth (Walras Law). The implication of this is that given a budget set $B_{p,w}$, the consumer will choose a bundle on the boundary of the budget set, sometimes called the budget frontier. The equation for the budget frontier is the set of all commodity bundles that cost exactly w . Thus, Walras' Law implies:

$$p \cdot x(p, w) \equiv w.$$

When a consumer's demand function $x(p, w)$ satisfies this identity for all values of p and w , we say that the consumer's demand satisfies Walras' Law. Thus the formal statement for "consumers always spend all of their wealth" is that "demand functions satisfy Walras' Law."

⁸Although it would be nice to get a more precise measurement of the effects of changes in the exogenous parameters, often we are only able to draw implications about the sign of the effect, unless we are willing to impose additional restrictions on consumer demand.

⁹The term "comparative statics" is meant to convey the idea that, while you analyze what happens before and after the change in the exogenous parameter, you don't analyze the process by which the change takes place.

2.5.1 What's the Funny Equals Sign All About?

Notice that in the expression of Walras' Law, I wrote a funny, three-lined equals sign. Contrary to popular belief, this doesn't mean "really, really equal." What it means is that, no matter what values of p and w you choose, this relationship holds. For example, consider the equality:

$$2z = 1.$$

This is true for exactly one value of z , namely $z = \frac{1}{2}$. However, think about the following equality:

$$2z = a.$$

Suppose I were to ask you, for any value of a , tell me a value of z that makes this equality hold. You could easily do this: $z = \frac{a}{2}$. Suppose I denote this by $z(a) = \frac{a}{2}$. That is, $z(a)$ is the value of z that makes $2z = a$ true, given any value of a . If I substitute the function $z(a)$ into the expression $2z = a$, I get the following equation:

$$2z(a) = a.$$

Note that this expression is no longer a function of z . If I tell you a , you tell me $z(a)$ (which is $\frac{a}{2}$), and no matter what value of a I choose, when I plug $z(a)$ in on the left side of the equals, the equality relation holds. Thus

$$2z(a) = a$$

holds for any value of a . We call an expression that is true for any value of the variable (in this case a) an identity, and we write it with the fancy, three-lined equals sign in order to emphasize this.

$$2z(a) \equiv a.$$

Why should we care if something is an equality or an identity? In a nut-shell, you can differentiate both sides of an identity and the two sides remain equal. You can't do this with an equality. In fact, it doesn't even make sense to differentiate both sides of an equality. To illustrate this point, think again about the equality: $2x = 1$. What happens if you increase x by a small amount (i.e. differentiate with respect to x)? If you differentiate both sides with respect to x , you get $2 = 0$, which is not true.

On the other hand, think about $2z(a) \equiv a$. We can ask the question what happens to z if you increase a . We can answer this by differentiating both sides of the identity with respect to a . If

you do this, you get

$$2 \frac{dz(a)}{da} = 1$$

$$\frac{dz}{da} = \frac{1}{2}$$

That is, if you increase a by 1, z increases by $\frac{1}{2}$. (If you don't believe me, plug in some numbers and confirm.)

It may seem to you like I'm making a big deal out of nothing, but this is really a critical point. We are interested in determining how endogenous variables change in response to changes in exogenous variables. In this case, z is our endogenous variable and a is our exogenous variable. Thus, we are interested in things like $\frac{dz}{da}(a)$. The only way we can determine these things is to get identities that depend only on the exogenous variables and then differentiate them. Even if you don't quite believe me, you should keep this in mind. Eventually, it will become clear.

2.5.2 Back to Walras' Law: Choice Response to a Change in Wealth

As we said, Walras' Law is defined by the identity:

$$p \cdot x(p, w) \equiv w$$

or

$$\sum_{i=1}^L p_i x_i(p, w) \equiv w.$$

where the vector $x(p, w)$ describes the bundle chosen:

$$x(p, w) = (x_1(p, w), \dots, x_L(p, w))$$

Suppose we are interested in what happens to the bundle chosen if w increases a little bit. In other words, how does the bundle the consumer chooses change if the consumer's income increases by a small amount? Since we have an identity defined in terms of the exogenous variables p and w , we can differentiate both sides with respect to w :

$$\frac{d}{dw} \sum_{i=1}^L p_i x_i(p, w) \equiv \frac{d}{dw} w$$

$$\sum_{i=1}^L p_i \frac{\partial x_i}{\partial w} \equiv 1. \quad (2.1)$$

So, now we have an expression relating the changes in the amount of commodities demanded in response to a change in wealth. What does it say? The left hand side is the change in expenditure due to the increase in wealth, and the right-hand side is the increase in wealth. Thus this expression says that if wealth increases by 1 unit, total expenditure on commodities increases by 1 unit as well. Thus the latter expression just restates Walras' Law in terms of responses to changes in wealth. Any change in wealth is accompanied by an equal change in expenditure. If you think about it, this is really the only way that the consumer could satisfy Walras' Law (i.e. spend all of her money) both before and after the increase in wealth.

Based only on this expression, $\sum_i p_i \frac{\partial x_i(p,w)}{\partial w} \equiv 1$, what else can we say about the behavior of the consumer's choices in response to a change in wealth? The answer depends on what kind of commodity this is. Ordinarily, we think that if your wealth increases you will want to consume more of a good. This is certainly true of goods like trips to the movies, meals at fancy restaurants, and other "normal goods." In fact, this is so much the normal case that we just go ahead and call such goods - which have $\frac{\partial x_i(p,w)}{\partial w} > 0$ - "normal goods." But, you can also think about goods you want to consume less of as your wealth goes up - cheap cuts of meat, cross-country bus trips, nights in cheap motels, etc. All of these are things that, the richer you get, the less you want to consume them. We call goods for which $\frac{\partial x_i(p,w)}{\partial w} < 0$ "inferior goods." Since $x(p, w)$ depends on w , $\frac{\partial x_i(p,w)}{\partial w}$ depends on w as well, which means that a good may be inferior at some levels of wealth but normal at others.

So, what can we say based on $\sum_i p_i \frac{\partial x_i(p,w)}{\partial w} \equiv 1$? Well, this identity tells us that there is always at least one normal good. Why? If all goods are inferior, then the terms on the left hand side are all negative, and no matter how many negative terms you add together, they'll never sum to 1.

2.5.3 Testable Implications

We can use this observation about normal goods to derive a testable implication of our theory. Put simply, we have assumed that consumers spend all of the money they have on commodities. Based on this, we conclude that following any change in wealth, total expenditure on goods should increase by the same amount as wealth. If we knew prices and how much of the commodities the consumer buys before and after the wealth change, we could directly test this. But, suppose that we don't observe prices. However, we believe that prices do not change when wealth changes. What should we conclude if we observe that consumption of all commodities decreases following

an increase in wealth? Unfortunately, the only thing we can conclude is that our theory is wrong. People aren't spending all of their wealth on commodities 1 through L .

Based on this observation, there are a number of possible directions to go. One possible explanation is that there is another commodity, $L + 1$, that we left out of our model, and if we had accounted for that then we would see that consumption increased in response to the wealth increase and everything would be right in the world. Another possible explanation is that in the world we are considering, it is not the case that there is always something that the consumer would like more of (which, you'll recall, is the implicit assumption behind Walras' Law). This would be the case, for example, if the consumer could become satiated with the commodities, meaning that there is a level of consumption beyond which you wouldn't want to consume more even if you could. A final possibility is that there is something wrong with the data and that if consumption had been properly measured we would see that consumption of one of the commodities did, in fact, increase. In any case, the next task of the intrepid economist is to determine which possible explanation caused the failure of the theory and, if possible, develop a theory that agrees with the data.

2.5.4 Summary: How Did We Get Where We Are?

Let's review the comparative statics methodology. First, we develop an identity that expresses a relationship between the endogenous variables (consumption bundle) and the exogenous variable of interest (wealth). The identity is true for all values of the exogenous variables, so we can differentiate both sides with respect to the exogenous variables. Next, we totally differentiate the identity with respect to a particular exogenous variable of interest (wealth). By rearranging, we derive the effect of a change in wealth on the consumption bundle, and we try to say what we can about it. In the previous example, we were able to make inferences about the sign of this relationship. This is all there is to comparative statics.

2.5.5 Walras' Law: Choice Response to a Change in Price

What are other examples of comparative statics analysis? Well, in the consumer model, the endogenous variables are the amounts of the various commodities that the consumer chooses, $x_i(p, w)$. We want to know how these things change as the restrictions placed on the consumer's choices change. The restriction put on the consumer's choice by Walras' Law takes the form of the budget constraint, and the budget constraint is in turn defined by the exogenous variables — the prices of the various commodities and wealth. We already looked at the comparative statics of wealth

changes. How about the comparative statics of a price change?

Return to the Walras' Law identity:

$$\sum_{i=1}^L p_i x_i(p, w) \equiv w.$$

Since this is an identity, we can differentiate with respect to the price of one of the commodities, p_j :

$$x_j(p, w) + p_j \sum_{i=1}^L x_i \frac{\partial x_i(p, w)}{\partial p_j} = 0. \quad (2.2)$$

How does spending change in response to a price change? Well, if p_j increases, spending on good j increases, assuming that you continue to consume the same amount. This is captured by the first term in (2.2). Of course, in response to the price change, you will also rearrange the products you consume, purchasing more or less of the other products depending on whether they are gross substitutes for good j or gross complements to good j .¹⁰ The effect of this rearrangement on total expenditure is captured by the terms after the summation. Thus the meaning of (2.2) is that once you take into account the increased spending in good j and the changes in spending associated with rearranging the consumption bundle, total expenditure does not change. This is just another way of saying that the consumer's demand satisfies Walras' Law.

2.5.6 Comparative Statics in Terms of Elasticities

The goal of comparative statics analysis is to determine the change in the endogenous variable that results from a change in an exogenous variable. Sometimes it is more useful to think about the percentage change in the endogenous variable that results from a percentage change in the exogenous variable. Economists refer to the ratio of percentage changes as elasticities. Equations (2.1) and (2.2), which are somewhat difficult to interpret in their current state, become much more meaningful when written in terms of elasticities.

A price elasticity of demand gives the percentage change in quantity demanded associated with a 1% change in price. Mathematically, price elasticity elasticity is defined as:

$$\epsilon_{ip_j} = \frac{\% \Delta x_i}{\% \Delta p_j} = \frac{\partial x_i}{\partial p_j} \cdot \frac{p_j}{x_i}$$

Read ϵ_{ip_j} as "the elasticity of demand for good i with respect to the price of good j ."¹¹

¹⁰ The term 'gross' refers to the fact that wealth is held constant. It contrasts with the situation where utility is held constant, where we drop the gross. All will become clear eventually.

¹¹ Technically, the second equals sign in the equation above should be a limit, as $\% \Delta \rightarrow 0$.

Now recall equation (2.2):

$$x_j(p, w) + \sum_{i=1}^L x_i \frac{\partial x_i(p, w)}{\partial p_j} = 0.$$

The terms that are summed look almost like elasticities, except that they need to be multiplied by $\frac{p_j}{x_i}$. Perform the following sneaky trick. Multiply everything by $\frac{p_j}{w}$, and multiply each term in the summation by $\frac{x_i}{x_i}$ (we can do this because $\frac{x_i}{x_i} = 1$ as long as $x_i \neq 0$).

$$\begin{aligned} \frac{p_j x_j(p, w)}{w} + \sum_{i=1}^L \frac{p_j x_i}{w x_i} \frac{\partial x_i(p, w)}{\partial p_j} &= 0 \\ \frac{p_j x_j(p, w)}{w} + \sum_{i=1}^L \frac{p_i x_i}{w x_i} \frac{\partial x_i(p, w)}{\partial p_j} &= 0 \\ b_j(p, w) + \sum_{i=1}^L b_i(p, w) \epsilon_{ip_j} &= 0 \end{aligned} \tag{2.3}$$

where $b_j(p, w)$ is the share of total wealth the consumer spends on good j , known as the budget share.

What does (2.3) mean? Consider raising the price of good j , p_j , a little bit. If the consumer did not change the bundle she consumes, this price change would increase the consumer's total spending by the proportion of her wealth she spends on good x_j . This is known as a "wealth effect" since it is as if the consumer has become poorer, assuming she does not change behavior. The wealth effect is the first term, $b_j(p, w)$. However, if good j becomes more expensive, the consumer will choose to rearrange her consumption bundle. The effect of this rearrangement on total spending will have to do with how much is spent on each of the goods, $b_i(p, w)$, and how responsive that good is to changes in p_j , as measured by ϵ_{ip_j} . Thus the terms after the sum represent the effect of rearranging the consumption bundle on total consumption - these are known as substitution effects. Hence the meaning of (2.3) is that when you combine the wealth effect and the substitution effects, total expenditure cannot change. This, of course, is exactly what Walras' Law says.

2.5.7 Why Bother?

In the previous section, we rearranged Walras' Law by differentiating it and then manipulating the resulting equation in order to get something that means exactly the same thing as Walras' Law. Why, then, did we bother? Hopefully, seeing Walras' Law in other equations forms offers some insight into what our model predicts for consumer behavior. Furthermore, many times it is easier for economists to measure things like budget shares and elasticities than it is to measure actual

quantities and prices. In particular, budget shares and elasticities do not depend on price levels, but only on relative prices. Consequently it can be much easier to apply Walras' Law when it is written as (2.3) than when it is written as (2.2).

2.5.8 Walras' Law and Changes in Wealth: Elasticity Form

Not to belabor the point, but we can also write (2.1) in terms of elasticities, this time using the

wealth elasticity, $\epsilon_{iw} = \frac{\partial x_i}{\partial w} \cdot \frac{w}{x_i}$. Multiplying (2.1) by $\frac{x_i w}{x_i w}$ yields:

$$x_i \frac{p_i x_i w}{w x_i} \frac{\partial x_i(p, w)}{\partial w} \equiv 1 \quad (2.4)$$

$$\sum_i b_i(p, w) \epsilon_{iw} = 1.$$

The wealth elasticity ϵ_{iw} gives the percentage change in consumption of good i induced by a 1% increase in wealth. Thus, in response to an increase in wealth, total spending changes by ϵ_{iw} weighted by the budget share $b_i(p, w)$ and summed over all goods. In other words, if wealth increases by 1, total expenditure must also increase by 1. Thus, equation (2.4) is yet another statement of the fact that the consumer always spends all of her money.

2.6 Requirement 2 Revisited: Demand is Homogeneous of Degree Zero.

The second requirement for consumer choices is that "only real opportunities matter." In mathematical terms this means that "demand is homogeneous of degree zero," or:

$$x(\alpha p, \alpha w) \equiv x(p, w)$$

Note that this is an identity. Thus it holds for any values of p and w . In words what it says is that if the consumer chooses bundle $x(p, w)$ when prices are p and income is w , and you multiply all prices and income by a factor, $\alpha > 0$, the consumer will choose the same bundle after the multiplication as before, $x(\alpha p, \alpha w) = x(p, w)$. The reason for this is straightforward. If you multiply all prices and income by the same factor, the budget set is unchanged. $B_{p,w} = \{x : p \cdot x \leq w\} = \{x : \alpha p \cdot x \leq \alpha w\} = B_{\alpha p, \alpha w}$. And, since the set of bundles that the consumer could choose is not changed, the consumer should choose the same bundle.

There are two important points that come out of this:

1. This is an expression of the belief that changes in behavior should come from changes in the set of available alternatives. Since the rescaling of prices and income do not affect the budget set, they should not affect the consumer's choice.

2. The second point is that nominal prices are meaningless in consumer theory. If you tell me that a loaf of bread costs \$10, I need to know what other goods cost before I can interpret the first statement. And, in terms of analysis, this means that we can always "normalize" prices by arbitrarily setting one of them to whatever we like (often it is easiest to set it equal to 1), since only the real prices matter and fixing one commodity's nominal price will not affect the relative values of the other prices.

Exercise 3 If you don't believe me that this change doesn't affect the budget set, you should go back to the two-commodity example, plug in the numbers and check it for yourself. If you can't do it with the general scaling factor α , you should let $\alpha = 2$ and try it for that. Most of the time, things that are hard to understand with general parameter values like α, p, w are simple once you plug in actual numbers for them and churn through the algebra.

2.6.1 Comparative Statics of Homogeneity of Degree Zero

We can also perform a comparative statics analysis of the requirement that demand be homogeneous of degree zero, i.e. only real opportunities matter. What does this imply for choice behavior?

The homogeneity assumption applies to proportional changes in all prices and wealth:

$$x_i(\alpha p, \alpha w) \equiv x_i(p, w) \text{ for all } i, \alpha > 0.$$

To make things clear, let the initial price vector be denoted $p^0 = p_1^0, \dots, p_L^0$ and let w^0 original wealth, and (for the time being) assume that $L = 2$. For example, $p^0 = (1, 7)$ and $w^0 = 7$. Before we differentiate, I want to make sure that we're clear on what is going on. So, rewrite the above expression as:

$$x_i(\alpha p_1^0, \alpha p_2^0, \alpha w^0) \equiv x_i(p_1^0, p_2^0, w^0) \tag{2.5}$$

Now, notice that on the left-hand side for any $\alpha > 0$ the price of good 1 is $p_1 = \alpha p_1^0$, the price of good 2 is $p_2 = \alpha p_2^0$, and wealth is $w = \alpha w^0$. That is, given α . We are interested in what happens to demand as α changes, so it is important to recognize that the prices and wealth are functions of α .

We are interested in what happens to demand when, beginning at original prices p^0 and wealth w^0 , we scale up all prices and wealth proportionately. To do this, we want to see what happens when we increase α , starting at $\alpha = 1$. Because the prices and wealth are functions of α , we have to use the Chain Rule in evaluating the derivative of (2.5) with respect to α . Differentiating (2.5) with respect to α yields:

$$\frac{\partial x_i}{\partial p_1} \frac{\partial p_1}{\partial \alpha} + \frac{\partial x_i}{\partial p_2} \frac{\partial p_2}{\partial \alpha} + \frac{\partial x_i}{\partial w} \frac{\partial w}{\partial \alpha} \equiv 0.$$

Since $p_1 = \alpha p_1^0$, $\frac{\partial p_1}{\partial \alpha} = p_1^0$, and similarly $\frac{\partial p_2}{\partial \alpha} = p_2^0$, and $\frac{\partial w}{\partial \alpha} = w^0$, so this expression becomes:

$$\frac{\partial x_i}{\partial p_1} p_1^0 + \frac{\partial x_i}{\partial p_2} p_2^0 + \frac{\partial x_i}{\partial w} w^0 \equiv 0. \tag{2.6}$$

Notice that the first line takes the standard Chain Rule form: for each argument (p_1 , p_2 , and w), take the partial derivative of the function with respect to that argument and multiply it by the derivative with respect to α of “what’s inside” the argument.¹²

Finally, notice that (2.6) has prices and wealth p_1^0, p_2^0, w^0 . We are asking the question “what happens to x_i when prices and wealth begin at p_1^0, p_2^0, w^0 and are all increased slightly by the same proportion?” In order to make sure we are answering this question, we need to set $\alpha = 1$, so that the partial derivatives are evaluated at the original prices and wealth. Evaluating the last expression at $\alpha = 1$ yields the following expression in terms of the original price-wealth vector (s_1, s_2, v) :

$$\frac{\partial x_i}{\partial p_1} p_1^0 + \frac{\partial x_i}{\partial p_2} p_2^0 + \frac{\partial x_i}{\partial w} w^0 \equiv 0. \tag{2.7}$$

Generalizing the previous argument to the case where L is any positive number, expression (2.7) becomes:

$$\frac{\partial x_i}{\partial w} w^0 + \sum_{j=1}^L \frac{\partial x_i}{\partial p_j} p_j^0 = 0 \text{ for all } i. \tag{2.8}$$

This is where we need to face an ugly fact. Economists are terrible about notation, which makes this stuff harder to learn than it needs to be. When you see (2.8) written in a textbook, it will look like this:

$$\frac{\partial x_i}{\partial w}(p, w) + \sum_{j=1}^L \frac{\partial x_i}{\partial p_j}(p, w) p_j = 0 \text{ for all } i.$$

¹² If you are confused, see the next subsection for further explanation.

That is, they drop the superscript “0” that denotes the original price vector. But, notice that the symbol “w” in this expression has two different meanings. The “w” in “ ∂w ” in the denominator of the first term says “we’re differentiating with respect to the wealth argument,” while the “w” in “ $\partial x_i(p, w)$ ” and the “w” multiplying this term refer to the original wealth level, i.e., the wealth level at which the expression is being evaluated. Similarly, “ p_j ” also has two different meanings in this expression. To make things worse, economists frequently skip steps in derivations.¹³

It is straightforward to get an elasticity version of (2.8). Just divide through by $x_i(p, w)$:

$$\varepsilon_{iw} + \sum_{j=1}^L \varepsilon_{ip_j} = 0. \quad (2.9)$$

Elasticities ε_{iw} and ε_{ip_j} give the elasticity of the consumer’s demand response to changes in wealth and the price of good j , respectively. The total percentage change in consumption of good i is given by summing the percentage changes due to changes in wealth and in each of the prices. Homogeneity of degree zero says that in response to proportional changes in all prices and wealth the total change in demand for each commodity should not change. This is exactly what (2.9) says.

2.6.2 A Mathematical Aside ...

If this is unfamiliar to you, the computation may seem strange. If it doesn’t seem strange, then skip on to the next section.

If you’re still here, let’s try it one more time. This time, we’ll let $L = 2$, and choose specific values for the prices and wealth. Let good 1’s price be 5, good 2’s price be 3, and wealth be 10 initially. Then, (2.5) writes as:

$$x_i(5\alpha, 3\alpha, 10\alpha) \equiv x_i(5, 3, 10).$$

Now, starting at prices (5, 3) and wealth 10, we are interested in what happens to demand for x_i as we increase all prices and wealth proportionately. To do this, we will first increase α by a small amount (i.e., differentiate with respect to α), and then we’ll evaluate the resulting expression at $\alpha = 1$. This will give us an expression for the effect of a small increase in α . So, totally

differentiate both sides with respect to α :

$$\frac{\partial x_i(5\alpha, 3\alpha, 10\alpha)}{\partial p_1} \frac{d(5\alpha)}{d\alpha} + \frac{\partial x_i(5\alpha, 3\alpha, 10\alpha)}{\partial p_2} \frac{d(3\alpha)}{d\alpha} + \frac{\partial x_i(5\alpha, 3\alpha, 10\alpha)}{\partial w} \frac{d(10\alpha)}{d\alpha} \equiv 0$$

$$\frac{\partial x_i}{\partial p_1} 5 + \frac{\partial x_i}{\partial p_2} 3 + \frac{\partial x_i}{\partial w} 10 \equiv 0.$$

¹³ These are a couple of the main reasons why documents such as these are needed.

Again, the partial derivatives $\frac{\partial x_i}{\partial p_j}$ denote the partial derivative of function $x_i(p, w)$ with respect to the “ p_j slot,” i.e., the j^{th} argument of the function. And, since we are interested in what happens when you increase all prices and wealth proportionately beginning from prices (5, 3) and wealth 10, we would like the left-hand side to be evaluated at (5, 3, 10). To get this, set $\alpha = 1$:

$$\frac{\partial x_i(5, 3, 10)}{\partial p_1} \cdot 5 + \frac{\partial x_i(5, 3, 10)}{\partial p_2} \cdot 3 + \frac{\partial x_i(5, 3, 10)}{\partial w} \cdot 10 \equiv 0 \quad (2.10)$$

Comparing this expression with (2.7) shows that the role of p_1^* , p_2^* , and w^* are played by 5, 3, and 10, respectively in (2.10), as is expected.

The source of confusion in understanding this derivation seems to lie in confusing the partial derivative of x_i with respect to the p_1 argument (for example) with the particular price of good 1, which is 5α in this example and αp_1^* in the more general derivation above. The key is to notice that, in applying the chain rule, you always differentiate the function (e.g., $x_i()$) with respect to its argument (e.g., p_1), and then differentiate the function that is in the argument’s “slot” (e.g., 5α or αp_1 or αp_1 if you are an economist) with respect to α .

2.7 Requirement 3 Revisited: The Weak Axiom of Revealed Preference

The third requirement that we will place on consumer choices is that they satisfy the Weak Axiom of Revealed Preference (WARP). To remind you of the informal definition, WARP is a requirement of consistency in decision-making. It says that if a consumer chooses z when y was also affordable, this choice reveals that the consumer prefers z to y . Since we assume that consumer preferences are constant and we have modeled all of the relevant constraints on consumer behavior and preferences, if we ever observe the consumer choose y , it must be that z was not available (since if it were, the consumer would have chosen z over y since she had previously revealed her preference for z). We now turn to the formal definition.

Definition 4 Consider any two distinct price-wealth vectors (p, w) and $(p^0, w^0) \neq (p, w)$. Let $z = x(p, w)$ and $y = x(p^0, w^0)$. The consumer’s demand function satisfies WARP if whenever $p \cdot y \leq w$, $p^0 \cdot z > w^0$.

We can restate the last part of the definition as: if $y \in B_{p,w}$, then $z \notin B_{p^0, w^0}$. If y could have been chosen when z was chosen, then the consumer has revealed that she prefers z to y . Therefore

if you observe her choose y , it must be that z was not available. I apologize for repeating the same definition over and over, but a) it helps to attach words to the math, and b) if you wanted math without explanation you could read a textbook.

In its basic form, WARP does not generate any predictions that can immediately be taken to the data and tested. But, if we rearrange the statement a little bit, we can get an easily testable prediction. So, let me ask the WARP question a different way. Suppose the consumer chooses z when prices and wealth are (p, w) , and z is affordable when prices and wealth are (p^0, w^0) . What does WARP tell us about which bundles the consumer could choose when prices are (p^0, w^0) ?

There are two choices to consider: either $x(p^0, w^0) = z$. This is perfectly admissible under WARP. The other choice is that $x(p^0, w^0) = y \neq z$. In this case, WARP will place restrictions on which bundles y can be chosen. What are they? By virtue of the fact that z was chosen when prices and wealth were (p, w) , we know that $y \notin B_{p,w}$, since if it were there would be a violation of WARP. Thus it must be that if the consumer chooses a bundle y different than x at (p^0, w^0) , y must not have been affordable when prices and wealth were (p, w) .

This is illustrated graphically in figure 2.F.1 in MWG (p. 30). In panel a, since $x(p^0, w^0)$ is chosen at (p^0, w^0) , when prices are (p^{00}, w^{00}) the consumer must either choose $x(p^0, w^0)$ again or a bundle $x(p^{00}, w^{00})$ that is not in B_{p^0, w^0} . If we assume that demand satisfies Walras' Law as well, $x(p^{00}, w^{00})$ must lay on the frontier. Thus if $x(p^0, w^0)$ is as drawn, it cannot be chosen at prices (p^{00}, w^{00}) . The chosen bundle must lay on the segment of $B_{p^{00}, w^{00}}$ below and to the right of the intersection of the two budget lines, as does $x(p^{00}, w^{00})$. Similar reasoning holds in panel b. The chosen bundle cannot lay within B_{p^0, w^0} if WARP holds. Panel c depicts the case where $x(p^0, w^0)$ is affordable both before and after the change in prices and wealth. In this case, $x(p^0, w^0)$ could have been chosen after the price change. But, if it is not chosen at (p^{00}, w^{00}) , then the chosen bundle must lay outside of $B_{p^{00}, w^{00}}$, as does $x(p^{00}, w^{00})$. In panels d and e, $x(p^{00}, w^{00}) \in B_{p^0, w^0}$, and thus this behavior does not satisfy WARP.

2.7.1 Compensated Changes and the Slutsky Equation

Panel c in MWG Figure 2.F.1 suggests a way in which WARP can be used to generate predictions about behavior. Imagine two different price-wealth vectors, (p, w) and (p^0, w^0) , such that bundle $z = x(p, w)$ lies on the frontier of both $B_{p,w}$ and B_{p^0, w^0} . This corresponds to the following hypothetical situation. Suppose that originally prices are (p, w) and you choose bundle $z = x(p, w)$. I tell you that I am going to change the price vector to p^0 . But, I am fair, and so I tell you that in

order to make sure that you are not made worse off by the price change, I am also going to change your wealth to w^0 , where w^0 is chosen so that you can still just afford bundle z at the new prices and wealth (p^0, w^0) . Thus $w^0 = p^0 \cdot z$. We call this a compensated change in price, since I change your wealth to compensate you for the effects of the price change.

Since you can afford z before and after the price change, we know that:

$$p \cdot z = w \text{ and } p^0 \cdot z = w^0.$$

Let $y = x(p^0, w^0)$ be the bundle chosen at (p^0, w^0) . Since you actually choose y at price-wealth (p^0, w^0) , assuming your demand satisfies Walras Law we know that $p^0 \cdot y = w^0$ as well. Thus

$$0 = w^0 - w^0 = p^0 \cdot y - p^0 \cdot z$$

so, $p^0 \cdot (y - z) = 0$.

Further, since z is affordable at (p^0, w^0) , by WARP it must be that y was not affordable at (p, w) :

$$p \cdot y > w$$

$$p \cdot y - p \cdot z > 0$$

$$p \cdot (y - z) > 0.$$

Finally, subtracting $p \cdot (y - z) > 0$ from $p^0 \cdot (y - z) = 0$ yields:

$$p^0 \cdot (y - z) - p \cdot (y - z) < 0 \tag{2.11}$$

Equation (2.11) captures the idea that, following a compensated price change, prices and demand move in opposite directions. Although this takes a little latitude since prices and bundles are vectors, you can interpret (2.11) as saying that if prices increase, demand decreases.¹⁴ To put it another way, let $\Delta p = p^0 - p$ denote the vector of price changes and $\Delta x = x(p^0, w^0) - x(p, w)$ denote the vector of quantity changes. (2.11) can be rewritten as

$$\Delta p \cdot \Delta x^c \leq 0$$

where we have replaced the strict inequality with a weak inequality in recognition that it may be the case that $y = z$. Note that the superscript c on Δx^c is to remind us that this is the compensated change in x . This is a statement of the Compensated Law of Demand (CLD): If the price of

¹⁴ This is especially true in the case where p and p^0 differ only in the price of good j , which changes by an amount dp_j . In this case, $p^0 - p = (0, 0, \dots, dp_j, 0, \dots, 0)$, and $(p^0 - p) \cdot (y - z^0) = dp_j dx_j$.

a commodity goes up, you demand less of it. If we take a calculus view of things, we can rewrite this in terms of differentials: $dp \cdot dx^C \leq 0$.

We're almost there. Now, what does it mean to give the consumer a compensated price change? Let \hat{x} be the initial consumption bundle, i.e., $\hat{x} = x(p, w)$, where p and w are the original prices and wealth. A compensated price change means that at any price, p , bundle \hat{x} is still affordable. Hence, after the price change, wealth is changed to $\hat{w} = p \cdot \hat{x}$. Note that the \hat{x} in this expression is the original consumption bundle, not the choice function $x(p, w)$. Consider the consumer's demand for good i

$$x_i^C = x_i(p, p \cdot \hat{x})$$

following a compensated change in the price of good j :

$$\frac{d}{dp_j} (x_i(p, p \cdot \hat{x})) = \frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} \frac{\partial (p \cdot \hat{x})}{\partial p_j}$$

$$\frac{dx_i^C}{dp_j} = \frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} \hat{x}_j$$

Since $\hat{x}_j = x_j(p, w)$, we'll just drop the "hat" from now on. If we write the previous equation as a differential, this is simply:

$$dx_i^C = \frac{\partial x_i}{\partial p_j} dp_j + \frac{\partial x_i}{\partial w} x_j dp_j = s_{ij} dp_j$$

where $s_{ij} = \frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} x_j$. If we change more than one p_j , the change in x_i^C would simply be the sum of the changes due to the different price changes:

$$dx_i^C = \sum_{j=1}^L \left(\frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} x_j \right) dp_j = s_i \cdot dp$$

where $s_i = (s_{i1}, \dots, s_{ij}, \dots, s_{iL})$ and $dp = (dp_1, \dots, dp_L)$ is the vector of price changes. Finally, we can arrange the dx_i^C into a vector by stacking these equations vertically. This gives us:

$$dx^C = Sdp$$

where S is an $L \times L$ matrix with the element in the i th the row and j th column being s_{ij} .

Now, return to the statement of WARP:

$$dp \cdot dx^C \leq 0$$

Substituting in $dx^C = Sdp$ yields

$$dp \cdot Sdp \leq 0. \tag{2.12}$$

Inequality (2.12) has a mathematical significance: It implies that matrix S , which we will call the substitution matrix, is negative semi-definite. What this means is that if you pre- and post-multiply S by the same vector, the result is always a non-positive number. This is important because mathematicians have figured out a bunch of nice properties of negative semi-definite matrices. Among them are:

1. The principal-minor determinants of S follow a known pattern.
2. The diagonal elements s_{ij} are non-positive. (Generally, they will be strictly negative, but we can't show that based on what we've done so far).
3. Note that WARP does not imply that S is symmetric. This is the chief difference between the choice-based approach and the preference-based approach we will consider later.

All I want to say about #1 is this. Basically, it amounts to knowing that the second-order conditions for a certain maximization problem are satisfied. But, in this course we aren't going to worry about second-order conditions. So, file it away that if you ever need to know anything about the principal minors of S , you can look it up in a book.

Item #2 is a fundamental result in economics, because it says that the change in demand for a good in response to a compensated price increase is negative. In other words, if price goes up, demand goes down. This is the Compensated Law of Demand (CLD). You may be thinking that it was a lot of work to derive something so obvious, but the fact that the CLD is derived from WARP and Walras' Law is actually quite important. If these were not sufficient for the CLD, which we know from observation to be true, then that would be a strong indicator that we have left something out of our model.

The fact that $s_{ij} \leq 0$ can be used to help explain an anomaly of economic theory, the Giffen good. Ordinarily, we think that if the price of a good increases, holding wealth constant, the demand for that good will decrease. This is probably what you thought of as the "Law of Demand," even though it isn't always true. Theoretically, it is possible that when the price of a good increases, a consumer actually chooses to consume more of it. By way of motivation, think of the following story. A consumer spends all of her money on two things: food and trips to Hawaii. Suppose the price of food increases. It may be that after the increase, the consumer can no longer afford the trip to Hawaii and therefore spends all of her money on food. The result is that the consumer actually buys more food than she did before the price increase.

How does this story manifest itself in the theory we have learned up until now? We know that:

$$s_{ii} = \frac{\partial x_i}{\partial p_i} + \frac{\partial x_i}{\partial w} X_i$$

Rearranging it:

$$\frac{\partial x_i(p, w)}{\partial p_i} = s_{ii} - \frac{\partial x_i}{\partial w} X_i$$

We know that $s_{ii} \leq 0$ since S is negative semi-definite. Clearly, $x_i \geq 0$. But, what happens if x_i

is a strongly inferior good? In this case, $\frac{\partial x_i}{\partial w} < 0$, meaning $-\frac{\partial x_i}{\partial w} X_i > 0$. And, if the magnitude of $-\frac{\partial x_i}{\partial w} X_i$ is greater than s_{ii} , it can be that $\frac{\partial x_i}{\partial p_i} > 0$, which is what it means to be a Giffen good.

What does the theory tell us? Well, it tells us that in order for a good to be a Giffen good, it must be a strongly inferior good. Or, to put it the other way, a normal good cannot be a Giffen good.¹⁵

Before going on, let me give one more aside on why we bother with all of this stuff. Remember when I started talking about increasing prices, and I said that I was fair, so I was going to also change your wealth? Well, it turns out that a good measure of the impact of a price change on a consumer is given by the change in wealth it would take to compensate you for a price change. So, if we could observe the s_{ij} terms, this would help us to measure the impact of price changes on consumers. But, the problem is that we never observe compensated price changes, we only observe the uncompensated ones, $\frac{\partial x_i(p, w)}{\partial p_i}$. But, the relationship above gives us a way to recover s_{ii} from

observations on uncompensated price changes $\frac{\partial x_i}{\partial p_i}$, wealth changes, $\frac{\partial x_i}{\partial w}$, and actual consumption, x_i . Thus the importance of the relationship $s_{ii} = \frac{\partial x_i}{\partial p_i} + \frac{\partial x_i}{\partial w} X_i$ is that it allows us to recover an unobservable quantity that we are interested in, s_{ii} , from observables.

2.7.2 Other Properties of the Substitution Matrix

Based on what we know about demand functions, we can also determine a couple of additional properties of the Substitution matrix. They are:

$$\begin{aligned} p \cdot S(p, w) &= 0 \\ S(p, w) p &= 0 \end{aligned}$$

¹⁵ Technically, it is not the goods that are Giffen. Rather, the consumer's behavior at a particular price-wealth combination is Giffen. For example, it has been shown that very poor consumers in China exhibit Giffen behavior: their demand curve for rice slopes upward in the price of rice. But, non-poor consumers do not exhibit Giffen behavior: their demand curve slopes downward. See R. Jensen and N. Miller (2001).

These can be derived from the comparative statics implications of Walras' Law and homogeneity of degree zero. Their effect is to impose additional restrictions on the set of admissible demand

functions. So, suppose you get some estimates of $\frac{\partial x_i}{\partial p_j}$, p , w , and $\frac{\partial x_i}{\partial w}$, which can all be computed

from data, and you are concerned with whether you have a good model. One thing you can do is compute S from the data, and check to see if the two equations above hold. If they do, you're doing okay. If they don't, this is a sign that your data do not match up with your theory. This could be due to data problems or to theory problems, but in either case it means that you have work to do.¹⁶

¹⁶ The usual statistical procedure in this instance is to impose these conditions as restrictions on the econometric model and then test to see if they are valid. I leave it to people who know more econometrics than I do to explain how.

Chapter 3

The Traditional Approach to Consumer Theory

In the previous section, we considered consumer behavior from a choice-based point of view. That is, we assumed that consumers made choices about which consumption bundle to choose from a set of feasible alternatives, and, using some rather mild restrictions on choices (homogeneity of degree zero, Walras' law, and WARP), were able to make predictions about consumer behavior. Notice that our predictions were entirely based on consumer behavior. In particular, we never said anything about why consumers behave the way they do. We only hold that the way they behave should be consistent in certain ways.

The traditional approach to consumer behavior is to assume that the consumer has well-defined preferences over all of the alternative bundles and that the consumer attempts to select the most preferred bundle from among those bundles that are available. The nice thing about this approach is that it allows us to build into our model of consumer behavior how the consumer feels about trading off one commodity against another. Because of this, we are able to make more precise predictions about behavior. However, at some point people started to wonder whether the predictions derived from the preference-based model were in keeping with the idea that consumers make consistent choices, or whether there could be consistent choice-based behavior that was not derived from the maximization of well-defined preferences. It turns out that if we define consistent choice making as homogeneity of degree zero, Walras' law, and WARP, then there are consistent choices that cannot be derived from the preference-based model. But, if we replace WARP with a slightly stronger but still reasonable condition, called the Strong Axiom of Revealed Preference (SARP),

then any behavior consistent with these principles can be derived from the maximization of rational preferences.

Next, we take up the traditional approach to consumer theory, often called “neoclassical” consumer theory.

3.1 Basics of Preference Relations

We’ll continue to assume that the consumer chooses from among L commodities and that the commodity space is given by $X \subset \mathbb{R}_+^L$. The basic idea of the preference approach is that given any two bundles, we can say whether the first is “at least as good as” the second. The “at-least-as-good-as” relation is denoted by the curvy greater-than-or-equal-to sign: \succsim . So, if we write $x \succsim y$, that means that “ x is at least as good as y .”

Using \succsim , we can also derive some other preference relations. For example, if $x \succsim y$, we could also write $y \preceq x$, where \preceq is the “no better than” relation. If $x \succsim y$ and $y \succsim x$, we say that a consumer is “indifferent between x and y ,” or symbolically, that $x \sim y$. The indifference relation is important in economics, since frequently we will be concerned with indifference sets. The indifference curve I_y is defined as the set of all bundles that are indifferent to y . That is, $I_y = \{x \in X \mid y \sim x\}$. Indifference sets will be very important as we move forward, and we will spend a great deal of time and effort trying to figure out what they look like, since the indifference sets capture the trade-offs the consumer is willing to make among the various commodities. The final preference relation we will use is the “strictly better than” relation. If x is at least as good as y and y is not at least as good as x , i.e., $x \succsim y$ and not $y \succsim x$ (which we could write $y \prec x$), we say that $x \hat{\succ} y$, or x is strictly better than (or strictly preferred to) y .

Our preference relations are all examples of mathematical objects called binary relations. A binary relation compares two objects, in this case, two bundles. For instance, another binary relation is “less-than-or-equal-to,” \leq . There are all sorts of properties that binary relations can have. The first two we will be interested in are called completeness and transitivity. A binary relation is complete if, for any two elements x and y in X , either $x \succsim y$ or $y \succsim x$. That is, any two elements can be compared. A binary relation is transitive if $x \succsim y$ and $y \succsim z$ imply $x \succsim z$. That is, if x is at least as good as y , and y is at least as good as z , then x must be at least as good as z .

The requirements of completeness and transitivity seem like basic properties that we would like any person’s preferences to obey. This is true. In fact, they are so basic that they form economists’

very definition of what it means to be rational. That is, a preference relation \succsim is called rational if it is complete and transitive.

When we talked about the choice-based approach, we said that there was implicit in the idea that demand functions satisfy Walras Law the idea that “more is better.” This idea is formalized in terms of preferences by making assumptions about preferences over one bundle or another. Consider the following property, called monotonicity:

Definition 5 A preference relation \succsim is monotone if $x \succsim y$ for any x and y such that $x_l > y_l$ for $l = 1, \dots, L$. It is strongly monotone if $x_l \geq y_l$ for all $l = 1, \dots, L$ and $x_j > y_j$ for some $j \in \{1, \dots, L\}$ implies that $x \succ y$.

Monotonicity and strong monotonicity capture two different notions of “more is better.” Monotonicity says that if every component of x is larger than the corresponding component of y , then x is strictly preferred to y . Strong monotonicity is the requirement that if every component of x is at least as large (but not necessarily strictly larger) than the corresponding component of y and at least one component of x is strictly larger, x is strictly preferred to y .

The difference between monotonicity and strong monotonicity is illustrated by the following example. Consider the bundles $x = (1, 1)$ and $y = (1, 2)$. If \succsim is strongly monotone, then we can say that $y \succ x$. However, if \succsim is monotone but not strongly monotone, then it need not be the case that y is strictly preferred to x . Since preference relations that are strongly monotone are monotone, but preferences that are monotone are not necessarily strongly monotone, strong monotonicity is a more restrictive (a.k.a. “stronger”) assumption on preferences.

If preferences are monotone or strongly monotone, it follows immediately that a consumer will choose a bundle on the boundary of the Walrasian budget set. Hence an assumption of some sort of monotonicity must have been in the background when we assumed consumer choices obeyed Walras’ Law. However, choice behavior would satisfy Walras’ Law even if preferences satisfied the following weaker condition, called local nonsatiation.

Condition 6 A preference relation \succsim satisfies local nonsatiation if for every x and every $\varepsilon > 0$ there is a point y such that $\|x - y\| \leq \varepsilon$ and $y \succ x$.

That is, for every x , there is always a point “nearby” that the consumer strictly prefers to x , and this is true no matter how small you make the definition of “nearby.” Local nonsatiation allows for the fact that some commodities may be “bads” in the sense that the consumer would sometimes

prefer less of them (like garbage or noise). However, it is not possible for all goods to always be bads if preferences are non-satiated. (Why?)

It's time for a brief discussion about the practice of economic theory. Recall that the object of doing economic theory is to derive testable implications about how real people will behave. But, as we noted earlier, in order to derive testable implications, it is necessary to impose some restrictions on (make assumptions about) the type of behavior we allow. For example, suppose we are interested in the way people react to wealth changes. We could simply assume that people's behavior satisfies Walras' Law, as we did earlier. This allows us to derive testable implications. However, it provides little insight into why they satisfy Walras' Law. Another option would be to assume monotonicity — that people prefer more to less. Monotonicity implies that people will satisfy Walras' Law. But, it rules out certain types of behavior. In particular, it rules out the situation where people prefer less of an object to more of it. But, introspection tells us that sometimes we do prefer less of something. So, we ask ourselves if there is a weaker assumption that allows people to prefer less to more, at least sometimes, that still implies Walras' Law. It turns out that local nonsatiation is just such an assumption. It allows for people to prefer less to more — even to prefer less of everything — the only requirement is that, no matter which bundle the consumer currently selects, there is always a feasible bundle nearby that she would rather have.

By selecting the weakest assumption that leads to a particular result, we accomplish two tasks. First, the weaker the assumptions used to derive a result, the more “robust” it is, in the sense that a greater variety of initial conditions all lead to the same conclusion. Second, finding the weakest possible condition that leads to a particular conclusion isolates just what is needed to bring about the conclusion. So, all that is really needed for consumers to satisfy Walras' Law is for preferences to be locally nonsatiated — but not necessarily monotonic or strongly monotonic.

The assumptions of monotonicity or local nonsatiation will have important implications for the way indifference sets look. In particular, they ensure that $I_x = \{y \in X \mid y \sim x\}$ are downward sloping and “thin.” That is, they must look like Figure 3.1.

If the indifference curves were thick, as in Figure 3.2, then there would be points such as x , where in a neighborhood of x (the dotted circle) all points are indifferent to x . Since there is no strictly preferred point in this region, it is a violation of local-nonsatiation (or monotonicity).

In addition to the indifference set I_x defined earlier, we can also define upper-level sets and lower-level sets. The upper level set of x is the set of all points that are at least as good as x , $U_x = \{y \in X \mid y \succeq x\}$. Similarly, the lower level set of x is the set of all points that are no




Figure 3.1: Thin Indifferent Sets




Figure 3.2: Thick Indifference Sets

better than x , $L_x = \{y \in X | x \succ y\}$. Just as monotonicity told us something about the shape of indifference sets, we can also make assumptions that tell us about the shape of upper and lower level sets.

Recall that a set of points, X , is convex if for any two points in the set the (straight) line segment between them is also in the set.¹ Formally, a set X is convex if for any points x and x^0 in X , every point z on the line joining them, $z = tx + (1 - t)x^0$ for some $t \in [0, 1]$, is also in X . Basically, a convex set is a set of points with no holes in it and with no “notches” in the boundary. You should draw some pictures to figure out what I mean by no holes and no notches in the set.

Before we move on, let’s do a thought experiment. Consider two possible commodity bundles, x and x^0 . Relative to the extreme bundles x and x^0 , how do you think a typical consumer feels about an average bundle, $z = tx + (1 - t)x^0$, $t \in (0, 1)$? Although not always true, in general, people tend to prefer bundles with medium amounts of many goods to bundles with a lot of some things and very little of others. Since real people tend to behave this way, and we are interested in modeling how real people behave, we often want to impose this idea on our model of preferences.²

Exercise 7 Confirm the following two statements: (1) If \succ is convex, then if $y \succ x$ and $z \succ x$, $ty + (1 - t)z \succ x$ as well. (2) Suppose $x \sim y$. If \succ is convex, then for any $z = ty + (1 - t)x$, $z \succ x$.

Another way to interpret convexity of preferences is in terms of a diminishing marginal rate of substitution (MRS), which is simply the slope of the indifference curve. The idea here is that if you are currently consuming a bundle x , and I offer to take some x_1 away from you and replace it with some x_2 , I will have to give you a certain amount of x_2 to make you exactly indifferent for the loss of x_1 . A diminishing MRS means that this amount of x_2 I have to give you increases the more x_2 that you already are consuming - additional units of x_2 aren’t as valuable to you.

The upshot of the convexity and local non-satiation assumptions is that indifference sets have to be thin, downward sloping, and be “bowed upward.” There is nothing in the definition of convexity

¹This is the definition of a convex set. It should not be confused with a convex function, which is a different thing altogether. In addition, there is such thing as a concave function. But, there is no such thing as a concave set. I sympathize with the fact that this terminology can be confusing. But, that’s just the way it is.

My advice is to focus on the meaning of the concepts, i.e., “a set with no notches and no holes.”

²It is only partly true that when we assume preferences are convex we do so in order to capture real behavior. In addition, the basic mathematical techniques we use to solve our problems often depend on preferences being convex. If they are not (and one can readily think of examples where preferences are not convex), other, more complicated techniques have to be used.

that prevents flat regions from appearing on indifference curves. However, there are reasons why we want to rule out indifference curves with flat regions. Because of this, we strengthen the convexity assumption with the concept of strict convexity. A preference relation is strictly convex if for any distinct bundles y and z ($y \neq z$) such that $y \succsim x$ and $z \succsim x$, $ty + (1 - t)z \succ x$. Thus imposing strict convexity on preferences strengthens the requirement of convexity (which actually means that averages are at least as good as extremes) to say that averages are strictly better than extremes.

3.2 From Preferences to Utility

In the last section, we said a lot about preferences. Unfortunately, all of that stuff is not very useful in analyzing consumer behavior, unless you want to do it one bundle at a time. However, if we could somehow describe preferences using mathematical formulas, we could use math techniques to analyze preferences, and, by extension, consumer behavior. The tool we will use to do this is called a utility function.

A utility function is a function $U(x)$ that assigns a number to every consumption bundle $x \in X$. Utility function $U(\cdot)$ represents preference relation \succsim if for any x and y , $U(x) \geq U(y)$ if and only if $x \succsim y$. That is, function U assigns a number to x that is at least as large as the number it assigns to y if and only if x is at least as good as y . The nice thing about utility functions is that if you know the utility function that represents a consumer's preferences, you can analyze these preferences by deriving properties of the utility function. And, since math is basically designed to derive properties of functions, it can help us say a lot about preferences.

Consider a typical indifference curve map, and assume that preferences are rational. We also need to make a technical assumption, that preferences are continuous. For our purposes, it isn't worth derailing things in order to explain why this is necessary. But, you should look at the example of lexicographic preferences in MWG to see why the assumption is necessary and what can go wrong if it is not satisfied.

The line drawn in Figure 3.3 is the line $x_2 = x_1$, but any straight line would do as well. Notice that we could identify the indifference curve I_x by the distance along the line $x_2 = x_1$ you have to travel before intersecting I_x . Since indifference curves are downward sloping, each I_x will only intersect this line once, so each indifference curve will have a unique number associated with it. Further, since preferences are convex, if $x \succ y$, I_x will lay above and to the right of I_y (i.e. inside I_y), and so I_x will have a higher number associated with it than I_y .




Figure 3.3: Ranking Indifference Curves

We will call the number associated with I_x the utility of x . Formally, we can define a function $u(x_1, x_2)$ such that $u(x_1, x_2)$ is the number associated with the indifference curve on which (x_1, x_2) lies. It turns out that in order to ensure that there is a utility function corresponding to a particular preference relation, you need to assume that preferences are rational and continuous. In fact, this is enough to guarantee that the utility function is a continuous function. The assumption that preferences are rational agrees with how we think consumers should behave, so it is no problem. The assumption that preferences are continuous is what we like to call a technical assumption, by which we mean that it is needed for the arguments to be mathematically rigorous (read: true), but it imposes no real restrictions on consumer behavior. Indeed, the problems associated with preferences that are not continuous arise only if we assume that all commodities are infinitely divisible (or come in infinite quantities). Since neither of these is true of real commodities, we do not really harm our model by assuming continuous preferences.

3.2.1 Utility is an Ordinal Concept

Notice that the numbers assigned to the indifference curves in defining the utility function were essentially arbitrary. Any assignment of numbers would do, as long as the order of the numbers assigned to various bundles is not disturbed. Thus if we were to multiply all of the numbers by 2, or add 6 to them, or take the square root, the numbers assigned to the indifference curves after the transformation would still represent the same preferences. Since the crucial characteristic of a utility function is the order of the numbers assigned to various bundles, but not the bundles themselves, we say that utility is an ordinal concept.

This has a number of important implications for demand analysis. The first is that if $U(x)$ represents \succsim and $f(\cdot)$ is a monotonically increasing function (meaning the function is always increasing as its argument increases), then $V(x) = f(U(x))$ also represents \succsim . This is very valuable for the following reason. Consider the common utility function $u(x) = x_1^a x_2^{1-a}$, which is called the Cobb-Douglas utility function. This function is difficult to analyze because x_1 and x_2 have different exponents and are multiplied together. But, consider the monotonically increasing function $f(z) = \log(z)$, where “log” refers to the natural logarithm.³

$$V(x) = \log[x_1^a x_2^{1-a}] = a \log x_1 + (1-a) \log x_2$$

$V(\cdot)$ represents the same preferences as $U(\cdot)$. However, $V(\cdot)$ is a much easier function to deal with than $U(\cdot)$. In this way we can exploit the ordinal nature of utility to make our lives much easier. In other words, there are many utility functions that can represent the same preferences. Thus it may be in our interest to look for one that is easy to analyze.

A second implication of the ordinal nature of utility is that the difference between the utility of two bundles doesn't mean anything. For example, if $U(x) - U(y) = 7$ and $U(z) - U(a) = 14$, it doesn't mean that going from consuming z to consuming a is twice as much of an improvement than going from x to y . This makes it hard to compare things such as the impact of two different tax programs by looking at changes in utility. Fortunately, however, we have developed a method for dealing with this, using compensated changes similar to those used in the derivation of the Slutsky matrix in the section on consumer choice.

3.2.2 Some Basic Properties of Utility Functions


If preferences are convex, then the indifference curves will be convex, as will the upper level sets. When a function's upper-level sets are always convex, we say that the function is (sorry about this) quasiconcave. The importance of quasiconcavity will become clear soon. But, I just want to drill into you that quasiconcave means convex upper level sets. Keep that in mind, and things will be much easier.

For example, consider a special case of the Cobb-Douglas utility function I mentioned earlier.

$$u(x_1, x_2) = x_1^{\frac{1}{4}} x_2^{\frac{1}{4}}.$$

Figure 3.4 shows a three-dimensional (3D) graph of this function.

³Despite what you are used to, economists always use log to refer to the natural log, \ln , since we don't use base 10 logs at all.

Figure 3.4: Function $u(x)$





Figure 3.5: Level sets of $u(x)$

Notice the curvature of the surface. Now, consider Figure 3.5, which shows the level sets (I_x) for various utility levels. Notice that the indifference curves of this utility function are convex. Now, pick an indifference curve. Points offering more utility are located above and to the right of it. Notice how the contour map corresponds to the 3D utility map. As you move up and to the right, you move “uphill” on the 3D graph.

Quasiconcavity is a weaker condition than concavity. Concavity is an assumption about how the numbers assigned to indifference curves change as you move outward from the origin. It says that the increase in utility associated with an increase in the consumption bundle decreases as you move away from the origin. As such, it is a cardinal concept. Quasiconcavity is an ordinal concept. It talks only about the shape of indifference curves, not the numbers assigned to them. It can be shown that concavity implies quasiconcavity but a function can be quasiconcave without being concave (can you draw one in two dimensions). It turns out that for the results on utility

Figure 3.6: Function $v(x)$

maximization we will develop later, all we really need is quasiconcavity. Since concavity imposes cardinal restrictions on utility (which is an ordinal concept) and is stronger than we need for our maximization results, we stick with the weaker assumption of quasiconcavity.⁴

Here's an example to help illustrate this point. Consider the following function, which is also of the Cobb-Douglas form:

$$v(x) = x_1^3 x_2^3 .$$

Figure 3.6 shows the 3D graph for this function. Notice that $v(x)$ is “curved upward” instead of downward like $u(x)$. In fact, $v(x)$ is not a concave function, while $u(x)$ is a concave function.⁵ But, both are quasiconcave. We already saw that $u(x)$ was quasiconcave by looking at its level

⁴As in the case of convexity and strict convexity, a strictly quasiconcave function is one whose upper level sets are strictly convex. Thus a function that is quasiconcave but not strictly so can have flat parts on the boundaries of its indifference curves.

⁵See Simon and Blume or Chiang for good explanations of concavity and convexity in multiple dimensions.




Figure 3.7: Level sets of $v(x)$.

sets.⁶ To see why $v(x)$ is quasiconcave, let's look at the level sets of $v(x)$ in Figure 3.7. Even though $v(x)$ is curved in the other direction, the level sets of $v(x)$ are still convex. Hence $v(x)$ is quasiconcave. The important point to take away here is that quasiconcavity is about the shape of level sets, not about the curvature of the 3D graph of the function.

Before going on, let's do one more thing. Recall $u(x) = x^4 x^4$ and $v(x) = x^2 x^2$. Now, consider

the monotonic transformation $f(u) = u^6$. We can rewrite $v(x) = x^4 x^4 = (x^4 x^4)^{1/2} = f(u(x))$. Hence utility functions $u(x)$ and $v(x)$ actually represent the same preferences! Thus we see that utility and preferences have to do with the shape of indifference curves, not the numbers assigned to them. Again, utility is an ordinal, not cardinal, concept.

Now, here's an example of a function that is not quasiconcave.

$$h(x) = x^2 + y^2 - \frac{1}{4} \sin \left(\frac{3}{8} \arctan \frac{y}{x} \right)$$

⁶That isn't a proof, just an illustration.




Figure 3.8: Function $h(x)$

Don't worry about where this comes from. Figure 3.8 shows the 3D plot of $h(x)$.

Figure 3.9 shows the isoquants for this utility function. Notice that the level sets are not convex. Hence, function $h(x)$ is not quasiconcave. After looking at the mathematical analysis of the consumer's problem in the next section, we'll come back to why it is so hard to analyze utility functions that look like $h(x)$.

3.3 The Utility Maximization Problem (UMP)

Now that we have defined a utility function, we are prepared to develop the model in which consumers choose the bundle they most prefer from among those available to them.⁷ In order to

⁷Notice that in the choice model, we never said why consumers make the choices they do. We only said that those choices must appear to satisfy homogeneity of degree zero, Walras' law, and WARP. Now, we say that the consumer acts to maximize utility with certain properties.




Figure 3.9: Level Sets of $h(x)$

ensure that the problem is “well-behaved,” we will assume that preferences are rational, continuous, convex, and locally nonsatiated. These assumptions imply that the consumer has a continuous utility function $u(x)$, and the consumer’s choices will satisfy Walras’ Law. In order to use calculus techniques, we will assume that $u(\cdot)$ is differentiable in each of its arguments. Thus, in other words, we assume indifference curves have no “kinks.”

The consumer’s problem is to choose the bundle that maximizes utility from among those available. The set of available bundles is given by the Walrasian budget set $B_{p,w} = \{x \in X \mid p \cdot x \leq w\}$. We will assume that all prices are strictly positive (written $p \gg 0$) and that wealth is strictly positive as well. The consumer’s problem can be written as:

$$\begin{array}{ll} \max_{x \geq 0} & u(x) \\ \text{s.t.} & p \cdot x \leq w. \end{array}$$

The first question we should ask is: Does this problem have a solution? Since $u(x)$ is a continuous function and $B_{p,w}$ is a closed and bounded (i.e., compact) set, the answer is yes by the Weierstrass theorem - a continuous function on a compact set achieves its maximum. How do we find the solution? Since this is a constrained maximization problem, we can use Lagrangian methods. The Lagrangian can be written as:

$$L = u(x) + \lambda (w - p \cdot x)$$

Which implies Kuhn-Tucker first-order conditions (FOC’s):

$$\begin{aligned} u_i(x^*) - \lambda^* p_i &\leq 0 \text{ and } x_i (u_i(x^*) - \lambda^* p_i) = 0 \text{ for } i = (1, \dots, L) \\ w - p \cdot x^* &\geq 0 \text{ and } \lambda^* (w - p \cdot x^*) = 0 \end{aligned}$$

Note that the optimal solution is denoted with an asterisk. This is because the first-order conditions don’t hold everywhere, only at the optimum. Also, note that the value of the Lagrange multiplier λ is also derived as part of the solution to this problem.

Now, we have a system with $L + 1$ unknowns. So, we need $L + 1$ equations in order to solve for the optimum. Since preferences are locally non-satiated, we know that the consumer will choose a consumption bundle that is on the boundary of the budget set. Thus the constraint must bind. This gives us one equation.

The conditions on x_i are complicated because we must allow for the possibility that the consumer chooses to consume $x_i^* = 0$ for some i at the optimum. This will happen, for example, if the relative

price of good i is very high. While this is certainly a possibility, “corner solutions” such as these are not the focus of the course, so we will assume that $x_i^* > 0$ for all i for most of our discussion. But, you should be aware of the fact that corner solutions are possible, and if you come across a corner solution, it may appear to behave strangely.

Generally speaking, we will just assume that solutions are interior. That is, that $x_i^* > 0$ for all commodities i . In this case, the optimality condition becomes

$$u_i(x_i^*) - \lambda^* p_i = 0. \quad (3.1)$$

Solving this equation for λ^* and doing the same for good j yields:

$$\frac{u_i(x_i^*)}{p_i} = \frac{u_j(x_j^*)}{p_j} \quad \text{for all } i, j \in \{1, \dots, L\}.$$

This turns out to be an important condition in economics. The condition on the right is the slope of the budget line, projected into the i and j dimensions. For example, if there are two

commodities, then the budget line can be written $x_2 = -\frac{p_1}{p_2}x_1 + \frac{w}{p_2}$. The left side, on the other hand, is the slope of the utility indifference curve (also called an isoquant or isoutility curve).

To see why $-\frac{u_i(x_i^*)}{u_j(x_j^*)}$ is the slope of the isoquant, consider the following identity: $u(x_1, x_2(x_1)) \equiv k$, where k is an arbitrary utility level and $x_2(x_1)$ is defined as the level of x_2 needed to guarantee the consumer utility k when the level of commodity 1 consumed is x_1 . Differentiate this identity with respect to x_1 :⁸

$$u_1 + u_2 \frac{dx_2}{dx_1} = 0$$

$$\frac{dx_2}{dx_1} = -\frac{u_1}{u_2}$$

So, at any point (x_1, x_2) , $-\frac{u_1(x_1, x_2)}{u_2(x_1, x_2)}$ is the slope of the implicitly defined curve $x_2(x_1)$. But, this

curve is exactly the set of points that give the consumer utility k , which is just the indifference curve. As mentioned earlier, we call the slope of the indifference curve the marginal rate of substitution (MRS): $MRS = -\frac{u_1}{u_2}$.

Thus the optimality condition is that at the optimal consumption bundle, the MRS (the rate that the consumer is willing to trade good x_2 for good x_1 , holding utility constant) must equal the ratio of the prices of the two goods. In other words, the slope of the utility isoquant is the same as the slope of the budget line. Combine this with the requirement that the optimal bundle be on

⁸ Here, we adopt the common practice of using subscripts to denote partial derivatives, $\frac{\partial u(x)}{\partial x_i} \equiv u_i$.




Figure 3.10: Tangency Condition

the budget line, and this implies that the utility isoquant will be tangent to the budget line at the optimum.

In Figure 3.10, x^* is found at the point of tangency between the budget set and one of the utility isoquants. Notice that because the level sets are convex, there is only one such point. If the level sets were not convex, this might not be the case. Consider, for example, Figure 3.9. Here, for any budget set, there will be many points of tangency between utility isoquants and the budget set. Some will be maximizers, and some won't. The only way to find out whether a point is a maximizer is to go through the long and unpleasant process of checking the second-order conditions. Further, even once a maximizer is found, it may behave strangely. We discuss this point further in Section 3.3.1 below.

So, we are looking for the point of tangency between the budget set and a utility isoquant. One way to do this would be to use the following procedure:

1. Choose a point on the budget line, call it z . Find its upper level set U_z . Find $U_z \cap B_{p,w}$. This gives you the set of points that are feasible and at least as good as z .
2. If this set contains only z , you are done: z is the utility maximizing point. If this set contains more than just z , choose an arbitrary point on the budget line that is also inside U_z and repeat the process. Keep going until the only point that is in both the upper level set and the budget set is that point itself. This point is the optimum.

The problem with this procedure is that it could potentially take a very long time to find the optimal point. The calculus approach allows us to do it much faster by finding the point along the budget line that has the same slope as the indifference curve. This is a much easier task, but it turns out that it is really just a shortcut for the procedure outlined above.

3.3.1 Walrasian Demand Functions and Their Properties

So, suppose that we have found the utility maximizing point, x^* . What have we really found? Notice that if the prices and wealth were different, the utility maximizing point would have been different. For this reason, we will write the endogenous variable x^* as a function of prices and wealth, $x(p, w) = (x_1(p, w), x_2(p, w), \dots, x_L(p, w))$. This function gives the utility maximizing bundle for any values of p and w . We will call $x(p, w)$ the consumer's Walrasian demand function, although it is also sometimes called the Marshallian or ordinary demand function. This is to distinguish it from another type of demand function that we will study later.

As a consequence of what we have done, we can immediately derive some properties of the Walrasian demand function:

1. Walras' Law: $p \cdot x(p, w) \equiv w$ for all p and w . This follows from local nonsatiation. Recall the definition of local non-satiation: For any $x \in X$ and $\varepsilon > 0$ there exists a $y \in X$ such that $\|x - y\| < \varepsilon$ and $y \succ x$. Thus the only way for x to be the most preferred bundle is if there the nearby point that is better is not in the budget set. But, this can only happen if x satisfies $p \cdot x(p, w) \equiv w$.
2. Homogeneity of degree zero in (p, w) . The definition of homogeneity is the same as always. $x(\alpha p, \alpha w) = x(p, w)$ for all p, w and $\alpha > 0$. Just as in the choice based approach, the budget set does not change: $B_{p,w} = B_{\alpha p, \alpha w}$. Now consider the first-order condition:

$$-\frac{u_i(x_i^*)}{x_j^*} = -\frac{p_i}{p_j} \text{ for all } i, j \in \{1, \dots, L\}.$$

Suppose we multiply all prices by $\alpha > 0$. This makes the right hand side $-\frac{\alpha p_i}{\alpha p_j} = -\frac{p_i}{p_j}$, which is just the same as before. So, since neither the budget constraint nor the optimality condition are changed, the optimal solution must not change either.

3. Convexity of $x(p, w)$. Up until now we have been assuming that $x(p, w)$ is a unique point. However, it need not be. For example, if preferences are convex but not strictly convex, the

isoquants will have flat parts. If the budget line has the same slope as the flat part, an entire region may be optimal. However, we can say that if preferences are convex, the optimal region will be a convex set. Further, we can add that if preferences are strictly convex, so that $u(\cdot)$ is strictly quasiconcave, then $x(p, w)$ will be a single point for any (p, w) . This is because strict quasiconcavity rules out flat parts on the indifference curve.

A Note on Optimization: Necessary Conditions and Sufficient Conditions

Notice that we derived the first-order conditions for an optimum above. However, while these conditions are necessary for an optimum, they are not generally sufficient - there may be points that satisfy them that are not maxima. This is a technical problem that we don't really want to worry about here. To get around it, we will assume that utility is quasiconcave and monotone (and some other technical conditions that I won't even mention). In this case we know that the first-order conditions are sufficient for a maximum.

In most courses in microeconomic theory, you would be very worried about making sure that the point that satisfies the first-order conditions is actually a maximizer. In order to do this you need to check the second order conditions (make sure the function is "curved down"). This is a long and tedious process, and, fortunately, the standard assumptions we will make, strict quasiconcavity and monotonicity, are enough to make sure that any point that satisfies the first-order conditions is a maximizer (at least when the constraint is linear). Still, you should be aware that there is such things as second-order conditions, and that you either need to check to make sure they are satisfied or make assumptions to ensure that they are satisfied. We will do the latter, and leave the former to people who are going to be doing research in microeconomic theory.

A Word on Nonconvexities

It is worthwhile to spend another moment on what can happen if preferences are not convex, i.e. utility is not quasiconcave. We already mentioned that with nonconvex preferences it becomes necessary to check second-order conditions to determine if a point satisfying the first-order conditions is really a maximizer. There can also be other complications. Consider a utility function where the isoquants are not convex, shown in Figure 3.11.

When the budget line is given by line 1, the optimal point will be near x . When the budget line is line 2, the optimal points will be either x or y . But, none of the points between x and y on line




Figure 3.11: Nonconvex Isoquants

2 are as good as x or y (a violation of the idea that averages are better than extremes). Finally, if the budget line is given by line 3, the only optimal point will be near y . Thus the optimal point jumps from x to y without going through any intermediate values.

Now, lines 1, 2, and 3 can be generated by a series of compensated decreases in the price of good 1 (plotted on the horizontal axis). And, intuitively, it seems like people's behavior should change by a small amount if the price changes by a small amount. But, if the indifference curves are non-convex, behavior could change a lot in response to small changes in the exogenous parameters. Since non-convexities result in predictions that do not accord with how we feel consumers actually behave, we choose to model consumers as having convex preferences. In addition, non-convexities add complications to solving and analyzing the consumer's maximization problem that we are very happy to avoid, so this provides another reason why we assume preferences are convex.

Actually, the same sort of problem can arise when preferences are convex but not strictly convex. It could be that behavior changes a lot in response to small changes in prices (although it need not do so). In order to eliminate this possibility and ensure a unique maximizing bundle, we will generally assume that preferences are strictly convex and that utility is strictly quasiconcave (i.e., has strictly convex upper level sets).

3.3.2 The Lagrange Multiplier

You may recall that the optimal value of the Lagrange multiplier is the shadow value of the constraint, meaning that it is the increase in the value of the objective function resulting from a slight relaxation of the constraint. If you don't remember this, you should acquaint yourself with the

point by looking in the math appendix of your favorite micro text or “math for economists” book.⁹ If you still don’t believe this is true, I present you with the following derivation. In addition to showing this fact about the value of λ^* , it also illustrates a common method of proof in economics.

Consider the utility function $u(x)$. If we substitute in the demand functions, we get

$$u(x(p, w))$$

which is the utility achieved by the consumer when she chooses the best bundle she can at prices p and wealth w . The constraint in the problem is:

$$p \cdot x \leq w.$$

So, relaxing the constraint means increasing w by a small amount. If this is unfamiliar to you, think about why it is so: The budget set $B_{p, w+dw}$ strictly includes the budget set $B_{p, w}$, and so any bundle that could be chosen before the wealth increase could also be chosen after. Since there are more feasible points, the constraint after the wealth increase is a relaxation of the constraint before. We can analyze the effect of this by differentiating $u(x(p, w))$ with respect to w :

$$\begin{aligned} \frac{d}{dw} u(x(p, w)) &= \sum_{i=1}^L \frac{du_i}{dx_i} \frac{dx_i}{dw} \\ &= \sum_{i=1}^L \lambda p_i \frac{dx_i}{dw} \\ &= \lambda \sum_{i=1}^L p_i \frac{dx_i}{dw} \\ &= \lambda. \end{aligned}$$

The transition from the first line to the second line is accomplished by substituting in the first-order condition: $\frac{du_i}{dx_i} - \lambda p_i = 0$. The transition from the second line to the third line is trivial (you can factor out λ since it is a constant). The transition from the third line to the fourth line comes from the comparative statics of Walras’ Law that we derived in the choice section. Since $p \cdot x(p, w) \equiv w$, $\sum_{i=1}^L p_i \frac{dx_i}{dw} = 1$ (you could rederive this by differentiating the identity with respect to w if you want).

3.3.3 The Indirect Utility Function and Its Properties

The Walrasian demand function $x(p, w)$ gives the commodity bundle that maximizes utility subject to the budget constraint. If we substitute this bundle into the utility function, we get the utility

⁹If you don’t have a favorite, I recommend “Mathematics for Economists” by Simon and Blume.

that is earned when the consumer chooses the bundle that maximizes utility when prices are p and wealth is w . That is, define the function $v(p, w)$ as:

$$v(p, w) \equiv u(x(p, w)).$$

We call $v(p, w)$ the indirect utility function. It is indirect because while utility is a function of the commodity bundle consumed, x , indirect utility function $v(p, w)$ is a function of p and w . Thus it is indirect because it tells you utility as a function of prices and wealth, not as a function of commodities. You can think of it this way. Given prices p and wealth w , $x(p, w)$ is the commodity bundle chosen and $v(p, w)$ is the utility that results from consuming $x(p, w)$. But, if I know $v(p, w)$, then given any prices and wealth I can calculate utility without first having to solve for $x(p, w)$.

Just as $x(p, w)$ had certain properties, so does $v(p, w)$. In fact, most of them are inherited from the properties of $x(p, w)$. Suppose that preferences are locally nonsatiated. The indirect utility function corresponding to these preferences $v(p, w)$ has the following properties:

1. Homogeneity of degree zero: Since $x(p, w) = x(\alpha p, \alpha w)$ for $\alpha > 0$,

$$v(\alpha p, \alpha w) = u(x(\alpha p, \alpha w)) = u(x(p, w)) = v(p, w).$$

In other words, since the bundle you consume doesn't change when you scale all prices and wealth by the same amount, neither does the utility you earn.

2. $v(p, w)$ is strictly increasing in w and non-increasing in p_j . If $x_j > 0$, $v(p, w)$ is strictly decreasing in p_j . Indirect utility is strictly increasing in w by local non-satiation. If $x(p, w)$ is optimal and preferences are locally non-satiated, there must be a point just on the other side of the budget line that the consumer strictly prefers. If w increases a little bit, this point will become feasible, and the consumer will earn higher utility. Indirect utility is non-increasing in p_j since an increase in p_j shrinks the feasible region. After p_j increases, the budget line lies inside of the old budget set. Since the consumer could have chosen these points but didn't, the consumer can be no better off than before the price increase. Note that the consumer will do strictly worse unless it is the case that $x_j(p, w) = 0$ and the consumer still chooses to consume $x_j = 0$ after the price increase. In this case, the consumer's consumption bundle does not change, so neither does her utility. This is a subtle point having to do with corner solutions. But, a carefully drawn picture should make it all clear (See Figure 3.12).




Figure 3.12: $V(p, w)$ is decreasing in p

3. $v(p, w)$ is quasiconvex in (p, w) . In other words, the set $\{(p, w) \mid v(p, w) \leq \bar{v}\}$ is convex for all \bar{v} . Consider two distinct price-wealth vectors (p^0, w^0) and (p^{00}, w^{00}) such that $v(p^0, w^0) = v(p^{00}, w^{00})$. Since the consumer chooses her most preferred consumption bundle at each price, $x(p^0, w^0)$ is preferred to all other bundles in B_{p^0, w^0} and $x(p^{00}, w^{00})$ is preferred to all other bundles in $B_{p^{00}, w^{00}}$. Now, consider the budget set formed at an average price $p^a = ap^0 + (1-a)p^{00}$ and wealth $w^a = aw^0 + (1-a)w^{00}$. Every bundle in B_{p^a, w^a} will lie within either B_{p^0, w^0} or $B_{p^{00}, w^{00}}$. Hence the utility of any bundle in B_{p^a, w^a} can be no larger than the utility of the chosen bundle, $v(p^0, w^0)$. Thus $v(p^0, w^0) \geq v(p^a, w^a)$, which proves the result. Note: see the diagram on page 57 of MWG.¹⁰ A question for you: What is the intuitive meaning of this?
4. $v(p, w)$ is continuous in p and w . Small changes in p and w result in small changes in utility. This is especially clear in the case where indifference curves are strictly convex and differentiable.

3.3.4 Roy's Identity

Consider the indirect utility function: $v(p, w) = \max_{x \in B_{p, w}} u(x)$. The function $v(p, w)$ tells you how much utility the consumer earns when prices are p and wealth is w . Thanks to a very clever bit of mathematics, we can exploit this in order to figure out the relationship between the indirect utility function and the demand functions $x(p, w)$.

¹⁰ This is a slightly informal argument. Formally, we must show for any (p^0, w^0) and (p^{00}, w^{00}) , $v(p^a, w^a) \leq \max\{v(p^0, w^0), v(p^{00}, w^{00})\}$. But, the same intuition continues to hold.

The definition of the indirect utility function implies that the following identity is true:

$$v(p, w) \equiv u(x(p, w)).$$

Differentiate both sides with respect to p_i :

$$\frac{\partial v}{\partial p_i} = \sum_{i=1}^L \frac{\partial u}{\partial x_i} \frac{\partial x_i}{\partial p_i}.$$

But, based on the first-order conditions for utility maximization, which we know hold when $u(\cdot)$ is evaluated at the optimal x , $x(p, w)$ (see equation (3.1)):

$$\frac{\partial u}{\partial x_i} = \lambda p_i.$$

And, we also know (from Section 3.3.2) that the Lagrange multiplier is the shadow price of the constraint: $\lambda = \frac{\partial v}{\partial w}$. Hence:

$$\begin{aligned} \frac{\partial v}{\partial p_i} &= \sum_{i=1}^L \frac{\partial u}{\partial x_i} \frac{\partial x_i}{\partial p_i} = \sum_{i=1}^L \lambda p_i \frac{\partial x_i}{\partial p_i} \\ &= \lambda \sum_{i=1}^L p_i \frac{\partial x_i}{\partial p_i} = \frac{\partial v}{\partial w} \sum_{i=1}^L p_i \frac{\partial x_i}{\partial p_i}. \end{aligned}$$

Now, recall the comparative statics result of Walras' Law with respect to a change in p_i :

$$x_i(p, w) = - \sum_{i=1}^L p_i \frac{\partial x_i}{\partial p_i}.$$

Substituting this in yields:

$$\begin{aligned} \frac{\partial v}{\partial p_i} &= - \frac{\partial v}{\partial w} x_i(p, w) \\ x_i(p, w) &= - \frac{\partial p_i \frac{\partial v}{\partial p_i}}{\partial w}. \end{aligned}$$

The last equation, known as Roy's identity, allows us to derive the demand functions from the indirect utility function. This is useful because in many cases it will be easier to estimate an indirect utility function and derive the direct demand functions via Roy's identity than to derive $x(p, w)$ directly. Estimating Roy's identity involves estimating a single equation. Estimating $x(p, w)$, on the other hand, amounts to finding for every value of p and w the solution to a set of $L + 1$ first-order equations, which themselves may have unknown parameters.

3.3.5 The Indirect Utility Function and Welfare Evaluation

Consider the situation where the price of good 1 increases from p_1 to p_1^0 . What is the impact of this price change on the consumer? One way to measure it is in terms of the indirect utility function:

$$\text{impact} = v(p_1^0, w) - v(p, w).$$

That is, the impact of the price change is equal to the difference in the consumer's utility at prices p_1^0 and p . While this is certainly a measure of the impact of the price change, it is essentially useless. There are a number of reasons why, but they all hinge on the fact that utility is an ordinal, not a cardinal concept. As you recall, the only meaning of the numbers assigned to bundles by the utility function is that $x \succ y$ if and only if $u(x) > u(y)$. In particular, if $u(x) = 2u(y)$, this does not mean that the consumer likes bundle x twice as much as bundle y . Also, if $u(x) - u(z) > u(s) - u(t)$, this doesn't mean that the consumer would rather switch from bundle z to x than from t to s . Because of this, there is really nothing we can make of the numerical value of $v(p_1^0, w) - v(p, w)$. The only thing we can say is that if this difference is positive, the consumer likes (p_1^0, w) more than (p, w) . But, we can't say how much more.

Another problem with using the change in the indirect utility function as a measure of the impact of a policy change is that it cannot be compared across consumers. Comparing the change in two different utility functions is even more meaningless than comparing the change in a single person's utility function. This is because even if both utility functions were cardinal measures of the benefit to a consumer (which they aren't), there would still be no way to compare the scales of the two utility functions. This is the "problem of interpersonal comparison of utility," which arises in many aspects of welfare economics.

As a possible solution to the problem, consider the following thought experiment. Initially, prices and wealth are given by (p, w) . I am interested in measuring the impact of a change in prices to p_1^0 . So, I ask you the following question: By how much would I have to change your wealth so that you are indifferent between (p_1^0, w) and (p, w^0) ? That is, for what value w^0 does

$$v(p_1^0, w) = v(p, w^0).$$

The change in wealth, $w^0 - w$, in essence gives a monetary value for the impact of this change in price. And, this monetary value is a better measure of the impact of the price change than the utility measurement, because it is, at least to a certain extent, comparable.¹¹ You can compare

¹¹ I say to a certain extent, because the value of an additional Δw dollars of wealth will depend on the initial state.

the impact of two different changes in prices by looking at the associated changes in wealth needed to compensate the consumer. Similarly, you can compare the impact of price changes on different consumers by comparing the changes in wealth necessary to leave them just as well off.¹²

Finally, although using the amount of money needed to compensate the consumer is an imperfect measure of the impact of a policy decision, it has one huge benefit for the neoclassical economist, and that is that it is observable, at least in principle. There is nothing we can do to observe utility scales. However, we can often elicit from people the amount of money they would find equivalent to a certain policy change, either through experiments, surveys, or other estimation techniques.

3.4 The Expenditure Minimization Problem (EMP)

In the previous section, I argued that a good measure of the impact of a change in prices was the change in wealth necessary to make the consumer as well off at the old prices and new wealth as she was at the new prices and old wealth. However, this is not an easy exercise when all you have to work with is the indirect utility function. If we had a function that tells you how much wealth you would need to have in order to achieve a certain level of utility, then we would be able to do this much more efficiently. There is such a function. It is called the expenditure function, and in this section we will develop it.

The expenditure minimization problem (EMP) asks the question, if prices are p , what is the minimum amount the consumer would have to spend to achieve utility level u ? That is:

$$\begin{array}{ll} \min_x & p \cdot x \\ \text{s.t.} & u(x) \geq u. \end{array}$$

Before we go on, let's take a moment to figure out what the endogenous and exogenous variables are here. The exogenous variables are prices p and the reservation (or target) utility level u . The endogenous variable is x , the consumption bundle. So, in words, the expenditure minimization bundle amounts to finding the bundle x that minimizes the cost of achieving utility u when prices are p .

The Lagrangian for this problem is given by:

$$LEMP = p \cdot x - \lambda (u(x) - u).$$

For instance, poor people presumably value the same wealth increment more than rich people.

¹² Again, this measure is imperfect because it assumes that the two consumers have the same marginal utility of wealth.

Assuming an interior solution, the first-order conditions are given by:¹³

$$\begin{aligned} p_i - \lambda u_i(x) &= 0 \text{ for } i \in \{1, \dots, L\} \\ \lambda (u(x) - u) &= 0 \end{aligned} \quad (3.2)$$

If $u(\cdot)$ is well behaved (e.g., quasiconcave and increasing in each of its arguments), then the constraint will bind, and the second condition can be written as $u(x) = u$. Further, a unique solution to this problem will exist for any values of p and u . We will denote the value of the solution to this problem by $h(p, u) \in X$. That is, $h(p, u)$ is an L dimensional vector whose l^{th} component, $h_l(p, u)$ gives the amount of commodity l that is consumed when the consumer minimizes the cost of achieving utility u at prices p . The function $h(p, u)$ is known as the Hicksian (or compensated) demand function.¹⁴ It is a demand function because it specifies a consumption bundle. It differs from the Walrasian (or ordinary) demand function in that it takes p and u as its arguments, whereas the Walrasian demand function takes p and w as its arguments.

In other words, $h(p, u)$ and $x(p, w)$ are the answers to two different but related problems. Function $x(p, w)$ answers the question, "Which commodity bundle maximizes utility when prices are p and wealth is w ?" Function $h(p, u)$ answers the question, "Which commodity bundle minimizes the cost of attaining utility u when prices are p ?" We'll return to the difference between the two types of demand shortly.

Since $h(p, u)$ solves the EMP, substitution of $h(p, u)$ into the first-order conditions for the EMP yields the identities (assuming the constraint binds):

$$\begin{aligned} p_i - \lambda u_i(h(p, u)) &\equiv 0 \text{ for } i \in \{1, \dots, L\} \\ u(h(p, u)) - u &\equiv 0 \end{aligned}$$

Further, just as we defined the indirect utility function as the value of the objective function of the UMP, $u(x)$, evaluated at the optimal consumption bundle, $x(p, w)$, we can also define such a function for the EMP. The expenditure function, denoted $e(p, u)$, is defined by:

$$e(p, u) \equiv p \cdot h(p, u)$$

and is equal to the minimum cost of achieving utility u , for any given p and u .

¹³ Again, remember that if $f(y)$ is a function with a vector y as its argument, the notation f_i will frequently be used as shorthand notation for $\frac{\partial f}{\partial y_i}$. Thus u_i denotes $\frac{\partial u}{\partial x_i}$.

¹⁴ We'll return to why $h(p, u)$ is called the compensated demand function in a while.

3.4.1 A First Note on Duality

Consider the first-order conditions (from (3.2)) for x_i and x_j . Solving each for λ yields:

$$\begin{aligned} \frac{p_i}{u_i} &= \lambda = \frac{p_j}{u_j} \\ \frac{u_i}{u_j} &= \frac{p_i}{p_j} . \end{aligned} \tag{3.3}$$

Recall that this is the same tangency condition we derived in the UMP. What does this mean? Consider a price vector p and wealth w . The bundle that solves the UMP, $x^* = x(p, w)$ is found at the point of tangency between the budget line and the consumer's utility isoquant. The consumer's utility at this point is given by $u^* = u(x^*)$. Thus x^* is the point of tangency between the line $p \cdot x = w$ and the curve $u(x) = u^*$.

Now, consider the EMP when the target utility level is given by u^* . The bundle that solves the EMP is the bundle that achieves utility u^* at minimum cost. This is located by finding the point of tangency between the curve $u(x) = u^*$ and a budget line (which is what (3.3) says). But, we already know from the UMP that the curve $u(x) = u^*$ is tangent to the budget line $p \cdot x = w$ at x^* (and is tangent to no other budget line). Hence x^* must solve the EMP problem when the target utility level is u^* ! Further, since x^* lies on the budget line, $p \cdot x^* = w$. So the minimum cost of achieving utility u^* is w . Thus the UMP and the EMP pick out the same point.

Let me restate what I've just argued. If x^* solves the UMP when prices are p and wealth is w , then x^* solves the EMP when prices are p and the target utility level is $u(x^*)$. Further, maximal utility in the UMP is $u(x^*)$ and minimum expenditure in the EMP is w . This result is called the "duality" of the EMP and the UMP.

The UMP and the EMP are considered dual problems because the constraint in the UMP is the objective function in the EMP and vice versa. This is illustrated by looking at the graphical solutions to the two problems. In the UMP, shown in Figure 3.13, you keep increasing utility until you find the one that is tangent to the budget line. In the EMP, on the other hand, shown in Figure 3.14, you keep decreasing expenditure (which is like shifting a budget line toward the origin) until you find the expenditure line that is tangent to $u(x) = u^*$. Although the process of finding the optimal point is different in the UMP and EMP, they both pick out the same point because they are looking for the same basic relationship, as expressed in equation (3.3).

The duality relationship between the EMP and the UMP is captured by the following identities,




Figure 3.13: The Utility Maximization Problem




Figure 3.14: The Expenditure Minimization Problem

to which we will return later:

$$h(p, v(p, w)) \equiv x(p, w)$$

$$x(p, e(p, u)) \equiv h(p, u) .$$

These identities restate the principles discussed previously. The first says that the commodity bundle that minimizes the cost of achieving the maximum utility you can achieve when prices are p and wealth is w is the bundle that maximizes utility when prices are p and wealth is w . The second says that the bundle that maximizes utility when prices are p and wealth is equal to the minimum amount of wealth needed to achieve utility u at those prices is the same as the bundle that minimizes the cost of achieving utility u when prices are p .

Similar identities can be written using the indirect utility function and expenditure function:

$$u \equiv v(p, e(p, u))$$

$$w \equiv e(p, v(p, w)) .$$

Note to MWG readers: There is a mistake in Figure 3.G.3. The relationships on the horizontal line connecting $v(p, w)$ and $e(p, u)$ should be the ones written directly above.

The main implication of the previous analysis is this: The expenditure function contains the exact same information as the indirect utility function. And, since the indirect utility function can be used (by Roy's identity) to derive the Walrasian demand functions, which can, in turn, be used to recover preferences, the expenditure function contains the exact same information as the utility function. This means that if you know the consumer's expenditure function, you know her utility function, and vice versa. No information is lost along the way. This is another expression of what people mean when they say that the UMP and EMP are dual problems - they contain exactly the same information.

3.4.2 Properties of the Hicksian Demand Functions and Expenditure Function

In this section, we refer both to function $u(x)$ and to a particular level of utility, u . In order to be clear, let's put a bar over the u when we are talking about a level of utility, i.e., \bar{u} . Just as we derived the properties of $x(p, w)$ and $v(p, w)$, we can also derive the properties of the Hicksian demand functions $h(p, \bar{u})$ and expenditure function $e(p, \bar{u})$. Let's begin with $h(p, \bar{u})$. We will assume that $u(\cdot)$ is a continuous utility function representing a locally non-satiated preference relation.

Properties of the Hicksian Demand Functions

The Hicksian demand functions have the following properties:

1. Homogeneity of degree zero in p : $h(\alpha p, u^-) \equiv h(p, u^-)$ for p, u, \bar{u} and $\alpha > 0$. NOTE: THIS IS HOMOGENEITY IN P , NOT HOMOGENEITY IN P AND U ! Homogeneity of degree zero is best understood in terms of the graphical presentation of the EMP. The solution to the EMP is the point of tangency between the utility isoquant $u(x) = u^-$ and one of the budget lines. This is determined by the slope of the expenditure lines (lines of the form $p \cdot x = k$, where k is any constant). Any change that doesn't affect the slope of the budget lines should not affect the cost-minimizing bundle (although it will affect the expenditure on the cost minimizing bundle). Since the slope of the expenditure line is determined by relative prices and since scaling all prices by the same amount does not affect relative prices, the solution should not change. More formally, the EMP at prices αp is

$$\begin{aligned} & \min_x \alpha p \cdot x \\ \text{s.t.} & : u(x) \geq u^- \end{aligned}$$

But, this problem is formally equivalent to:

$$\min \alpha (p \cdot x) : \text{s.t.} : u(x) \geq x$$

which is equivalent to:

$$\alpha \min_x p \cdot x : \text{s.t.} : u(x) \geq x$$

which is just the same as the EMP when prices are p , except that total expenditure is multiplied by α , which doesn't affect the cost minimizing bundle.

2. No excess utility: $u(h(p, u^-)) = u^-$. This follows from the continuity of $u(\cdot)$. Suppose $u(h(p, u^-)) > u^-$. Then consider a bundle h^0 that is slightly smaller than $h(p, u^-)$ on all dimensions. By continuity, if h^0 is sufficiently close to $h(p, u^-)$, then $u(h^0) > u^-$ as well. But, then h^0 is a bundle that achieves utility u^- at lower cost than $h(p, u^-)$, which contradicts the assumption that $h(p, u^-)$ was the cost minimizing bundle in the first place.¹⁵ From this we can conclude that the constraint always binds in the EMP.

¹⁵ This type of argument - called "Proof by Contradiction" - is quite common in economics. If you want to show a implies b , assume that b is false and show that if b is false then a must be false as well. Since a is assumed to be true, this implies that b must be true as well.

3. If preferences are convex, then $h(p, u^-)$ is a convex set. If preferences are strictly convex (i.e. $u(\cdot)$ is strictly quasiconcave), then $h(p, u^-)$ is single valued.

Properties of the Expenditure Function

Based on the properties of $h(p, u^-)$, we can derive properties of the expenditure function, $e(p, u^-)$.

1. Function $e(p, u^-)$ is homogeneous of degree one in p : Since $h(p, u^-)$ is homogeneous of degree zero in p , this means that scaling all prices by $\alpha > 0$ does not affect the bundle demanded. Applying this to total expenditure:

$$e(\alpha p, u^-) = \alpha p \cdot h(\alpha p, u^-) = \alpha p \cdot h(p, u^-) = \alpha e(p, u^-).$$

In words, if all prices change by a factor of α , the same bundle as before achieves utility level u^- at minimum cost, only it now costs you twice as much as it used to. This is exactly what it means for a function to be homogeneous of degree one.

2. Function $e(p, u^-)$ is strictly increasing in u^- and non-decreasing in $|p|$ for any l . I'll

give the argument here to show that $e(p, u^-)$ cannot be decreasing in u^- . There are a few more details to show that it cannot stay constant either, but most of the intuition of the argument is contained in showing that $e(p, u^-)$ cannot be strictly decreasing. The argu-

ment is by contradiction. Suppose that for $u^{-0} > u^-$, $e(p, u^-) > e(p, u^{-0})$. But, then $h(p, u^{-0})$

satisfies the constraint $u(x) \geq u^-$ and does so at lower cost than $h(p, u^-)$, which contradicts the assumption that $h(p, u^-)$ is the cost minimizing bundle that achieves utility level u^- .

The argument that $e(p, u^-)$ is nondecreasing in $|p|$ uses another method which is quite common,

a method I call "feasible but not optimal." Let p and p^0 differ only in component l , and let $p^0_l > p_l$. From the definition of the expenditure function, $e(p^0, u^-) = p^0 \cdot h(p^0, u^-) = p \cdot h(p^0, u^-) + e(p, u^-)$. The first equality follows from the definition of the expenditure function, the first

\geq follows from the fact that $p^0 > p$ (note: $p^0 \cdot h(p^0, u^-) > p \cdot h(p, u^-)$ if $h_l(p^0, u^-) > 0$), and the second \geq follows from the fact that $h(p^0, u^-)$ achieves utility level u^- but does not necessarily do so at minimum cost (i.e. $h(p^0, u^-)$ is feasible in the EMP for (p, u^-) but not necessarily optimal).

3. Function $e(p, u^-)$ is concave in p .¹⁶ Consider two price vectors p and p^0 , and let $p^a =$

¹⁶ Recall the definition of concavity. Consider y and y^0 such that $y \succ y^0$. Function $f(y)$ is concave if, for any $a \in [0, 1]$, $f(ay + (1-a)y^0) \geq af(y) + (1-a)f(y^0)$.

$$ap + (1 - a)p^0.$$

$$\begin{aligned} e(p^a, u^-) &= p^a \cdot h(p^a, u^-) \\ &= ap \cdot h(p^a, u^-) + (1 - a)p \cdot h(p^a, u^-) \\ &\geq ap \cdot h(p, u^-) + (1 - a)p^0 \cdot h(p^0, u^-) \end{aligned}$$

where the first line is the definition of $e(p, u^-)$, the second follows from the definition of p^a , and the third follows from the fact that $h(p^a, u^-)$ is feasible but not optimal in the EMP for (p, u^-) and (p^0, u^-) .

The following heuristic explanation is also helpful in understanding the concavity of $e(p, u^-)$. Suppose prices change from p to p^0 . If the consumer continued to consume the same bundle at the old prices, expenditure would increase linearly:

$$\Delta \text{expenditure} = p^0 \cdot h(p, u^-) - p \cdot h(p, u^-).$$

But, in general the consumer will not continue to consume the same bundle after the price change. Rather, he will rearrange his bundle in order to minimize the cost of achieving u^- at the new prices, p^0 . Since this will save the consumer some money, total expenditure will decrease at less than a linear rate. And, an alternate definition of concavity is that the function always increases at less than a linear rate. In other words, $f(x)$ is concave if it always lies below its tangent lines.¹⁷

3.4.3 The Relationship Between the Expenditure Function and Hicksian Demand

Just as there was a relationship between the indirect utility function $v(p, w)$ and the Walrasian demand functions $x(p, w)$, there is also a relationship between the expenditure function $e(p, u^-)$ and the Hicksian demand function $h(p, u^-)$. In fact, it is even more straightforward for $e(p, u^-)$ and $h(p, u^-)$. Let's start with the derivation

$$e(p, u^-) \equiv p \cdot h(p, u^-)$$

Since this is an identity, differentiate it with respect to p_i :

$$\frac{\partial e}{\partial p_i} \equiv h_i(p, u^-) + \sum_j p_j \frac{\partial h_j}{\partial p_i}.$$

¹⁷ This explanation will be clearer once we show that $h_i(p, u) = \frac{\partial e(p, u)}{\partial p_i}$, i.e. that $h_i(p, u)$ is exactly the rate of

increase in expenditure if p_i increases by a small amount. Thus $e(p^0, u) - e(p, u) \leq h(p, u) \cdot (p^0 - p)$ is exactly the definition of concavity.

Now, substitute in the first-order conditions, $p_j = \lambda u_j$

$$\frac{\partial e}{\partial p_i} \equiv h_i(p, u^-) + \lambda \sum_j u_j \frac{\partial h_j}{\partial p_i} \quad (3.4)$$

Since the constraint binds at any optimum of the EMP,

$$u(h(p, u^-)) \equiv u^-$$

Differentiate with respect to p_i :

$$\sum_j u_j \frac{\partial h_j}{\partial p_i} = 0$$

and substituting this into (3.4) yields:

$$\frac{\partial e}{\partial p_j} \equiv h_j(p, u^-) \quad (3.5)$$

That is, the derivative of the expenditure function with respect to p_j is just the Hicksian demand for commodity j .

The importance of this result is similar to the importance of Roy's identity. Frequently, it will be easier to measure the expenditure function than the Hicksian demand function. Since we are able to derive the Hicksian demand function from the expenditure function, we can derive something that is hard to observe from something that is easier to observe.

From (3.5) we can derive several additional properties (assuming $u(\cdot)$ is strictly quasiconcave and $h(\cdot)$ is differentiable):

1. (a) $\frac{\partial h_i}{\partial p_j} = \frac{\partial^2 e}{\partial p_i \partial p_j}$. This one follows directly from the fact that (3.5) is an identity. Let

$D_{ph}(p, u^-)$ be the matrix whose i th row and j th column is $\frac{\partial h_i}{\partial p_j}$. This property is thus the same as saying that $D_{ph}(p, u^-) \equiv D_p^2 e(p, u^-)$, where $D_p^2 e(p, u^-)$ is the matrix of second derivatives (Hessian matrix) of $e(p, u^-)$.

- (b) $D_{ph}(p, u^-)$ is a negative semi-definite (n.s.d.) matrix. This follows from the fact that

$e(p, u^-)$ is concave, and concave functions have Hessian matrices that are n.s.d. The main implication is that the diagonal elements are non-positive, i.e., $\frac{\partial h_i}{\partial p_i} \leq 0$.

- (c) $D_{ph}(p, u^-)$ is symmetric. This follows from Young's Theorem (that it doesn't matter what order you take derivatives in): $\frac{\partial h_i}{\partial p_j} = \frac{\partial^2 e}{\partial p_i \partial p_j} = \frac{\partial h_j}{\partial p_i}$. The implication is that the

cross-effects are the same — the effect of increasing p_j on h_i is the same as the effect of increasing p_i on h_j .




Figure 3.15: Compensated Demand

- (d) $P \frac{\partial h_i}{\partial p_j} p_j = 0$ for all i . This follows from the homogeneity of degree zero of $h(p, u^-)$ in p . Consider the identity:

$$h(ap, u^-) \equiv h(p, u^-).$$

Differentiate with respect to a and evaluate at $a = 1$, and you have this result.

The Hicksian demand curve is also known as the compensated demand curve. The reason for this is that implicit in the definition of the Hicksian demand curve is the idea that following a price change, you will be given enough wealth to maintain the same utility level you did before the price change. So, suppose at prices p you achieve utility level u^- . The change in Hicksian demand for good i following a change to prices p^0 is depicted in Figure 3.15.

When prices are p , the consumer demands bundle x , which has total expenditure $p \cdot x = w$. When prices are p^0 , the consumer demands bundle x^0 , which has total expenditure $p^0 \cdot x^0 = w^0$. Thus implicit in the definition of the Hicksian demand curve is the idea that when prices change from p to p^0 , the consumer is compensated by changing wealth from w to w^0 so that she is exactly as well off in utility terms after the price change as she was before.

Note that since $\frac{\partial h_i}{\partial p_i} \leq 0$, this is another statement of the compensated law of demand (CLD).

When the price of a good goes up and the consumer is compensated for the price change, she will not consume more of the good. The difference between this version and the previous version we saw (in the choice based approach) is that here, the compensation is such that the consumer can achieve the same utility before and after the price change (this is known as Hicksian substitution), and in the previous version of the CLD the consumer was compensated so that she could just afford the same bundle as she did before (this is known as Slutsky compensation). It turns out that the

two types of compensation yield very similar results, and, in fact, for differential changes in price, they are identical.

3.4.4 The Slutsky Equation

Recall that the whole point of the EMP was to generate concepts that we could use to evaluate welfare changes. The purpose of the expenditure function was to give us a way to measure the impact of a price change in dollar terms. While the expenditure function does do this (you can just look at $e(p^0, u) - e(p, u)$), it suffers from another problem. The expenditure function is based on the Hicksian demand function, and the Hicksian demand function takes as its arguments prices and the target utility level u . The problem is that while prices are observable, utility levels certainly are not. And, while we can generate some information by asking people over and over again how they compare certain bundles, this is not a very good way of doing welfare comparisons.

To summarize our problem: The Walrasian demand functions are based on observables (p, w) but cannot be used for welfare comparisons. The Hicksian demand functions, on the other hand, can be used to make welfare comparisons, but are based on unobservables.

The solution to this problem is to somehow derive $h(p, u)$ from $x(p, w)$. Then we could use our observations of p and w to derive $h(p, u)$, and use $h(p, u)$ for welfare evaluation. Fortunately, we can do exactly this. Suppose that $u(x(p, w)) = u$ (which implies that $e(p, u) = w$), and consider the identity:

$$h_i(p, u) \equiv x_i(p, e(p, u)) .$$

Differentiate both sides with respect to p_j :

$$\begin{aligned} \frac{\partial h_i}{\partial p_j} &\equiv \frac{\partial x_i(p, e(p, u))}{\partial p_j} + \frac{\partial x_i(p, e(p, u))}{\partial e(p, u)} \frac{\partial e(p, u)}{\partial p_j} \\ &\equiv \frac{\partial x_i(p, w)}{\partial p_j} + \frac{\partial x_i(p, w)}{\partial w} h_j(p, u) \\ &\equiv \frac{\partial x_i(p, w)}{\partial p_j} + \frac{\partial x_i(p, w)}{\partial w} x_j(p, e(p, u)) \\ &\equiv \frac{\partial x_i(p, w)}{\partial p_j} + \frac{\partial x_i(p, w)}{\partial w} x_j(p, w) . \end{aligned}$$

The equation

$$\frac{\partial h_i(p, v(p, w))}{\partial p_j} \equiv \frac{\partial x_i(p, w)}{\partial p_j} + \frac{\partial x_i(p, w)}{\partial w} x_j(p, w)$$

is known as the Slutsky equation. Note that it provides the link between the Walrasian demand functions $x(p, w)$ and the Hicksian demand functions, $h(p, u)$. Thus if we estimate the right-hand

side of this equation, which is a function of the observables p and w , then we can derive the value of the left-hand side of the equation, even though it is based on unobservable u .

Recall that implicit in the idea of the Hicksian demand function is the idea that the consumer's wealth would be adjusted so that she can achieve the same utility after a price change as she did before. This idea is apparent when we look at the Slutsky equation. It says that the change in demand when the consumer's wealth is adjusted so that she is as well off after the change as she

was before is made up of two parts. The first, $\frac{\partial x_i(p,w)}{\partial p_j}$, is equal to how much the consumer would change demand if wealth were held constant. The second, $\frac{\partial x_i(p,w)}{\partial w} x_i(p,w)$, is the additional change in demand following the compensation in wealth.

For example, consider an increase in the price of gasoline. If the price of gasoline goes up by one unit, consumers will tend to consume less of it, if their wealth is held constant (since it is not a Giffen good). However, the fact that gasoline has become more expensive means that they will have to spend more in order to achieve the same utility level. The amount by which they will have to be compensated is equal to the change in price multiplied by the amount of gasoline the consumer buys, $x_i(p,w)$. However, when the consumer is given $x_i(p,w)$ more units of wealth to spend, she will adjust her consumption of gasoline further. Since gasoline is normal, the consumer will increase her consumption. Thus the compensated change in demand (sometimes called the

pure substitution effect), $\frac{\partial h_i(p,v(p,w))}{\partial p_j}$, will be the sum of the uncompensated change (also known as the substitution effect), $\frac{\partial x_i(p,w)}{\partial p_j}$, and the wealth effect, $\frac{\partial x_i(p,w)}{\partial w} x_i(p,w)$.¹⁸

In order to make this clear, let's rearrange the Slutsky equation and go through the intuition again.

$$\frac{\partial x_i(p,w)}{\partial p_j} \equiv \frac{\partial h_i(p,v(p,w))}{\partial p_j} - \frac{\partial x_i(p,w)}{\partial w} x_j(p,w)$$

Here, we are interested in explaining an uncompensated change in demand in terms of the compensated change and the wealth effect. Think about the effect of an increase in the price of bananas on a consumer's Walrasian demand for bananas. If the price of bananas were to go up, and my wealth were adjusted so that I could achieve the same amount of utility before and after the change, I would consume fewer bananas. This follows directly from the CLD: $\frac{\partial h_i}{\partial p_i} \leq 0$. However, the change in

compensated demand assumes that the consumer will be compensated for the price change. Since an increase in the price of bananas is a bad thing, this means that $\frac{\partial h_i}{\partial p_i}$ has built into it the idea

¹⁸ This latter term is often called the "income effect," which is not quite right. Variable w stands for total wealth, which is more than just income. When people call this the income effect (as I sometimes do), they are just being sloppy.

that income will be increased in order to compensate the consumer. But, in reality consumers are not compensated for price changes, so we are interested in the uncompensated change in demand

$\frac{\partial x_i}{\partial p_i}$. This means that we must remove from the compensated change in demand the effect of the compensation. Since $\frac{\partial h_i}{\partial p_i}$ assumed an increase in wealth, we must impose a decrease in wealth,

which is just what the terms $-\frac{\partial x_i}{\partial w}(p, w) x_j(p, w)$ are. The decrease in wealth is given by $-x_i(p, w)$, and the effect of this decrease on demand for bananas is given by $\frac{\partial x_j}{\partial w}$.¹⁹

3.4.5 Graphical Relationship of the Walrasian and Hicksian Demand Functions

Demand functions are ordinarily graphed with price on the vertical axis and quantity on the horizontal axis, even though this is technically “backward.” But, we will follow with tradition and draw our graphs this way as well.

The difference between the compensated demand response to a price change and the uncom-pensated demand response to a price change is equal to the wealth effect:

$$\frac{\partial h_i(p, v(p, w))}{\partial p_j} \equiv \frac{\partial x_i(p, w)}{\partial p_j} + \frac{\partial x_i(p, w)}{\partial w} x_j(p, w)$$

Since $\frac{\partial h_i}{\partial p_i}$ is negative, when the wealth effect is positive (i.e., good i is normal) this means that the Hicksian demand curve will be steeper than the Walrasian demand curve at any point where they cross.²⁰ If, on the other hand, the wealth effect is negative (i.e. good i is inferior), this means that the Hicksian demand curve will be less steep than the Walrasian demand curve (see MWG Figure 3.G.1).

Let’s go into a bit more detail in working out the relative slopes of the Walrasian and Hicksian demand curves and determining how changes in u shift the Hicksian demand curve (depending on whether the good is normal or inferior).

In this subsection we discuss how changes in the exogenous parameters, p and u , affect the Hicksian Demand curve when it is drawn on the typical P - Q axes.

¹⁹ We are interested in the total change in consumption of bananas when the price of bananas goes up. In the real world, we don’t compensate people when prices change. But, the Slutsky equation tells us that the total (uncompensated) effect of a change in the price of bananas is a combination of the substitution effect (compensated effect) and the wealth effect. This result should be familiar to you from your intermediate micro course. If it isn’t, you may want to take a look at the (less abstract) treatment of this point in an intermediate micro text, such as Varian’s Intermediate Microeconomics. Test of understanding: If bananas are a normal good, could demand for bananas ever rise when the price increases (i.e. could bananas be a Giffen good)? Answer using the Slutsky equation.

²⁰ Remember that the graphs are backwards, so a less negative slope $\frac{\partial h_i}{\partial p_j}$ is actually steeper.

Part 0: Recall that $\frac{\partial h_i}{\partial p_i} \leq 0$ because the Slutsky matrix is negative semi-definite. We'll assume, as is typical, that $\frac{\partial h_i}{\partial p_i} < 0$. To keep things simple, I'll omit the subscripts for the rest of the subsection, since we're always talk about a single good.

Part 1: Relationship between Hicksian and Walrasian Demand.

By duality, we know that through any point on Walrasian demand there is a Hicksian demand curve through that point. How do their slopes compare? This is given by the Slutsky equation. But, keep in mind two things. First, since we put p on the vertical axis and x on the horizontal axis, when we draw a graph, the derivatives of the demand functions aren't slopes. They're inverse slopes. That is, the slope of the Walrasian demand is $\frac{1}{\frac{\partial x}{\partial p}}$ and the slope of the Hicksian is $\frac{1}{\frac{\partial h}{\partial p}}$. Second, these derivatives are usually negative. So, we have to be a bit careful about thinking about quantities that are larger (i.e., further to the right on the number line) and quantities that are larger in magnitude (i.e., are further from zero on the number line). Since slopes are negative, a "larger" slope corresponds to a flatter curve. You'll see why this is important in a minute.

To figure out whether $x(p, w)$ or $h(p, u)$ through a point is steeper, use the Slutsky equation.

$$\frac{\partial h}{\partial p} = \frac{\partial x}{\partial p} + \frac{\partial x}{\partial w} x.$$

The answer will depend on whether x is normal or inferior. So, begin by considering a normal good. In this case, $\frac{\partial x}{\partial p} x > 0$, so:

$$\begin{aligned} \frac{\partial h}{\partial p} &> \frac{\partial x}{\partial p} \\ \frac{\partial h}{\partial p} &< \frac{\partial x}{\partial p} \\ \frac{1}{\frac{\partial h}{\partial p}} &> \frac{1}{\frac{\partial x}{\partial p}} \end{aligned}$$

The first line comes from the Slutsky equation and the fact that the income effect is positive. The second comes from the fact that, for a normal good, both sides are negative, and hence if $\frac{\partial h}{\partial p} > \frac{\partial x}{\partial p}$, $\frac{\partial h}{\partial p}$ is smaller in magnitude (absolute value) than $\frac{\partial x}{\partial p}$. The third follows from the second since if $x < y$ and both are positive, then $1/x > 1/y$. Hence, for normal goods, the Hicksian Demand through a point is steeper than the Walrasian Demand through that point.

For an inferior good, things reverse. To simplify, suppose x is inferior but not Giffen (so that

$\frac{\partial x}{\partial p} < 0$ — you can do the Giffen case on your own). In this case $\frac{\partial x}{\partial p} x < 0$, so:

$$\begin{aligned} -\frac{\frac{\partial h}{\partial p}}{\frac{\partial h}{\partial x}} &< -\frac{\frac{\partial x}{\partial p}}{\frac{\partial x}{\partial x}} \\ -\frac{1}{\frac{\partial h}{\partial x}} &< -\frac{1}{\frac{\partial x}{\partial x}} \end{aligned}$$

and so Hicksian demand is flatter than Walrasian Demand.

Part 2: Dependence of Hicksian demand on u .

How does changing u shift the Hicksian Demand curve? Again, the answer depends on whether the good is normal or inferior. To see how, use duality:

$$h(p, u) \equiv x(p, e(p, u)),$$

and differentiate both sides with respect to u :

$$\frac{\partial h}{\partial u} \equiv \frac{\partial x}{\partial w} \frac{\partial e}{\partial u}.$$

By the properties of the expenditure function, we know that $\frac{\partial e}{\partial u} > 0$ (see MWG Prop 3.E.2, p. 59), so that $\frac{\partial h}{\partial u}$ has the same sign as $\frac{\partial x}{\partial w}$. Hence, when the good is normal, increasing u increases Hicksian demand for any price. Thus, increasing u shifts the Hicksian demand curve to the right. Similarly, when the good is inferior, increasing u decreases Hicksian demand for any price, and thus increasing u shifts the Hicksian demand for an inferior good to the left. The intuition is that in order to achieve a higher utility level, the consumer must spend more, and consumption increases with expenditure for a normal good and decreases with expenditure for an inferior good.

A couple of pictures. These pictures depict Walrasian and Hicksian demand before and after a price decrease for a normal good and for an inferior good. Note that $p^1 < p^0$ so that $u^1 > u^0$. For the normal good, Hicksian demand is steeper than Walrasian, and shifts to the right when the price decreases. For the inferior good, Hicksian demand is flatter than Walrasian and shifts to the left when the price decreases.

Substitutes and Complements Revisited Remember when we studied the UMP, we said that goods i and j were gross complements or substitutes depending on whether $\frac{\partial x_i}{\partial p_j}$ was negative or positive? Well, notice that we could also classify goods according to whether $\frac{\partial h_i}{\partial p_j}$ is negative or positive. In fact, we will call goods i and j complements if $\frac{\partial h_i}{\partial p_j} < 0$ and substitutes if $\frac{\partial h_i}{\partial p_j} > 0$. That




Figure 3.16:

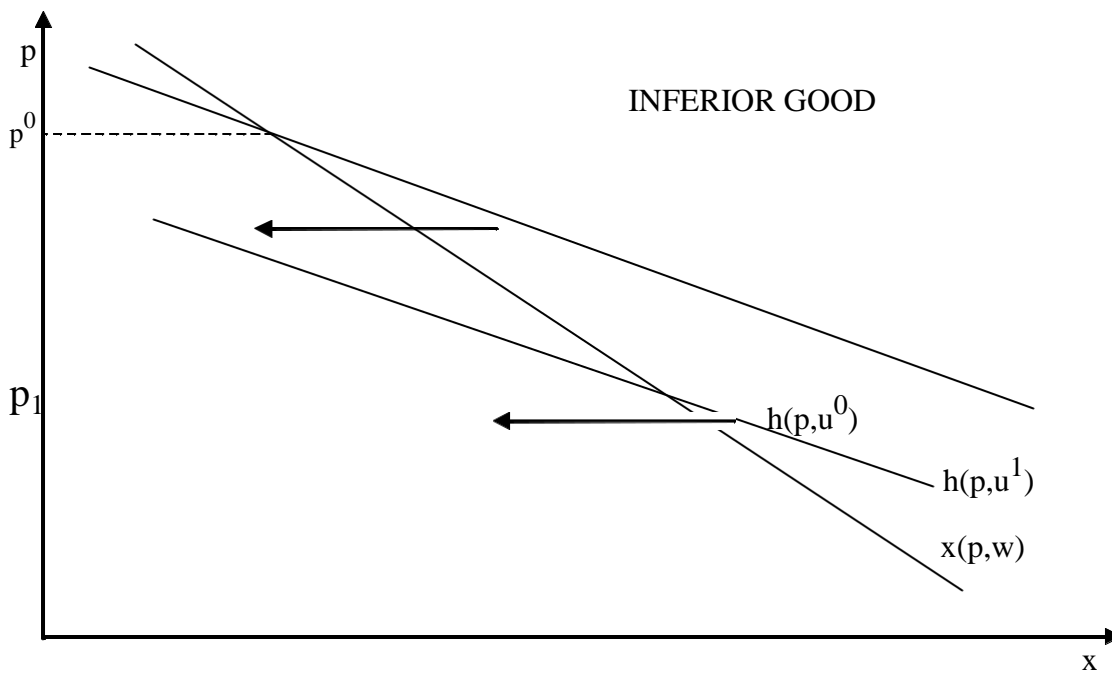


Figure 3.17:

is, we drop the “gross” when talking about the Hicksian demand function.²¹ In many ways, the Hicksian demand function is the proper function to use to talk about substitutes and complements since it separates the question of wealth effects and substitution effects. For example, it is possible that good j is a gross complement for good i while good i is a gross substitute for good j (if good i is normal and good j is inferior), but no such thing is possible when talking about (just plain) complements or substitutes since $\frac{\partial h_i}{\partial p_j} = \frac{\partial h_j}{\partial p_i}$.

3.4.6 The EMP and the UMP: Summary of Relationships

The relationships between all of the parts of the EMP and the UMP are summarized in Figure 3.G.3 of MWG and similar figures appear in almost any other micro theory book. So, I urge you to look it over (with the proviso about the typo that I mentioned earlier).

Here, I'll do it in words. Start with the UMP.

$$\begin{array}{ll} & \max u(x) \\ \text{s.t.} & : \\ & p \cdot x \leq w. \end{array}$$

The solution to this problem is $x(p, w)$, the Walrasian demand functions. Substituting $x(p, w)$ into $u(x)$ gives the indirect utility function $v(p, w) \equiv u(x(p, w))$. By differentiating $v(p, w)$ with respect to p_i and w , we get Roy's identity, $x_i(p, w) \equiv -\frac{v_{p_i}}{v_w}$.

Now the EMP.

$$\begin{array}{ll} & \min p \cdot x \\ \text{s.t.} & : \quad u(x) \geq u. \end{array}$$

The solution to this problem is the Hicksian demand function $h(p, u)$, and the expenditure function is defined as $e(p, u) \equiv p \cdot h(p, u)$. Differentiating the expenditure function with respect to p_j gets you back to the Hicksian demand, $h_j(p, u) \equiv \frac{\partial e(p, u)}{\partial p_j}$.

The connections between the two problems are provided by the duality results. Since the same bundle that solves the UMP when prices are p and wealth is w solves the EMP when prices are p and the target utility level is $v(p, w)$, we have that

$$\begin{aligned} x(p, w) &\equiv h(p, v(p, w)) \\ h(p, u) &\equiv x(p, e(p, u)). \end{aligned}$$

²¹ You can think of the ‘gross’ as referring to the fact that $\frac{\partial x_i}{\partial p_j}$ captures the effect of the price change before adding in the effect of compensation, sort of like how gross income is sales before adding in the effect of expenses.

Applying these identities to the expenditure and indirect utility functions yields more identities:

$$v(p, e(p, u)) \equiv u$$

$$e(p, v(p, w)) \equiv w.$$

Note: These last equations are where the mistake is in the book. Finally, from the relationship between $x(p, w)$ and $h(p, u)$ we can derive the Slutsky equation:

$$\frac{\partial h_i(p, v(p, w))}{\partial p_j} \equiv \frac{\partial x_i(p, w)}{\partial p_j} + \frac{\partial x_i(p, w)}{\partial w} x_j.$$

If you are really interested in such things, there is also a way to recover the utility function from the expenditure function (see a topic in MWG called “integrability”), but I’m not going to go into that here.

3.4.7 Welfare Evaluation

Underlying our approach to the study of preferences has been the ultimate goal of developing a tool for the welfare evaluation of policy changes. Recall that:

1. The UMP leads to $x(p, w)$ and $v(p, w)$, which are at least in principle observable. However, $v(p, w)$ is not a good tool for welfare analysis.
2. The EMP leads to $h(p, u)$ and $e(p, u)$, which are based on unobservables (u) but provide a good measure for the change in a consumer’s welfare following a policy change.
3. The Slutsky equation provides the link between the observable concepts, $x(p, w)$, and the useful concepts, $h(p, u)$.

In this section, we explore how these tools can be used for welfare analysis. The neoclassical preference-based approach to consumer theory gives us a measure of consumer well-being, both in terms of utility and in terms of the wealth needed to achieve a certain level of well-being. It turns out that this is crucial for welfare evaluation.

We will consider a consumer with “well-behaved” preferences (i.e. a strictly increasing, strictly quasiconcave utility function). The example we will focus on is the welfare impact of a price change.

Consider a consumer who has wealth w and faces initial prices p^0 . Utility at this point is given by

$$v(p^0, w).$$

If prices change to p^1 , the consumer's utility at the new prices is given by:

$$v_i^{\phi}(p^1, w).$$

Thus the consumer's utility increases, stays constant, or decreases depending on whether:

$$v_i^{\phi}(p^1, w) - v_i^{\phi}(p^0, w)$$

is positive, equal to zero, or negative.

While looking at the change in utility can tell you whether the consumer is better off or not, it cannot tell you how much better off the consumer is made. This is because utility is an ordinal concept. The units that utility is measured in are arbitrary. Thus it is meaningless to compare, for example, $v_i^{\phi}(p^1, w) - v_i^{\phi}(p^0, w)$ and $v(p_2, w) - v(p_3, w)$. And, if $v(\cdot)$ and $y(\cdot)$ are the indirect utility functions of two people, it is also meaningless to compare the change in v to the change in y .

However, suppose we were to compare, instead of the direct utility earned at a particular price-wealth pair, the wealth needed to achieve a certain level of utility at a given price-wealth pair. To see how this works, let

$$\begin{aligned} u^0 &= v_i^{\phi}(p^0, w^{\phi}) \\ u^1 &= v_i^{\phi}(p^1, w^{\phi}). \end{aligned}$$

We are interested in comparing the expenditure needed to achieve u^1 or u^0 . Of course, this will depend on the particular prices we use. It turns out that we have broad latitude to choose whichever set of prices we want, so let's call the reference price vector p^{ref} , and we'll assume that it is strictly greater than zero on all components.

The expenditure needed to achieve utility level u at prices p^{ref} is just

$$e(p^{\text{ref}}, u).$$

Thus, if we want to compare the expenditure needed to achieve utility u^0 and u^1 , this is given by:

$$\begin{aligned} e(p^{\text{ref}}, u^1) - e(p^{\text{ref}}, u^0) \\ e(p^{\text{ref}}, v_i^{\phi}(p^1, w^{\phi})) - e(p^{\text{ref}}, v_i^{\phi}(p^0, w^{\phi})). \end{aligned}$$

This expression will be positive whenever it takes more wealth to achieve utility u^1 at prices p^{ref} than to achieve u^0 . Hence this expression will also be positive, zero, or negative depending on

whether $u^1 > u^0$, $u^1 = u^0$, or $u^1 < u^0$. However, the units now have meaning. The difference is measured in dollar terms. Because of this, $e^{p^{ref}}$, $v(p, w)$ is often called a

indirect utility function. i $\$$ money metric

We can construct a money metric indirect utility function using virtually any strictly positive price as the reference price p^{ref} . However, there are two natural candidates: the original price, p^0 , and the new price, p^1 . When $p^{ref} = p^0$, the change in expenditure is equal to the change in wealth such that the consumer would be indifferent between the new price with the old wealth and the old price with the new wealth. Thus it asks what change in wealth would be equivalent to the

change in price. Formally, define the equivalent variation, $EV_{p^0, p^1, w^{\$}}$, as

$$EV_{p^0, p^1, w^{\$}} = e_{p^0, v_{p^1, w^{\$}}} - e_{p^0, v_{p^0, w^{\$}}} = e_{p^0, v_{p^1, w^{\$}}} - w.$$

that i i since $e_{p^0, v_{p^0, w}} = w$. Equivalent variation is illustrated in MWG Figure 3.1.2, panel a. Notice the compensation takes place at the old prices — the budget line shifts parallel to the one for p^0, w .

Since $w = e_{p^1, v_{p^1, w^{\$}}}$, an alternative definition of EV would be:

$$EV_{p^0, p^1, w^{\$}} = e_{p^0, v_{p^1, w^{\$}}} - e_{p^1, v_{p^1, w^{\$}}}.$$

In this form, EV asks how much more money does it take to achieve utility level $v_{p^1, w^{\$}}$ at p^0 than at p^1 . Note: if $EV < 0$, this means that it takes less money to achieve utility $v_{p^1, w^{\$}}$ at p^0

than p^1 (which means that prices have gone up to get to p^1 , at least on average). i $\$$

When considering the case where the price of only one good changes, EV has a useful interpretation in terms of the Hicksian demand curve. Applying the fundamental theorem of calculus and

the fact that $\frac{\partial e(p, u)}{\partial p_i} = h_i(p, u)$, if only the price of good 1 changes, we have:²²

$$e_{p^0, v_{p^1, w}} - e_{p^1, v_{p^1, w}} = \int_{p^0_1}^{p^1_1} h_1(s, v_{p^1, w}) ds$$

Thus the absolute value of EV is given by the area to the left of the Hicksian demand curve between p^0_1 and p^1_1 . If $p^0_1 < p^1_1$, EV is negative - a welfare loss because prices went up. If $p^0_1 > p^1_1$, EV is

²² Often when we are interested in a particular component of a vector - say, the price of good i - we will write the vector as (p_i, p_{-i}) , where p_{-i} consists of all the other components of the price vector. Thus, (p^*_i, p_{-i}) stands for the vector $(p_1, p_2, \dots, p_{i-1}, p^*_i, p_{i+1}, \dots, p_L)$. It's just a shorthand notation.

Another notational explanation - in an expression such as p^0_1 , the superscript refers to the timing of the price vector (i.e. new or old prices), and the subscript refers to the commodity. Thus, p^0_1 is the old price of good 1.

positive - a welfare gain because prices went down. The relevant area is depicted in MWG Figure 3.1.3, panel a.

The other case to consider is the one where the new price is taken as the reference price. When $p^{ref} = p^1$, the change in expenditure is equal to the change in wealth such that the consumer is indifferent between the original situation (p^0, w) and the new situation $(p^1, w + \Delta w)$. Thus it asks

how much wealth would be needed to compensate the consumer for the price change. Formally, define the compensating variation (depicted in MWG Figure 3.1.2, panel b)

$$CV(p^0, p^1, w) = e(p^1, v(p^1, w)) - e(p^1, v(p^0, w)) = w - e(p^1, v(p^0, w))$$

Again, when only one price changes, we can readily interpret CV in terms of the area to the left of a Hicksian demand curve. However, this time it is the Hicksian demand curve for the old utility level, u^0 . To see why, note that $w = e(p^0, v(p^0, w))$

of good 1 changes): $e(p^0, v(p^0, w)) = \int_{p_1^0}^{p_1^1} h_1(p_1, v(p_1, w)) dp_1$, and so (again assuming only the price

$$CV(p^0, p^1, w) = e(p^0, v(p^0, w)) - e(p^1, v(p^0, w)) = \int_{p_1^0}^{p_1^1} h_1(p_1, v(p_1, w)) dp_1$$

which is positive whenever $p_1^0 > p_1^1$ and negative whenever $p_1^0 < p_1^1$. The relevant area is illustrated in MWG Figure 3.1.3, panel b.

Recall that whenever good i is a normal good, increasing the target utility level u shifts

$h_i(p_i, p^{-i}, u)$ to the right in the (x_i, p_i) space. This is because in order to achieve higher utility the consumer will need to spend more wealth, and if the good is normal and the consumer spends more wealth, more of the good will be consumed. Thus when the good is normal, $EV \geq CV$.

On the other hand, if the good is inferior, then increasing u shifts $h_i(p_i, p^{-i}, u)$ to the left, and $CV \geq EV$. When there is no wealth effect on the good, i.e., $\frac{\partial x_i(p, w)}{\partial w} = 0$, then $CV = EV$.

Figure 3.1.3 also shows the Walrasian demand curve. In fact, it shows it crossing $h_1(p_1, p^{-1}, v(p_1, w))$ at p_1^0 and $h_1(p_1, p^{-1}, v(p_1, w))$ at p_1^1 . This results from the duality of utility

$$e(p, w) = \int_{p_1}^{p_1^0} h_1(p_1, v(p_1, w)) dp_1 = x_1(p, w)$$

$$e(p, w) = \int_{p_1}^{p_1^1} h_1(p_1, v(p_1, w)) dp_1 = x_1(p, w)$$

which each arise from the identity $h_i(p, v(p, w)) \equiv x_i(p, w)$. The result of this is that the Walrasian demand curve crosses the Hicksian demand curves at the two points mentioned above, and that the area to the left of the Walrasian demand curve lies somewhere between the EV and CV. There are a number of comments that must be made on this topic:

1. Although the area to the left of the Hicksian demand curve is equal to the change in the expenditure function, the area to the left of the Walrasian demand function has no ready interpretation.
2. The area to the left of the Walrasian demand curve is called the change in Marshallian consumer surplus, ΔCS , and is probably the notion of welfare change that you are used to from your intermediate micro courses.
3. Unfortunately, the change in Marshallian consumer surplus is a meaningless measure (see part 1) except for:
 - (a) If there are no wealth effects on the good whose price changes, then $EV = CV = \Delta CS$.
 - (b) Since ΔCS lies between EV and CV , it can sometimes be a good approximation of the welfare impact of a price change. This is especially true if wealth effects are small.
4. Some might argue that ΔCS is a useful concept because it is easier to compute than EV or CV since it does not require estimation of the Hicksian demand curves. But, if you know about the Slutsky equation (which you do), this isn't such a problem.

So Which is Better, EV or CV ?

Both EV and CV provide dollar measures of the impact of a price change on consumer welfare, and there are circumstances in which each is the appropriate measure to use. EV does have one advantage over CV , though, and that is that if you want to consider two alternative price changes, EV gives you a meaningful measure, while CV does not (necessarily). For example, consider

initial price p^U and two alternative price vectors p^a and p^D . The quantities $EV(p^U, p^a, w)$ and $CV(p^U, p^D, w)$ are both measured in terms of wealth at prices p^U and thus they can be compared. On the other hand, $CV(p^a, p^D, w)$ is in terms of wealth at prices p^a and $CV(p^U, p^D, w)$ is in terms of prices p^D , which cannot be readily compared.

of wealth needed at p^U and p^D . This distinction is important in policy issues such as deciding which commodity to tax. The impact of placing a tax on gasoline vs. the impact of placing a tax on electricity needs to be measured with respect to the same reference price if we want to compare the two in a meaningful way. This means using EV .

Example: Deadweight Loss of Taxation.

Suppose that the government is considering putting a tax of $t > 0$ dollars on commodity 1. The

current price vector is p^0 . Thus the new price vector is $p^1 = p^0 + t, p_2^0, \dots, p_L^0$.

After the tax is imposed, consumers purchase h_i^1, u_i^1 units of the good, where $u_i^1 = v_i(p^1, w)$.
 The tax revenue raised by the government is $T = t \sum_i h_i^1$. However, in order to raise

this T dollars, the government must increase the effective price p_i^1 of good 1. This makes consumers worse off, and the amount by which it makes consumers worse off is given by:

$$EV(p^0, p^1, w) = \int_{p^1}^{p^0} h_1^i(s, p^0, u^1) ds.$$

Since $p_1^1 > p_1^0$, EV is negative and gives the amount of money that consumers would be willing to pay in order to avoid the tax. Thus consumers are made worse off by $EV(p^0, p^1, w)$ dollars.

Since the tax raises T dollars, the net impact of the tax is $-EV(p^0, p^1, w) - T$.

$$-EV(p^0, p^1, w) - T.$$

The previous expression, known as the deadweight loss (DWL) of taxation, gives the amount by which consumers would have been better off, measured in dollar terms, if the government had just taken T dollars away from them instead of imposing a tax. To put it another way, consumers see the tax as equivalent to losing EV dollars of income. Since the tax only raises T dollars of income, $-EV - T$ is the dollar value of the consumers' loss that is not transferred to the government as tax revenue. It simply disappears.

Well, it doesn't really disappear. Consumers get utility from consuming the good. In response to the tax, consumers decrease their consumption of the good, and this decreases their utility and is the source of the deadweight loss. On the other hand, a tax that does not distort the price consumers must pay for the good would not change their compensated demand for the good. Consequently, it would not lead to a deadweight loss. This is one argument for lump-sum taxes instead of per-unit taxes. Lump-sum taxes (each consumer pays T dollars, regardless of the consumption bundle each one purchases) do not distort consumers' purchases, and so they do not lead to deadweight losses. However, lump-sum taxes have problems of their own. First, they are regressive, meaning that they impact the poor more than the rich, since everybody must pay the same amount. Second, lump-sum taxes do not charge the users of commodities directly. So, there is some question whether, for example, money to pay for building and maintaining roads should be raised by charging everybody the same amount or by charging a gasoline tax or by charging drivers

a toll each time they use the road. The lump-sum tax is non-distortionary, but it must be paid by people who don't drive, even people who can't afford to drive. The gasoline tax is paid by all drivers, including people who don't use the particular roads being repaired, and it is distortionary in the sense that people will generally reduce their driving in response to the tax, which induces a deadweight loss. Charging a toll to those who use the road places the burden of paying for repairs on exactly those who are benefiting from having the roads. But like the gasoline tax, it is also distortionary (since people will tend to avoid toll roads). And, since the tolls are focused on relatively few consumers, the tolls may have to be quite high in order to raise the necessary funds, imposing a large burden on those people who cannot avoid using the toll roads. These are just some of the issues that must be considered in deciding which commodities should be taxed and how.

3.4.8 Bringing It All Together

Recall the basic dilemma we faced. The UMP yields solution $x(p, w)$ and value function $v(p, w)$ that are based on observables but not useful for doing welfare evaluation since utility is ordinal. The EMP yields solution $h(p, u)$ and value function $e(p, u)$, which can be used for welfare evaluation but are based on u , which is unobservable. As I have said, the link between the two is provided by the Slutsky equation

$$\frac{\partial h_i(p, v(p, w))}{\partial p_j} = \frac{\partial x_i(p, w)}{\partial p_j} + \frac{\partial x_i(p, w)}{\partial w} x_j(p, w).$$

We now illustrate how this is implemented. Suppose the price of good 1 changes. EV is given by:

$$EV(p^0, p^1, w) = \int_{p^0_1}^{p^1_1} h_1(s, p^0_{-1}, u^1) ds.$$

We can approximate $h_1(s, p^0_{-1}, u^1)$ using a first-order Taylor approximation.

Recall, a first-order Taylor approximation to $f(x)$ is given by:

$$f(x) \approx f(x_0) + f'(x_0)(x - x_0).$$

This gives a linear approximation to $f(x)$ that is tangent to $f(x)$ at x_0 and a good approximation for x that are not too far from x_0 . But, the further x is away from x_0 , the worse the approximation will be. See Figure 3.18.

Now, the first-order Taylor approximation to $h_1(s, p^0_{-1}, u^1)$ is given by:

$$h_1(s, p^0_{-1}, u^1) \approx h_1(p^0_1, p^0_{-1}, u^1) + \frac{\partial h_1(p, v(p, w))}{\partial p_1} (s - p^0_1) \tag{3.6}$$

$$h_1(s, p^0_{-1}, u^1) \approx h_1(p^0_1, p^0_{-1}, u^1) + \frac{\partial h_1(p, v(p, w))}{\partial p_1} (s - p^0_1)$$




Figure 3.18: A First-Order Taylor Approximation

Note that we have taken as our original point $p = p^1_1, p^0_{-1}$. That is, the price vector after the price change? Why do we do this? The reason is that we are using the Hicksian demand curve for u^1 , the utility level after the price change. Because of that, we also want to use the price after the price change. We know that, at $p = p^1_1, p^U_{-1}, h^1_1, p^U_{-1}, u^1 = x^1_1, p^1_1, p^U_{-1}, w$. This fact, along

with the Slutsky equation, allows us

$$h^1_1(p^1_1, p^0_{-1}, u^1) = x^1_1(p^1_1, p^0_{-1}, w) + \tilde{A} \frac{\partial x^1_1(p^1_1, p^0_{-1}, w)}{\partial p^1_1} + \frac{\partial x^1_1(p^1_1, p^0_{-1}, w)}{\partial w} \quad (3.7)$$

The last equation provides an approximation for the Hicksian demand curve based only on observable quantities. That is, we have eliminated the need to know the (unobservable) target utility level. Finally, note that demand $x^1_1(p^1_1, p^0_{-1}, w)$ and derivatives $\frac{\partial x^1_1(p^1_1, p^0_{-1}, w)}{\partial p^1_1}$ and $\frac{\partial x^1_1(p^1_1, p^0_{-1}, w)}{\partial w}$ can

be observed or approximated using $x^1_1(p^1_1, p^0_{-1}, w)$ approximation and one based on the Walrasian demand curve is the addition of the wealth-effect term, $\frac{\partial x^1_1(p^1_1, p^0_{-1}, w)}{\partial w}$.

Figure 3.19

in our estimate to the Hicksian demand function is $p = p^1_1, p^0_{-1}$, estimated Hicksian demand h^1_1 is coincides with and it tangent to the actual Hicksian h^1_1 prices that are further away from p^1_1 , the approximation is less good. True EV is the area left of h^1_1 . Thus, the estimation "error", the difference between true EV and estimated EV, is given by the area between h^1_1 and h^1_1 between prices p^0_1 and p^1_1 .

In the case of CV, CV is computed as the area left of the Hicksian demand curve at the original

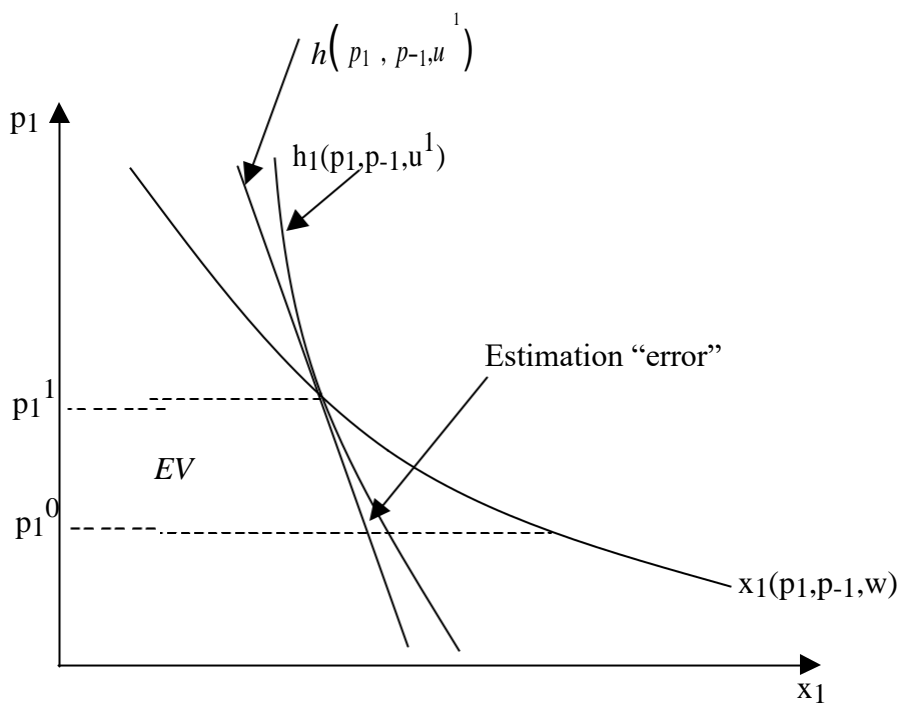


Figure 3.19: The First-Order Taylor Approximation to EV




Figure 3.20: The First-Order Taylor Approximation to CV

utility level, $h_1(p_1, p_{-1}, u^0)$. Because of this, we must use the original price as the “original point” in the Taylor approximation. Thus, for the purposes of CV, estimated Hicksian demand is given by:

$$h_1(p_1^1, p_{-1}^0, u^1) \approx h_1(p_1^0, p_{-1}^0, u^1) + \frac{\partial h_1(p_1, p_{-1}, u^1)}{\partial p_1} (p_1^1 - p_1^0) + \frac{\partial x_1(p_1, p_{-1}, w)}{\partial p_1} (p_1^1 - p_1^0) + \frac{\partial x_1(p_1, p_{-1}, w)}{\partial w} (w^1 - w^0)$$

The diagram for the Taylor approximation to CV corresponding to Figure 3.19 therefore looks like Figure 3.20:

Welfare Evaluation: An Example

Data: $x_1^0 = 100, p_1^0 = 10, \frac{\partial x_1}{\partial p_1} = -4, \frac{\partial x_1}{\partial w} = 0.02$. Note: there is no information on w .

Suppose the price of good 1 increases to $p_1^1 = 12.5$. How much should a public assistance program aimed at maintaining a certain standard of living be increased to offset this price increase?

To answer this question, we are looking for the CV of the price change. To compute this, we

need to approximate the Hicksian demand curve for the original utility level, $h_1(p_1, p_{-1}, u^0)$.




Figure 3.21:

1. We know that $h_{i, p-1, p-1, u^0} = x_i(10, p-1, w)$.
2. The slope of the $h_{i, p-1, p-1, u^0}$ can be approximated using the data and the Slutsky equation.

$$\begin{aligned} \frac{\partial h}{\partial p} &\equiv -\frac{\partial x}{\partial p} + \frac{\partial x}{\partial w} \frac{\partial w}{\partial p} \\ &= -4 + 0.02(100) \\ &= -2 \end{aligned}$$

3. So, at price 12.5, Hicksian demand is given by

$$\begin{aligned} h &= 100 + \frac{\partial h}{\partial p} dp \\ &= 100 + (-2) 2.5 = 95 \end{aligned}$$

4. To compute CV, compute the area of a trapezoid (or the area of a rectangle plus a triangle):

$$|CV| = (2.5) \frac{95 + 100}{2} = 243.75.$$

Since the price is increasing, we know that $CV < 0$, so $CV = -243.75$.

We could also estimate the change in Marshallian Consumer Surplus. This is just the area to the left of the Walrasian demand curve between the two prices. Hence $\Delta CS = - (2.5) \frac{90+100}{2} =$

-237.5 . Hence if we were to use the Marshallian consumer surplus in this case, we would not compensate the consumer enough for the price increase.

Another thing we could do is figure that the harm done to the consumer is just the change in price times the original consumption of this good, i.e., $2.5(100) = 250$. However, if we gave the consumer 250 additional dollars, we would be overcompensating for the price increase.

3.4.9 Welfare Evaluation for an Arbitrary Price Change

The basic analysis of welfare change using CV and EV considers the case of a single price change. However, what should we do if the policy change is not a single price change? For changes in multiple prices, we can just compute the CV for each of the changes (i.e., changing prices one by one and adding the CV (or EV) from each of the changes along a “path” from the original price to the new price). If price and wealth change, we can add the change in wealth to the CV (or EV) from the price changes (see below). But, what if the policy change involves something other than prices and wealth, such a change in environmental quality, roads, etc. How do we value such a change?

The answer is that, if we have good estimates of Walrasian demand, we can always represent the change as a change in a budget set. After doing so, we can compute the CV in the usual way.

Part 1: Any arbitrary policy change can be thought of as a simultaneous change in p and w .

To illustrate, suppose that we have a good estimate of consumers’ demand functions (i.e., we fit a flexible functional form for demand using high-quality data). Let $x(p, w)$ denote demand. Suppose that initially prices and wealth are p^0, w^0 and the consumer chooses bundle $x^0(p^0, w^0)$. Now, suppose that “something happens” that leads the consumer to choose bundle x instead of x^0 . What is the CV (or EV) of this change?

The first step is to note that, if demand is quasiconcave, there is some price-wealth vector for which x^0 and x are optimal choices. You can find these price-wealth vectors, which we’ll call (p^u, w^u) and (p^x, w^x) , by solving the equations $x^0 = x(p^u, w^u)$ and $x = x(p^x, w^x)$. (In reality you probably already know p^0, w^0 and have an observation of x or estimate of it.) Remember, we have a good estimate of $x(p, w)$. Once we find (p^u, w^u) and (p^x, w^x) , then we know that the change in the consumer’s utility in going from x^0 to x is just $v(p^x, w^x) - v(p^u, w^u)$, and so the impact of the policy change reduces to computing the EV or CV for this change and $v(p^x, w^x) - v(p^u, w^u) = u^x - u^u$.

Part 2: Compute the EV or CV for a simultaneous change in p and w .

So, we’ve recast the policy change as a change from (p^0, w^0) to (p^x, w^x) , letting u^x and u^u denote

the utility levels before and after the change. To compute EV, return to the definition of EV used before.

$$EV = e(p^0, u^0) - e(p^1, u^0)$$

Adding and subtracting $e(p^0, u^0)$, we get:

$$EV = e(p^0, u^0) - e(p^0, u^0) + e(p^0, u^0) - e(p^1, u^0)$$

But, note that $e(p^0, u^0) = w^0$ and $e(p^0, u^0) = w^0$, so

$$EV = e(p^0, u^0) - e(p^1, u^0) + w^0 - w^0 \tag{*}$$

and note that $e(p^0, u^0) - e(p^1, u^0)$ is as in the definition of EV when only a price changes. So, if only the price p_i changes, EV can be written as

$$EV = \int_{p_i^0}^{p_i^1} h_i(p_1, \dots, p_{i-1}, p_{i+1}, \dots, p_L, u^0) dp_i + w^0 - w^0$$

and this can be estimated in the usual way from the estimated Walrasian demand curve.

If multiple prices change, we change them one by one and add up the integral from each change, and then we add the change in wealth. That is, if prices change from p_0, p_1, \dots, p_L to (p_1, p_2, \dots, p_L)

and wealth changes from w^0 to w^1 , the EV is:

$$EV = \int_{p_1^0}^{p_1^1} h_1(p_1, p_2^0, \dots, p_L^0) dp_1 + \int_{p_2^0}^{p_2^1} h_2(p_1^0, p_2, p_3^0, \dots, p_L^0) dp_2 + \dots + \int_{p_L^0}^{p_L^1} h_L(p_1^0, p_2^0, \dots, p_{L-1}^0, p_L) dp_L + w^1 - w^0$$

If you replace each Hicksian demand with an estimate based on Marshallian demand and the

Slutsky equation, you can estimate this using only observables. It is tedious, but certainly possible.

This is a diagram that illustrates the whole thing. Suppose a policy change moves the consumer's consumption bundle from x^0 to x^1 . To compute the EV, the first thing you do is find the (p, w) for which $x^0 = x(p^0, w^0)$. This budget set is labeled $B(p^0, w^0)$. Then, you find the (p, w) for which x^1 is optimal, which we call (p^1, w^1) . This budget set (red) is labeled $B(p^1, w^1)$. Denote the initial utility level u^0 and the final utility level u^1 , and note that neither the utility levels nor the indifference curves (which are drawn in as dotted lines for illustration) are observed.

Next, we decompose the change from (p^0, w^0) to (p^1, w^1) into two parts. Part 1 is a change in wealth holding prices fixed at p^0 . Let y^0 denote the point the consumer chooses at (p^0, w^0) , and let u^0 denote the utility earned. This point and the associated budget set are in blue. Note that moving from budget set $B(p^0, w^0)$ to budget set $B(p^0, w^1)$ is just like losing $w^0 - w^1$ dollars (since in fact, exactly what happens). This is where the $w^0 - w^1$ term comes

prices don't change this is, y^0

from in expression (*) above. Distance $w^0 - w^0$ is denoted on the left. (Note that in the diagram, these distances are scaled by p_2 , since we are showing them on the x_2 -axis.)

Part 2 of the decomposition is the change in prices from p^0 to p^0 when wealth is w^0 . But, note that this just the kind of EV we computed in the simple case. That is, prices change but wealth remains constant. Let z denote the point that offers the same utility as x^0 but is chosen at prices

p^0 . That is, $z = x^0 + EV(p^0, p^0, w^0)$. The budget line supporting z is denoted in green, and the EV for the

from blue $B(p^0, w^0)$ to green $B(p^0, w^0 + EV(p^0, p^0, w^0))$ denoted $EV(p^0, p^0, w^0)$ on the left. Since

$$\begin{aligned} EV(p^0, p^0, w^0) &= e(p^0, u^0) - e(p^0, u^0) \\ &= e(p^0, u^0) - w^0 \end{aligned}$$

$$= e(p^0, u^0) - e(p^0, u^0),$$

this is just like the EV's we computed when only prices changed. This is where the $e(p^0, u^0) - e(p, u)$

comes from in expression (*) above.

The total EV is the sum of these two parts. The distance is denoted Total EV in the diagram.

Note that since the consumer ends up worse off overall, the total EV should be negative.

You could also do something similar for CV.

$$\begin{aligned} CV &= e(p^0, u^0) - e(p^0, u^0) \\ &= e(p^0, u^0) - e(p^0, u^0) + e(p^0, u^0) - e(p^0, u^0) \\ &= w^0 - w^0 + e(p^0, u^0) - e(p^0, u^0), \end{aligned}$$

and note that once again $e(p^0, u^0) - e(p^0, u^0)$ is as in our original definition of CV. So, this

term can be rewritten in $e(p^0, u^0) - e(p^0, u^0)$




Figure 3.22:

Chapter 4

Topics in Consumer Theory

4.1 Homothetic and Quasilinear Utility Functions

One of the chief activities of economics is to try to recover a consumer's preferences over all bundles from observations of preferences over a few bundles. If you could ask the consumer an infinite number of times, "Do you prefer x to y ?", using a large number of different bundles, you could do a pretty good job of figuring out the consumer's indifference sets, which reveals her preferences. However, the problem with this is that it is impossible to ask the question an infinite number of times.¹ In doing economics, this problem manifests itself in the fact that you often only have a limited number of data points describing consumer behavior.

One way that we could help make the data we have go farther would be if observations we made about one particular indifference curve could help us understand all indifference curves. There are a couple of different restrictions we can impose on preferences that allow us to do this.

The first restriction is called homotheticity. A preference relation is said to be homothetic if the slope of indifference curves remains constant along any ray from the origin. Figure 4.1 depicts such indifference curves.

If preferences take this form, then knowing the shape of one indifference curve tells you the shape of all indifference curves, since they are "radial blowups" of each other. Formally, we say a preference relation is homothetic if for any two bundles x and y such that $x \sim y$, then $\alpha x \sim \alpha y$ for any $\alpha > 0$.

We can extend the definition of homothetic preferences to utility functions. A continuous

¹In fact, to completely determine the indifference sets you would have to ask an uncountably infinite number of questions, which is even harder.




Figure 4.1: Homothetic Preferences

preference relation \succsim is homothetic if and only if it can be represented by a utility function that is homogeneous of degree one. In other words, homothetic preferences can be represented by a function $u(\cdot)$ such that $u(\alpha x) = \alpha u(x)$ for all x and $\alpha > 0$. Note that the definition does not say that every utility function that represents the preferences must be homogeneous of degree one — only that there must be at least one utility function that represents those preferences and is homogeneous of degree one.

EXAMPLE: Cobb-Douglas Utility: A famous example of a homothetic utility function is the Cobb-Douglas utility function (here in two dimensions):

$$u(x_1, x_2) = x_1^a x_2^{1-a} : a > 0.$$

The demand functions for this utility function are given by:

$$x_1(p, w) = \frac{aw}{p_1}$$

$$x_2(p, w) = \frac{(1-a)w}{p_2}$$

Notice that the ratio of x_1 to x_2 does not depend on w . This implies that Engle curves (wealth expansion paths) are straight lines (see MWG pp. 24-25). The indirect utility function is given by:

$$v(p, w) = \frac{aw}{p_1}^a \left(\frac{(1-a)w}{p_2}\right)^{1-a} = W \frac{a^a (1-a)^{1-a}}{p_1^a p_2^{1-a}}$$

Another restriction on preferences that can allow us to draw inferences about all indifference curves from a single curve is called quasilinearity. A preference relation is quasilinear if there is one commodity, called the numeraire, which shifts the indifference curves outward as consumption

of it increases, without changing their slope. Indifference curves for quasilinear preferences are illustrated in Figure 3.B.6 of MWG.

Again, we can extend this definition to utility functions. A continuous preference relation is quasilinear in commodity 1 if there is a utility function that represents it in the form $u(x) = x_1 + v(x_2, \dots, x_L)$.

EXAMPLE: Quasilinear utility functions take the form $u(x) = x_1 + v(x_2, \dots, x_L)$. Since we typically want utility to be quasiconcave, the function $v(\cdot)$ is usually a concave function such as $\log x$ or \sqrt{x} . So, consider:

$$u(x) = x_1 + \sqrt{x_2}.$$

The demand functions associated with this utility function are found by solving:

$$\max x_1 + x_2^{0.5}$$

$$\text{s.t. : } p \cdot x \leq w$$

or, since $x_1 = \frac{w}{p_1} - x_2 \frac{p_2}{p_1}$,

$$\max \frac{w}{p_1} - x_2 \frac{p_2}{p_1} + x_2^{0.5}.$$

The associated demand curves are

$$x_1(p, w) = \frac{w}{p_1} - \frac{1}{4} \frac{p_1}{p_2}$$

$$x_2(p, w) = \frac{1}{4} \frac{p_1}{p_2}$$

and indirect utility function:

$$v(p, w) = \frac{1}{4} \frac{p_1}{p_2} + \frac{w}{p_1}.$$

Isoquants of this utility function are drawn in Figure 4.2.

4.2 Aggregation

Our previous work has been concerned with developing the testable implications of the theory of the consumer behavior on the individual level. However, in any particular market there are large numbers of consumers. In addition, often in empirical work it will be difficult or impossible to collect data on the individual level. All that can be observed are aggregates: aggregate consumption of the various commodities and a measure of aggregate wealth (such as GNP). This raises the




Figure 4.2: Quasilinear Preferences

natural question of whether or not the implications of individual demand theory also apply to aggregate demand.

To make things a little more concrete, suppose there are N consumers numbered 1 through N , and the n^{th} consumer's demand for good i is given by $x_i^n(p, w^n)$, where w^n is consumer n 's initial wealth. In this case, total demand for good i can be written as:

$$D_i(p, w, \dots, w^N) = \sum_{n=1}^N x_i^n(p, w^n).$$

However, notice that $D_i(\cdot)$ gives total demand for good i as a function of prices and the wealth levels of the n consumers. As I said earlier, often we will not have access to information about individuals, only aggregates. Thus we may ask the question of when there exists a function $D_i(p, w)$, where $w = \sum_{n=1}^N w^n$ is aggregate wealth, that represents the same behavior as $D_i(p, w^1, \dots, w^N)$. A second question is when, given that there exists an aggregate demand function $D_i(p, w)$, the behavior it characterizes is rational. We ask this question in two ways: First, when will the behavior resulting from $D_i(p, w)$ satisfy WARP? Second, when will it be as if $D_i(p, w)$ were generated by a "representative consumer" who is herself maximizing preferences? Finally, we will ask if there is a representative consumer, in what sense is the well-being of the representative consumer a measure of the well-being of society?

4.2.1 The Gorman Form

The major theme that runs through our discussion in this section is that in order for demand to aggregate, each individual's utility function must have an indirect utility function of the Gorman

Form. So, let me take a moment to introduce the terminology before we need it. An indirect utility function for consumer n , $v^n(p, w)$, is said to be of the Gorman Form if it can be written in terms of functions $a^n(p)$, which may depend on the specific consumer, and $b(p)$, which does not depend on the specific consumer:

$$v^n(p, w) = a^n(p) + b(p) w^n.$$

That is, an indirect utility function of the Gorman form can be separated into a term that depends on prices and the consumer's identity but not on her wealth, and a term that depends on a function of prices that is common to all consumers that is multiplied by that consumer's wealth.

The special nature of indirect utility functions of the Gorman Form is made apparent by applying Roy's identity:

$$x_i^n(p, w^n) = - \frac{\frac{\partial v^n}{\partial p_i}}{\frac{\partial v^n}{\partial w^n}} = - \frac{a_i^n(p) + \frac{\partial b(p)}{\partial p_i} w^n}{b(p)} \tag{4.1}$$

Now consider the derivative of a particular consumer's demand for commodity i with respect to wealth w^n . From now on, we will let $b_i(p)$ = demand for commodity i :

$$\frac{\partial x_i^n(p, w^n)}{\partial w^n} = \frac{b_i(p)}{b(p)}$$

This implies that wealth-expansion paths are given

by:

$$\frac{\frac{\partial x_i^n(p, w^n)}{\partial w^n}}{\frac{\partial x_j^n(p, w^n)}{\partial w^n}} = \frac{b_i(p)}{b_j(p)}$$

not depend on wealth. Thus, as wealth increases, each consumer increases her consumption of the goods at a linear rate. The result is that each consumer's wealth-expansion paths are straight lines. Second, $\frac{\partial x_i^n(p, w^n)}{\partial w^n}$ is the same for all consumers, since $\frac{b_i(p)}{b(p)}$ does not depend on n . This implies that the wealth-expansion paths for different consumers are parallel (see MWG Figure 4.B.1).

Next, let's aggregate the demand functions of consumers with Gorman form indirect utility functions. Sum the individual demand functions from (4.1) across all n to get aggregate demand:

$$D_i(p, w^1, \dots, w^n) = \sum_n \frac{-a_i^n(p) - b_i(p) w^n}{b(p)} = \sum_n \frac{-a_i^n(p)}{b(p)} - \frac{b_i(p)}{b(p)} \sum_n w^n$$

$$= \sum_n \frac{-a_i^n(p)}{b(p)} - \frac{b_i(p)}{b(p)} W^{\text{total}}$$

Thus if all consumers have utility functions of the Gorman form, demand can be written solely as a function of prices and total wealth. In fact, this is a necessary and sufficient condition: Demand can be written as a function of prices and total wealth if and only if all consumers have indirect utility functions of the Gorman form (see MWG Proposition 4.B.1).

As a final note on the Gorman form, recall the examples of quasilinear and homothetic utility we did earlier. It is straightforward to verify (at least in the examples) that if all consumers have identical homothetic preferences or if consumers have (not necessarily identical) preferences that are quasilinear with respect to the same good, then their preferences will be representable by utility functions of the Gorman form.

4.2.2 Aggregate Demand and Aggregate Wealth

I find the notation in the book in this section somewhat confusing. So, I will stick with the notation used above. Let $x_i^n(p, w^n)$ be the demand by consumer n for good i when prices are p and wealth is w^n , and let $D_i(p, w^1, \dots, w^N)$ denote aggregate demand as a function of the entire vector of wealths.²

The general question we are asking here is whether or not the distribution of wealth among the consumers matters. If the distribution of wealth affects total demand for the various commodities, then we will be unable to write total demand as a function of prices and total wealth. On the other hand, if total demand does not depend on the distribution of wealth, we will be able to do so.

Let prices be given by p and the initial wealth for each consumer be given by w^{-n} . Let dw be a vector of wealth changes where dw^n represents the change in consumer n 's wealth and $\sum_{n=1}^N dw^n = 0$.

Thus dw represents a redistribution of wealth among the n consumers. If total demand can be written as a function of total wealth and prices, then

$$\frac{\partial x_i^n(p, w^{-n})}{\partial w^n} = 0$$

for all i . If this is going to be true for all initial wealth distributions w^{-1}, \dots, w^{-N} and all possible rearrangements dw , it must be the case that partial derivative of demand with respect to wealth is equal for every consumer and every distribution of wealth:

$$\frac{\partial x_i^n(p, w^n)}{\partial w^n} = \frac{\partial x_i^m(p, w^m)}{\partial w^m}$$

But, this condition is exactly the condition that at any price vector p , and for any initial distribution of wealth, the wealth effects of all consumers are the same. Obviously, if this is true then the changes

² One should be careful not to confuse the superscript with an exponent here. We are concerned with the question of when aggregate demand can be written as $D_i(p, \sum_{n=1}^N w^n)$, a function of prices and the total wealth of all consumers.

in demand as wealth is shifted from one consumer to another will cancel out. In other words, only total wealth (and not the distribution of wealth) will matter in determining total demand. And, this is equivalent to the requirement that for a fixed price each consumer's wealth expansion path

is a straight line (since $\frac{\partial x^n(p, w^n)}{\partial w^n}$ and $\frac{\partial x_j^n(p, w^n)}{\partial w^n}$ must be independent of w^n) and that the slope of the straight line must be the same for all consumers (since $\frac{\partial x^n(p, w^n)}{\partial w^n} = \frac{\partial x^m(p, w^m)}{\partial w^m}$).

And, as shown in the previous section, this property holds if and only if consumers' indirect utility functions take the Gorman form. Hence if we allow wealth to take any possible initial distribution, aggregate demand depends solely on prices and total wealth if and only if consumers' indirect utility functions take the Gorman form.

To the extent that we prefer to look at aggregate demand or are unable to look at individual demand (because of data problems), the previous result is problematic. There are a whole lot of utility functions that don't take the Gorman form. There a number of approaches that can be taken:

1. We can weaken the requirement that aggregate demand depend only on total wealth. For example, if we allow aggregate demand to depend on the empirical distribution of wealth but not on the identity of the individuals who have the wealth, then demand can be aggregated whenever all consumers have the same utility function.
2. We required that aggregate demand be written as a function of prices and total wealth for any distribution of initial wealth. However, in reality we will be able to put limits on what the distributions of initial wealth look like. It may then be possible to write aggregate demand as a function of prices and aggregate wealth when we restrict the initial wealth distribution. One situation in which it will always be possible to write demand as a function of total wealth and prices is when there is a rule that tells you, given prices and total wealth, what the wealth of each individual should be. That is, if for every consumer n , there exists a function $w^n(p, w)$ that maps prices p and total wealth w to individual wealth w^n . Such a rule would exist if individual wealth were determined by government policies that depend only on p and w . We call this kind of function a wealth distribution rule.
 - (a) An important implication of the previous point is that it always makes sense to think of aggregate demand when the vector of individual wealths is held fixed. Thus if we are only interested in the effects of price changes, it makes sense to think about their aggregate effects. (This is because $w^n(p, w) = w^{-n}$ for all p and w .)

4.2.3 Does individual WARP imply aggregate WARP?

The next aggregation question we consider is whether the fact that individuals make consistent choices implies that aggregate choices will be consistent as well. In terms of our discussion in Chapter 2, this involves the question of whether, when the Walrasian demand functions of the N consumers satisfy WARP, the resulting aggregate demand function will satisfy WARP as well. The answer to this question is, "Not necessarily."

To make things concrete, assume that there is a wealth distribution rule, so that it makes sense to talk about aggregate demand as $D(p, w) = (D_1(p, w), \dots, D_L(p, w))$. In fact, to keep things simple, assume that the wealth distribution rule is that $w^n(p, w) = a_n w$. Thus consumer n is assigned a fraction a_n of total wealth. Thus

$$D(p, w) = \sum_n x^n(p, w^n).$$

The aggregate demand function satisfies WARP if, for any two combinations of prices and aggregate wealth, (p, w) and (p^0, w^0) , if $p \cdot D(p^0, w^0) \leq w$ and $D(p, w) \neq D(p^0, w^0)$, then $p^0 \cdot D(p, w) > w^0$. This is the same definition of WARP as before.

The reason why individual WARP is not sufficient for aggregate WARP has to do with the Compensated Law of Demand (CLD). Recall that an individual's behavior satisfies WARP if and only if the CLD holds for all possible compensated price changes. The same is true for aggregate WARP. The aggregate will satisfy WARP if and only if the CLD holds in the aggregate for all possible compensated price changes. The problem is that just because a price change is compensated in the aggregate, it does not mean that the price change is compensated for each individual. Because of this, it does not necessarily follow from the fact that each individual's behavior satisfies the CLD that the aggregate will as well, since compensated changes in the aggregate need not imply compensated changes individually. See Example 4.C.1 and Figure 4.C.1 in MWG.

To make this a little more concrete without going into the details of the argument, think about how you would prove this statement: "If individuals satisfy WARP then the aggregate does as well." The steps would be:

1. Consider a compensated change in aggregate wealth.

2. This can be written as a sum of compensated changes in individual wealths.³

³Of course, this step is not true!

3. Individuals satisfy WARP if and only if they satisfy the CLD.
4. So, each individual change satisfies the CLD.
5. Adding over individual changes, the aggregate satisfies the CLD as well.

This proof is clearly flawed since step 2 is not valid. As shown above, it is not possible to write every price change that is compensated in the aggregate in terms of price changes that are compensated individual-by-individual. So, it turns out that satisfying WARP and therefore the CLD is not sufficient for aggregate decisions to satisfy WARP. However, if we impose stronger conditions on individual behavior, we can find a property that aggregates. That property is the Uncompensated Law of Demand (ULD). The ULD is similar to the CLD, but it involves uncompensated changes. Thus a demand function $x(p, w)$ satisfies the ULD if for any price change $p \rightarrow p^0$ the following holds:

$$p^0 - p \cdot x(p^0, w) - x(p, w) \leq 0.$$

Note the following:

1. If a consumer's demand satisfies the ULD, then it satisfies the CLD as well.
2. Unlike the CLD, the ULD aggregates. Thus if each consumer's demand satisfies the ULD, the aggregate demand function will as well.

Hence even though satisfaction of the CLD individually is not sufficient for the CLD in the aggregate, the ULD individually is sufficient for the ULD in the aggregate. So, the ULD individually implies WARP in the aggregate.

If we want to know which types of utility functions imply aggregate demand functions that satisfy WARP, we need to find those that satisfy the ULD. It turns out that homothetic preferences satisfy the ULD. Thus if each consumer has homothetic preferences, the implied aggregate demand will satisfy WARP.

In general, there is a calculus test to determine if a utility function satisfies the ULD property. It is given in MWG, and my advice is that if you ever need to know about such things, you look it up at that time. Basically, it has to do with making sure that wealth effects are not too strange (recall the example of the Giffen good — where the wealth effect leads to an upward sloping demand curve — the same sort of thing is a concern here).

4.2.4 Representative Consumers

The final question is when can the aggregate demand curve be used to make welfare measurements? In other words, when can we treat aggregate demand as if it is generated by a fictional “representative consumer,” and when will changes in the welfare of that consumer correspond to changes in the welfare of society as a whole?

The first part of this question is, when is there a rational preference relation \succsim such that the aggregate demand function corresponds to the Walrasian demand function generated by these preferences? If such a preference relation exists, we say that there is a positive representative consumer.

The first necessary condition for the existence of a positive representative consumer is that it makes sense to aggregate demand. Thus consumers must have indirect utility functions of the Gorman form (or wealth must be assigned by a wealth-assignment rule). In addition, the demand must correspond to that implied by the maximization of some rational preference relation. In essence, we need the Slutsky matrix to be negative semi-definite as well.

An additional question is whether the preferences of the positive representative consumer capture the welfare of society as a whole. This is the question of whether the positive representative consumer is normative as well. For example, suppose there is a social welfare function

$W(u_1, \dots, u_N)$ that maps utility levels for the N consumers to real numbers and such that utility vectors assigned higher numbers are better for the society than vectors assigned lower numbers. Thus $W(\cdot)$ is like a utility function for the society. Now suppose that for any level of aggregate wealth we assign wealth to the consumers in order to maximize W . That is, w^1, \dots, w^N solves

$$\begin{aligned} & \max_{w^1, \dots, w^N} W(v^1_i(p, w^1), \dots, v^N_i(p, w^N)) \\ \text{s.t. } & \sum_{n=1}^N w^n \leq w. \end{aligned}$$

Thus it corresponds to the situation where a benevolent dictator distributes wealth in the society in order to maximize social welfare. This defines a wealth assignment rule, so we know that aggregate demand can be represented as a function of p and total wealth w .

In the case where wealth is assigned as above, not only can demand be written as $D(p, w)$, but also these demand functions are consistent with the existence of a positive representative consumer. Further, if the aggregate demand functions are generated by solving the previous program, they have welfare significance and can be used to make welfare judgments (using the techniques from

Chapter 3).

An important social welfare function is the utilitarian social welfare function. The utilitarian social welfare function says that social welfare is the sum of the utilities of the individual consumers in the economy. Now, assume that all consumers have indirect utility functions of the Gorman Form: $v^n(p, w^n) = a^n(p) + b(p) w^n$. Using the utilitarian social welfare function implies that the social welfare maximization problem is:

$$\begin{aligned} & \max_{\{w^n\}} \sum_n v^n(p, w^n) \\ \text{s.t. : } & \sum_n w^n \leq w. \end{aligned}$$

But, this can be rewritten as:

$$\begin{aligned} & \max_{\{w^n\}} \sum_n a^n(p) + b(p) \sum_n w^n \\ \text{s.t. : } & \sum_n w^n \leq w, \end{aligned}$$

and any wealth assignment rule that fully distributes wealth, $\sum_n w^n = w$, solves this problem. The result is this: When consumers have indirect utility of the Gorman Form (with the same $b(p)$), aggregate demand can always be thought of as being generated by a normative representative consumer with indirect utility function $v(p, w) = \sum_n a^n(p) + b(p) w$, who represents the utilitarian social welfare function.

In fact, it can be shown that when consumers' preferences have Gorman Form indirect utility functions, then $v(p, w) = \sum_n a^n(p) + b(p) w$ is an indirect utility function for a normative representative consumer regardless of the form of the social welfare function.⁴ In addition, when consumers have Gorman Form utility functions, the indirect utility function is also independent of the particular wealth distribution rule that is chosen.⁵

This is all I want to say on the subject for now. The main takeaway message is that you should be careful about dealing with aggregates. Sometimes they make sense, sometimes they do not. And, just because they make sense in one way (i.e., you can write demand as $D(p, w)$), they may not make sense in another (i.e., there is a positive or normative consumer).

⁴This not generally true when consumers' preferences are not Gorman-form. The preferences of the normative representative consumer will depend on the particular social welfare function used to generate those preferences.

⁵Again, this property will not hold if consumers' preferences cannot be represented by a Gorman form utility function.

4.3 The Composite Commodity Theorem

There are many commodities in the world, but usually economists will only be interested in a few of them at any particular time. For example, if we are interested in studying the wheat market, we may divide the set of commodities into “wheat” and “everything else.” In a more realistic setting, an empirical economist may be interested in the demand for broad categories of goods such as “food,” “clothing,” “shelter,” and “everything else.” In this section, we consider the question of when it is valid to group commodities in this way.⁶

To make things simple, consider a three-commodity model. Commodity 1 is the commodity we are interested in, and commodities 2 and 3 are “everything else.” Denote the initial prices of goods 2 and 3 by p_2^0 and p_3^0 , and suppose that if prices change, the relative price of goods 2 and 3 remain fixed. That is, the price of goods 2 and 3 can always be written as $p_2 = t p_2^0$ and $p_3 = t p_3^0$, for $t \geq 0$. For example, if good 2 and good 3 are apples and oranges, this says that whenever the price of apples rises, the price of oranges also rises by the same proportion. Clearly, this assumption will be reasonable in some cases and unreasonable in others, but for the moment we will assume that this is the case.

The consumer’s expenditure minimization problem can be written as:

$$\begin{aligned} \min_{x \geq 0} & p_1 x_1 + t p_2^0 x_2 + t p_3^0 x_3 \\ \text{s.t.} & u(x) \geq u. \end{aligned}$$

Solving this problem yields Hicksian demand functions $h_i(p_1, t p_2^0, t p_3^0, u)$ and expenditure function $e_i(p_1, t p_2^0, t p_3^0, u)$.

Now, suppose that we are interested only in knowing how consumption of good 1 depends on t . In this case, we can make the following change of variables. Let $y = p_2^0 x_2 + p_3^0 x_3$. Thus y is equal to expenditure on goods 2 and 3, and t then corresponds to the “price” of this expenditure. As t increases, y becomes more expensive. Applying this change of variable to the $h(\cdot)$ and $e(\cdot)$ yields the new functions:

$$\begin{aligned} h^*(p_1, t, u) & \equiv h(p_1, t p_2^0, t p_3^0, u) \\ e^*(p_1, t, u) & \equiv e(p_1, t, t p_2^0, t p_3^0, u). \end{aligned}$$

⁶References: Silberberg, Section 11.3; Deaton and Muellbauer *Economics and Consumer Behavior*, pp. 120-122; Jehle and Reny, p. 266.

It remains to be shown that $h^*(\cdot)$ and $e^*(\cdot)$ satisfy the properties of well-defined compensated demand and expenditure functions (see Section 3.4). For $e^*(p_1, t, u)$, these include:

1. Homogeneity of degree 1 in p
2. Concavity in (p_1, t) (i.e. the Slutsky matrix is negative semi-definite)
3. $\frac{\partial e^*}{\partial t} = y$ (and the other associated derivative properties)

In fact, these relationships can be demonstrated. Hence we have the composite commodity theorem:

Theorem 8 When the prices of a group of commodities move in parallel, then the total expenditure on the corresponding group of commodities can be treated as a single good.

The composite commodity theorem has a number of important applications. First, the composite commodity theorem can be used to justify the two-commodity approach that is frequently used in economic models. If we are interested in the effect of a change in the price of wheat on the wheat market, assuming that all other prices remain fixed, the composite commodity theorem justifies treating the world as consisting of wheat and the composite commodity “everything else.”

A second application of the composite commodity theorem is to models of consumption over time, which we will cover later (see Section 4.6 of these notes). Since the prices of goods in future periods will tend to move together, application of the composite commodity theorem allows us to analyze consumption over time in terms of the composite commodities “consumption today,” “consumption tomorrow,” etc.

4.4 So Were They Just Lying to Me When I Studied Intermediate Micro?

Recall from your intermediate microeconomics course that you probably did welfare evaluation by looking at changes in Marshallian consumer surplus, the area to the left of the aggregate demand curve. But, I’ve told you that: a) consumer surplus is not a good measure of the welfare of an individual consumer; b) even if it were, it usually doesn’t make sense to think of aggregate demand as depending only on aggregate wealth (which it does in the standard intermediate micro model); and c) even if it did, looking at the equivalent variation (which is better than looking at the change

in consumer surplus) for the aggregate demand curve may not have welfare significance. So, at this point, most students are a little concerned that everything they learned in intermediate micro was wrong. The point of this interlude is to argue that this is not true. Although many of the assumptions made in order to simplify the presentation in intermediate micro are not explicitly stated, they can be explicitly stated and are actually pretty reasonable.

To begin, note that the point of intermediate micro is usually to understand the impact of changes on one or a few markets. For example, think about the change in the price of apples on the demand for bananas. It is widely believed that since expenditure on a particular commodity (like apples or bananas) is usually only a small portion of a consumer's budget, the income effects of changes in the prices of these commodities are likely to be small. In addition, since we are looking at only a few price changes and either holding all other prices constant or varying them in tandem, we can apply the composite commodity theorem and think of the consumer's problem as depending on the commodity in question and the composite commodity "everything else." Thus the consumer can be thought of as having preferences over apples, bananas, and everything else.

Now, since the income effects for apples and bananas are likely to be small, a reasonable way to represent the consumer's preferences is as being quasilinear in "everything else." That is, utility looks like:

$$u(a, b, e) = f(a, b) + e$$

where a = apples, b = bananas, and e = everything else. Once we agree that this is a reasonable representation of preferences for our purposes, we can point out the following:

1. Since there are no wealth effects for apples or bananas, the Walrasian and Hicksian demand curves coincide, and the change in Marshallian consumer surplus is the same as EV. Hence ΔCS is a perfectly fine measure of changes in welfare.
2. If all individual consumers in the market have utility functions that are quasilinear in everything else, then it makes sense to write demand as a function of aggregate wealth, since quasilinear preferences can be represented by indirect utility functions of the Gorman form.
3. Since all individuals have Gorman form indirect utility functions, then aggregate demand can always be thought of as corresponding to a representative consumer for a social welfare function that is utilitarian. Thus ΔCS computed using the aggregate demand curve has welfare significance.

Thus, by application of the composite commodity theorem and quasilinear preferences, we can save the intermediate micro approach. Of course, our ability to do this depends on looking at only a few markets at a time. If we are interested in evaluating changes in many or all prices, this may not be reasonable. As you will see later, this merely explains why partial equilibrium is a topic for intermediate micro and general equilibrium is a topic for advanced micro.

4.5 Consumption With Endowments

Until now we have been concerned with consumers who are endowed with initial wealth w . However, an alternative approach would be to think of consumers as being endowed with both wealth w and a vector of commodities $a = (a_1, \dots, a_L)$, where a_i gives the consumer's initial endowment of commodity i .⁷ In this case, the consumer's UMP can be written as:

$$\begin{aligned} & \max_x u(x) \\ \text{s.t.} \quad & p \cdot x \leq p \cdot a + w \end{aligned}$$

The value of the consumer's initial assets is given by the sum of her wealth and the value of her endowment, $p \cdot a$. Thus the mathematical approach is equivalent to the situation where the consumer first sells her endowment and then buys the best commodity bundle she can afford at those prices.

The first-order conditions for this problem are found in the usual way. The Lagrangian is given by:

$$L = u(x) + \lambda (p \cdot a + w - p \cdot x)$$

implying optimality conditions:

$$\begin{aligned} u_i - \lambda p_i &= 0 : i = 1, \dots, L. \\ p \cdot x - p \cdot a - w &\leq 0. \end{aligned}$$

Denote the solution to this problem as

$$x(p, w, a)$$

where w is non-endowment wealth and a is the consumer's initial endowment.

⁷ Reference: Silberberg (3rd edition), pp. 299-304.

We can also solve a version of the expenditure minimization problem in this context. Consider the problem:

$$\begin{aligned} \min_x & p \cdot x - p \cdot a \\ \text{s.t.} & : u(x) \geq u. \end{aligned}$$

The objective function in this model is non-endowment wealth. Thus it plays the role of w in the UMP, and the question asked by this problem can be stated as: How much non-endowment wealth is needed to achieve utility level u when prices are p and the consumer is endowed with a ?

The endowment a drops out of the Lagrangian when you differentiate with respect to x_i . Hence the non-endowment expenditure minimizing bundle (NEEMB) is not a function of a . We'll continue to denote it as $h(p, u)$. However, while the NEEMB does not depend on a , the non-endowment expenditure function does. Let

$$e^*(p, u, a) \equiv p \cdot (h(p, u) - a).$$

Again, $e^*(p, u, a)$ represents the non-endowment wealth necessary to achieve utility level u as a function of p and endowment a . By the envelope theorem (or the derivation for $\frac{\partial e}{\partial p_i} = h_i(p, u)$ we

did in Section 3.4.3) it follows that

$$\frac{\partial e^*}{\partial p_i} \equiv h_i(p, u) - a_i.$$

Thus the sign of $\frac{\partial e^*}{\partial p_i}$ depends on whether $h_i(p, u) > a_i$ or $h_i(p, u) < a_i$. If $h_i(p, u) > a_i$ the consumer is a net purchaser of good i , consuming more of it than her initial endowment. If this is the case, then an increase in p_i increases the cost of purchasing the good i from the market, and this increases total expenditure at a rate of $h_i(p, u) - a_i$. On the other hand, if $h_i(p, u) < a_i$, then the consumer is a net seller of good i , consuming less of it than her initial endowment. In this case, increasing p_i increases the revenue the consumer earns by selling the good to the market. The result is that the non-endowment wealth the consumer needs to achieve utility level u decreases at a rate of $|h_i(p, u) - a_i|$.

Now, let's rederive the Slutsky equation in this environment. The following identity relates $h(\cdot)$ and $x(\cdot)$:

$$h_i(p, u) \equiv x_i(p, e^*(p, u, a), a).$$

Differentiating with respect to p_j yields:

$$\begin{aligned} \frac{\partial h_i}{\partial p_j} &\equiv \frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} \frac{\partial e^*}{\partial p_j} \\ &\equiv \frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} (h_j(p, u) - a_j) \\ &\equiv \frac{\partial x_i}{\partial p_j} + \frac{\partial x_i}{\partial w} (x_j(p, w, a) - a_j). \end{aligned}$$

A useful reformulation of this equation is:

$$\frac{\partial x_i}{\partial p_j} = \frac{\partial h_i}{\partial p_j} - \frac{\partial x_i}{\partial w} (x_j(p, w, a) - a_j).$$

The difference between this version of the Slutsky equation and the standard form is in the wealth effect. Here, the wealth effect is weighted by the consumer's net purchase of good i .⁸ So, think about a consumer who is endowed with a_1 units of good 1 and faces an increase in p_1 . For concreteness, say that good 1 is gold, I am the consumer, and we are interested in my purchases of new ties (good 2) in response to a change in the price of gold. If the price of gold goes up, I will tend to purchase more ties if we assume that ties and gold are substitutes in my utility function.

This means that $\frac{\partial h_2}{\partial p_1} > 0$. However, an increase in the price of gold will also have a wealth effect.

Whether this effect is positive or negative depends on whether I am a net purchaser or net seller of gold. If I buy more gold than I sell, then the price increase will be bad for me. In terms of the Slutsky equation, this means $(x_1 - a_1) > 0$. For a normal good ($\frac{\partial x_2}{\partial w} > 0$), this means that $\frac{\partial x_2}{\partial p_1}$ will be smaller than $\frac{\partial h_2}{\partial p_1}$ — I shift consumption towards ties due to the price change, but the price increase in gold makes me poorer so I don't increase tie consumption quite as much as in a compensated price change.

If I am a net seller of gold, an increase in the price of gold has a positive effect on my wealth. Since I am selling gold to the market, increasing its price p_1 actually makes me wealthier in proportion to $(a_1 - x_1)$. And, since the price change makes me wealthier (because $x_1 - a_1 < 0$), the effect of the whole wealth/endowment term is subtracting a negative number (again, assuming ties are a normal good). Thus $\frac{\partial x_2}{\partial p_1}$ will be greater than $\frac{\partial h_2}{\partial p_1}$, and I will consume more ties due to both the price substitution effect and the fact that the price change makes me wealthier.

⁸Actually, there is no difference between this relationship and the standard Slutsky equation. The standard model is equivalent to this model where $a = (0, \dots, 0)$. If you insert these values into the Slutsky equation with endowments, you get the exactly the standard version of the Slutsky equation.




Figure 4.3: Labor-Leisure Choice

Thus the main difference between the standard model and the endowment model lies in this adjustment to the Slutsky matrix: The wealth effect must be adjusted by whether a consumer is a net purchaser or a net seller of the good in question. This has important applications in general equilibrium theory (which we'll return to much later), as well as applications in applied consumption models. We turn to one such example here.

4.5.1 The Labor-Leisure Choice

As an application of the previous section, consider a consumer's choice between labor and leisure. We are interested in the consumer's leisure decision, so we'll apply the composite commodity theorem and model the consumer as caring about leisure, l , and everything else, y . Let the consumer's utility function be

$$u(y, l).$$

If the wage rate is s , w is non-endowment wealth, and the price of "everything else" is normalized to 1, the consumer's budget constraint is given by:

$$y \leq s(24 - l) + w.$$

The solution to this problem is given by the point of tangency between the utility function and the budget set. This point is illustrated in the Figure 4.3.

Initially, the consumer is endowed with 24 hours of leisure per day. Since the consumer cannot consume more than 24 hours of leisure per day, at the optimum the consumer must be a net seller




Figure 4.4: A Wage Increase

of leisure. Thus an increase in the price of leisure, s , increases the consumer's wealth. Hence the compensation must be negative. A compensated increase in the price of leisure is illustrate in Figure 4.4. At the original wage rate the consumer maximizes utility by choosing the bundle at point A. Since the consumer is a net seller of leisure, the compensated change in demand for leisure is negative. So, when compensated for the price change, the consumer's choice moves from point A to point B, and she consumes less leisure at the higher wage rate. However, since the consumer is a net seller of leisure, the compensation is negative. Hence when going from the compensated change to uncompensated change we move from point B to point C. That is, the wealth effect here leads to the consumer consuming more leisure than before the compensation took place.

Let's think of this another way. Suppose that wages increase. Since you get paid more for every additional hour you work, you will tend to work more (which means that you will consume less leisure). However, since you make more for every hour you work, you also get paid more for all of the hours you are already working. This makes you wealthier, and because of it you will tend to want to work less (that is, consume more leisure, assuming it is a normal good). Thus the income effect and substitution effect work in opposite directions here precisely because the consumer is a net seller of leisure. This is in contrast with consumer theory without endowments, where you decrease consumption of a normal good whose price has increased, both because its relative price has increased and because this increase has made you poorer.

Note that it is also possible to get a Giffen-good like phenomenon here even though leisure is a normal good. This happens if the income effect is much larger than the substitution effect,




Figure 4.5: Positive Labor Supply Elasticity

as in Figure 4.5, where the arrows depict the large income effect (point B to point C). As an illustration, think of the situation in which a person earns minimum wage, let's say \$5 per hour, and chooses to work 60 hours per week. That gives total wages of \$300 per week. If the government raises the minimum wage by \$1 per hour, this increases the consumer's total wages to \$360, a 20% increase. The consumer likely has two responses to this. Since the consumer gets paid more for each additional hour of work, she may decide to work more hours (since she will be willing to give up more leisure at the higher wage rate). However, since the \$1 increase in wages has increased total wage revenue by 20% already, this may make the consumer work less, since she is already richer than before. In situations where the change in total wages is large relative to the wage rate (i.e., the consumer is working a lot of hours), the latter effect may swamp the former.

There have been many studies of this labor-leisure tradeoff in the U.S. They are frequently associated with worries over whether raising taxes on the wealthy will cause them to cut their labor supply. My understanding of the evidence (through conversations with labor economists mostly) is that labor supply elasticities are positive but small, similar to the depiction in Figure 4.5.⁹

4.5.2 Consumption with Endowments: A Simple Separation Theorem

The first half of the course deals with doing welfare evaluation correctly. The second half of the course deals with markets and market interactions. One of the themes I will try to bring out is

⁹In fact, labor supply elasticities tend to be pretty small for men, larger for women, but always positive (i.e. an increase in wages - or a cut in income taxes - leads people to work more).

the idea that when markets work well, people make good decisions. That is, they make decisions that maximize economic value, and then they use markets to buy and sell commodities to arrive at the consumption bundle they want. Remarkably, it turns out that when markets are perfect, this approach also maximizes the consumer's welfare. However, if markets do not function well, then people may be forced to distort their decisions to make up for the fact that they cannot use markets to modify their consumption bundle. For example, if markets are perfect and the price of bananas is high and the price of coconuts is low, even a farmer who hates bananas and loves coconuts will be best off by choosing to grow bananas. But, once the bananas are harvested he will sell them and use the proceeds to purchase coconuts. In this way, he will be able to eat more coconuts than if he grew them himself. If there are not good markets, then the farmer will be forced to eat what he grows. In this case, he will be forced to grow coconuts, and he will end up with fewer coconuts than he would have if the markets were better. In addition to the "micro" effects, there can also be larger scale effects. If everyone makes suboptimal decisions because markets are not well developed, then overall growth may be adversely impacted.

To illustrate, consider a consumer who must choose between endowment $a = (a_1, a_2)$ and endowment $b = (b_1, b_2)$. (Assume non-endowment wealth is $w = 0$ for the sake of the diagrams.) How should the consumer choose? If markets are perfect, the consumer should choose whichever bundle has the higher market value. After all, if bundle a has higher market value than bundle b , then the budget set for endowment a includes the budget set for endowment b , and therefore the consumer must be strictly better off at a than b . Thus, $a \hat{A} b$ if and only if $p \cdot a > p \cdot b$. Figure 1 illustrates. Note, however, that a critical assumption underlying this is that the price at which you can buy a commodity is the same as the price at which you can sell it. This is often not the case. In fact, it is the norm in markets for the "buy price" to be greater than the "sell price," and the extent to which the two differ is often interpreted as a sign of market development or competitiveness.

Next, consider the case where the buy price of a good is greater than the sell price. Let p^b_1 and p^b_2 denote the buy prices and p^s_1 and p^s_2 denote the sell prices, and suppose $p^b_1 > p^s_1$ and $p^b_2 > p^s_2$. In this case, the budget set will have a kink at the endowment point. Above the endowment point, the slope of the budget line is $-p^s_1$ (since over this range the consumer is selling x_1 and using the money he makes to buy x_2). Below the endowment point, the slope of the budget line is $-p^b_1$, since over this range the consumer is selling good 2 and using the profit to buy good 1. (See Figure 2.) Since $p^b_1 > p^s_1$ and $p^b_2 > p^s_2$, the slope is steeper when the consumer buys 1 and sells 2 than when




Figure 4.6:

he sells 1 and buys 2. (If you don't believe me, plug in some numbers for p^b_1 , p^s_1 , p^b_2 , and p^s_2 .)

Now, return to the question of whether the consumer should prefer endowment a or endowment b. Now the answer to the question is: "it depends." Consider Figure 3. In this case, whether the consumer prefers a or b will depend on the nature of his preferences. If he has a strong preference for good x_1 , he may choose endowment b. This is true even though the "market value" of a would be larger if there were no bid-ask spread.

Thus, the inability to increase consumption of x_1 through the market may lead the consumer to make decisions that maximize short run utility but have negative long-run consequences. For example, if x_1 is a food crop and x_2 is a non-food crop, then a represents focusing on the cash crop while b represents focusing on the subsistence crop. If markets are good, the consumer should grow the cash crop and use the market to purchase food. If markets are not good, then the consumer will have to grow the food crop, passing on the opportunity to increase welfare by growing the cash crop.

This is known as a "separation result" because it essentially says that if markets are perfect, then the consumer's production decision (which endowment to choose) and consumption decision (what to consume) can be separated. The consumer maximizes welfare by making the production




Figure 4.7:




Figure 4.8:

decision that maximizes the value of the endowment and then maximizes utility given the resulting budget set. We will see these kind of results in a wide variety of circumstances. Often, and especially in a developing context, the real importance of these results is not when they work (since markets are never perfect!), but when they don't. In this case, separation results suggest that improving markets can improve welfare. This is much more interesting and useful than the way the result is usually stated.

4.6 Consumption Over Time

Up until now we have been considering a model of consumption that is static. Time does not enter into our model at all. This model is very useful for modeling a consumer's behavior at a particular point in time. It is also useful for modeling the consumer's behavior in two different situations. This is what we called "comparative statics." However, as the name suggests, even though the consumer's behavior in two different situations can be compared using the static model, we are really just comparing two static situations: No attempt is made to model how the consumer's behavior evolves over time.

While the static model is useful for answering some questions, often we will be interested specifically in the consumer's consumption decisions over time. For example, will the consumer borrow or save? Will her consumption increase or decrease over time? How are these conclusions affected by changes in exogenous parameters such as prices, interest rates, or wealth?

Fortunately, we can adapt our model of static consumption to consider dynamic situations. There are two key features of the dynamic model that need to be addressed. First, the consumer may receive her wealth over the course of her lifetime. But, units of wealth today and units of wealth tomorrow are not worth the same to the consumer. Thus we must come up with a way to measure wealth received (or spent) at different times. Second, there are many different commodities sold and consumed during each time period. Explicitly modeling every commodity would be difficult, and it would make it harder to evaluate broad trends in the consumer's behavior, which is what we are ultimately interested in.

The solution to these problems is found in the applications of consumer theory that we have been developing. The first step is to apply the composite commodity theorem. Since prices at a particular time tend to move in unison, we can combine all goods bought at a particular time into a composite commodity, "consumption at time t ." We can then analyze the dynamic problem as a

static problem in which the commodities are “consumption today,” “consumption tomorrow,” etc. The problem of wealth being received over time is addressed by adding endowments to the static model. Thus the consumer’s income (addition to wealth) during period t can be thought of as the consumer’s endowment of the composite commodity “consumption at time t .” The final issue, that of capturing the fact that a unit of wealth today is worth more than a unit of wealth tomorrow, is addressed by assigning the proper prices to consumption in each period. This is done through a process known as discounting.

4.6.1 Discounting and Present Value

Suppose that you have \$1 today that you can put in the bank. The interest rate the bank pays is 10% per year. If you invest this dollar, you have \$1.10 at the end of the year. On the other hand, suppose that you need to have \$1 at the end of the year. How much should you invest today in order to make sure that you have \$1 at the end of the year? The answer to this question is given by the solution to the equation:

$$(1 + .1)y = 1$$

$$y = \frac{1}{1+.1} = .91.$$

Thus in order to make sure you have \$1 a year from now, you should invest 91 cents today.

To put the question of the previous paragraph another way, if I were to offer you \$1 a year from now or y dollars today, how large would y have to be so that you are just indifferent between the dollar in a year and y today? The answer is $y = 0.91$ (assuming the interest rate is still 10%).¹⁰ Thus we call \$0.91 the present value of \$1 a year from now because it is the value, in current dollars, of the promise of \$1 in a year.

In fact, we can think of the 91 cents in another way. We can also think of it as the price, in current dollars, of \$1 worth of consumption a year from now. In other words, if I were to offer to buy you \$1 worth of stuff a year from now and I wanted to break even, I should charge you a price of 91 cents.

The concept of present value can also be used to convert streams of wealth received over multiple years into their current-consumption equivalents. Suppose we call the current period 0, and that the world lasts until period T . If the consumer receives a_t dollars in period t , and the interest rate is r (and remains constant over time), then the present value of this stream of payments is given

¹⁰ This answer ignores the issue of impatience, which we will address shortly.

by:

$$P Va = a_0 + \sum_{t=1}^T \frac{a_t}{(1+r)^t} = \sum_{t=0}^T \delta^t a_t. \quad (4.2)$$

where $\delta = \frac{1}{1+r}$ is the discount factor.¹¹ But, this can also be thought of as a problem of consumption with endowments. Let the commodities be denoted by $x = (x_0, \dots, x_T)$, where x_t is consumption in period t (by application of the composite commodity theorem). Let a_t be the consumer's endowment of the consumption good in period t . Then, if we let the price of consumption in period t , denoted p_t , be $p_t = \frac{1}{(1+r)^t}$, the present value formula above can be written as:

$$P Va = \sum_{t=0}^T p_t a_t = p \cdot a$$

where $p = (p_0, \dots, p_T)$ and $a = (a_0, \dots, a_T)$. But, this is exactly the expression we had for endowment wealth in the model of consumption with endowments. This provides the critical link between the static model and the dynamic model.

4.6.2 The Two-Period Model

We now show how the approach developed in the previous section can be used to develop a model of consumption over time. Suppose that the consumer lives for two periods: today (called period 0) and tomorrow (called period 1). Let x_0 and x_1 be consumption in periods $t = 0$ and $t = 1$, respectively, and let a_0 and a_1 be income (endowment) in each period, measured in units of consumption. Suppose that the consumer can borrow or save at an interest rate of $r \geq 0$. Thus the price of consumption in period t (in terms of consumption in period 0) is given by $p_t = \frac{1}{(1+r)^t}$.

Assume that the consumer has preferences over consumption today and consumption tomorrow represented by utility function $u(x_0, x_1)$, and that this utility function satisfies all of the nice properties: It is strictly quasiconcave and strictly increasing in each of its arguments, and twice differentiable in each argument. The consumer's UMP can then be written as:

$$\begin{aligned} & \max_{x_0, x_1} u(x_0, x_1) \\ \text{s.t.} & \quad : \\ & x_0 + \frac{x_1}{1+r} \leq a_0 + \frac{a_1}{1+r} \end{aligned}$$

¹¹ In the event that the interest rate changes over time, the interest rate r can be replaced with the period-specific interest rate, r_t , and the discount rate is then $\delta_t = \frac{1}{1+r_t}$.

where, of course, the constraint is just another way of writing $p \cdot x \leq p \cdot a$, which just says that the present value of consumption must be less than the present value of the consumer's endowment. It is simply a dynamic version of the budget constraint.¹² Note that since $p_0 = 1$, the exogenous parameters in this problem are r and a . It is convenient to write them as $p_1 = \frac{1}{1+r}$ and a , however, and we will do this.

This problem can be solved using the standard Lagrangian methodology:

$$L = u(x_0, x_1) + \lambda \left[a_0 + \frac{a_1}{1+r} - x_0 - \frac{x_1}{1+r} \right].$$

Assuming an interior solution, first-order conditions are given by:

$$\frac{u_{x_1}}{u_{x_0}} = \frac{\lambda}{(1+r)^t} : t \in \{0, 1\}.$$

$$x_0 + \frac{x_1}{1+r} \leq a_0 + \frac{a_1}{1+r}.$$

Of course, as before, we know that the budget constraint will bind. This gives us our three equations in three unknowns, which can then be solved for the demand functions $x_t(p_1, a)$. The arguments of the demand functions are the exogenous parameters — interest rate r and endowment vector $a = (a_0, a_1)$. See Figure 12.1 in Silberberg for a graphical illustration — it's just the same as the standard consumer model, though.

We can also consider the expenditure minimization problem for the dynamic model. Earlier, we minimized the amount of non-endowment wealth needed to achieve a specified utility level. We do the same here, where non-endowment wealth is taken to be wealth in period 0.

$$\begin{aligned} \min a_0 &= x_0 + p_1(x_1 - a_1) \\ \text{s.t.} \quad & : u(x) \geq u. \end{aligned}$$

The Lagrangian is given by:

$$L = x_0 + p_1(x_1 - a_1) - \lambda(u(x) - u).$$

¹² Note that we could use the same model where the price of consumption in period t is not necessarily given by $p_t = \frac{1}{(1+r)^t}$. This approach will work whenever the price of consumption in period t in terms of consumption today is well-defined, even if it is not given by the above formula. The advantage of the discount-rate formulation is that it allows us to consider the impact of changes in the interest rate on consumption.

The first-order conditions are derived as in the standard EMP, and the solution can be denoted by $h_t(r, u)$.¹³ Let $a_0(p_1, a) = h_0(p_1, u) + p_1(h_1(p_1, u) - a_1)$ be the minimum wealth needed in period 0 to achieve utility level u when the interest rate is r .

Finally, we can link the solutions to the UMP and EMP in this context using the identity:

$$h_t(p_1, u) = x_t(p_1, a_0(p_1, u), a_1).$$

Differentiating with respect to p

$$\begin{aligned} \frac{\partial h_t}{\partial p_1} &= \frac{\partial x_t}{\partial p_1} + \frac{\partial x_t}{\partial a_0} \frac{\partial a_0}{\partial p_1} \\ &= \frac{\partial x_t}{\partial p_1} + \frac{\partial x_t}{\partial a_0} (h_1(p_1, u) - a_1) \\ &= \frac{\partial x_t}{\partial p_1} + \frac{\partial x_t}{\partial a_0} (x_1(p_1, a) - a_1). \end{aligned}$$

Using this version of the Slutsky equation, we can determine the effect of a change in the interest rate in each period. Let $t = 1$, and rewrite the Slutsky equation as:

$$\frac{\partial x_1}{\partial p_1} = \frac{\partial h_1}{\partial p_1} + \frac{\partial x_1}{\partial a_0} (a_1 - x_1).$$

If r decreases, the price of future consumption (p_1) increases. We know that the compensated change in demand for future consumption $\frac{\partial h_1}{\partial p_1} \leq 0$. In fact, it is most likely negative: $\frac{\partial h_1}{\partial p_1} < 0$.

The wealth effect depends on whether x_1 is normal or inferior (i.e., the sign of $\frac{\partial x_1}{\partial a_0}$) and whether the

consumer saves in period 0 (implying $a_1 < x_1$) or borrows in period 0 (meaning $a_1 > x_1$). Since x_1 is all consumption in period 1, it only makes sense to think of it as normal.

So, there are two cases to consider. Either the consumer borrows or saves in period 0. If the consumer saves in period 0, the wealth effect will reinforce the compensated change in demand. That is, when p_1 goes up, the consumer reduces period 1 consumption both because the relative price has gone up (the substitution effect) and because the price increase (interest rate decrease) makes the saver "poorer" in the future. Conversely, if the consumer borrows in period 0, then the substitution and wealth effects go in opposite directions. The substitution effect is still negative. However, the increase in p_1 (decrease in the interest rate) makes the consumer wealthier in the future since less consumption must be forfeited in period 1 to finance the same consumption in period 0. In this case, the effect on second period consumption is ambiguous: $\frac{\partial h_1}{\partial p_1}$ is negative, but

the wealth effect is positive.

¹³ Note that the endowment in period 1 (a_1) drops out when you differentiate. Hence the expenditure minimizing consumption bundle does not depend on a_1 (although the amount of period 0 wealth needed to purchase that consumption bundle will depend on a_1). Thus, h_t is not a function of a_1 , but e is.

4.6.3 The Many-Period Model and Time Preference

This section has three aims: 1) To extend the two-period model of the previous section to a many-period model; 2) To incorporate into our model the idea that people's attitudes toward intertemporal substitution remain constant over time - we call this idea dynamic consistency; 3) To incorporate into our model the idea that people are impatient.

Extending the model to multiple periods is straightforward. Define utility over consumption in periods 0 through T as $U(x_0, \dots, x_T)$. The UMP is then given by:¹⁴

$$\begin{aligned} & \max_{x_0, \dots, x_T} U(x_0, \dots, x_T) \\ \text{s.t. : } & \sum_{t=0}^T \frac{x_t}{(1+r)^t} \leq \sum_{t=0}^T \frac{a_t}{(1+r)^t} \end{aligned}$$

What does it mean for consumers to have dynamically consistent preferences, i.e., attitudes toward intertemporal substitution that remain constant over time? The idea is that your willingness to sacrifice a unit of consumption in period t_0 for a unit of consumption in period t_1 should depend only on the amount you are currently consuming in periods t_0 and t_1 and the amount of time between t_0 and t_1 : $t_1 - t_0$. For example, suppose the time period of consumption is 5 years, and that the consumer's current consumption path (which is not necessarily optimal) is given by:

	x_0	x_1	x_2	x_3	x_4
Consumption			10	20	
		5		10	20

If the consumer's attitudes toward intertemporal substitution remain constant, then the amount of consumption the consumer would be willing to give up in period 0 for an additional unit of consumption in period 1 should be the same as the amount of consumption the consumer is willing to give up in period 3 for an additional unit of consumption in period 4. This amount depends on the consumption in the two periods under consideration, 10 and 20 in each case, and on the amount of time between the periods, 1 in each case. Thus, for example, dynamic consistency implies that the consumer will prefer $x_0 = 11, x_1 = 19, x_2 = 5, x_3 = 10, x_4 = 20$ to the current consumption path if and only if she prefers $x_0 = 10, x_1 = 20, x_2 = 5, x_3 = 11,$ and $x_4 = 19$ to the current consumption path.

What we mean by impatience is this: Suppose I were to give you the choice between your favorite dinner today or the same dinner a year from now. Intuition about people as well as lots of

¹⁴ Note that utility over consumption paths here has been written as capital $U(x_0, \dots, x_T)$. There will be a function called small $u()$ in a minute.

experimental evidence tell us that almost everybody would rather have the meal today. Thus the meaning of impatience is that, all else being equal, consumers would rather consume sooner than later. Put another way, assume that you currently plan to consume the same amount today and tomorrow. The utility associated with an additional unit of consumption today is greater than the utility of an additional unit of consumption tomorrow.

Impatience and dynamic consistency of preferences are most easily incorporated into our consumer model by assuming that the consumer's utility function can be written as:

$$U(x_0, \dots, x_T) = \frac{\sum_{t=0}^T u(x_t)}{(1 + \rho)^t},$$

where $u(x_t)$ gives the consumer's utility from consuming x_t units of output in period t and $\rho > 0$ is the consumer's rate of time preference. Note that lower-case $u(x)$ gives utility of consuming x_t in a single period, while capital $U(x_0, \dots, x_T)$ is the utility from consuming consumption vector (x_0, \dots, x_T) .

We can confirm that this utility function exhibits impatience and dynamic consistency in a straightforward manner. Impatience is easy. Consider two periods t_0 and t_1 such that $t_1 > t_0$ and $x_{t_0} = x_{t_1} = x^*$. Marginal utility in periods t_0 and t_1 are given by:

$$U_{t_0} = \frac{u^0(x^*)}{(1 + \rho)^{t_0}}$$

$$U_{t_1} = \frac{u^0(x^*)}{(1 + \rho)^{t_1}}$$

And, $U_{t_0} - U_{t_1} = u^0(x^*) \left[\frac{1}{(1 + \rho)^{t_0}} - \frac{1}{(1 + \rho)^{t_1}} \right]$ which is positive whenever $t_1 > t_0$. Thus the consumer is impatient.

To check dynamic consistency, compute the consumer's marginal rate of substitution between two periods, t_0 and t_1 :

$$\frac{U_{t_1}}{U_{t_0}} = \frac{\frac{u^0(x_{t_1})}{(1 + \rho)^{t_1}}}{\frac{u^0(x_{t_0})}{(1 + \rho)^{t_0}}} = \frac{u^0(x_{t_1})}{u^0(x_{t_0})} (1 + \rho)^{t_0 - t_1}$$

Since the marginal rate of substitution depends only on the consumption in each period x_{t_1} and x_{t_0} and the amount of time between the two periods, $t_0 - t_1$, but not on the periods themselves, this utility function is also dynamically consistent.

Because it satisfies these two properties, we will use a utility function of the form:

$$U(x_0, \dots, x_T) = \frac{\sum_{t=0}^T u(x_t)}{(1 + \rho)^t},$$

for most of our discussion. We will assume that $U(x_0, \dots, x_T)$ is strictly quasiconcave, and increasing and differentiable in each of its arguments.

Question: Does this mean that $u()$ is concave? Answer: No!

In the multi-period version of the dynamic consumer model, the UMP can be written as:

$$\begin{aligned} & \max_{x_0, \dots, x_T} \sum_{t=0}^T \frac{u(x_t)}{(1+\rho)^t} \\ & \text{s.t. : } \sum_{t=0}^T \frac{x_t}{(1+r)^t} \leq \sum_{t=0}^T \frac{a_t}{(1+r)^t} \end{aligned}$$

The Lagrangian is set up in the usual way, and the first-order conditions for an interior solution are:

$$\frac{u'(x_t)}{(1+\rho)^t} - \frac{\lambda}{(1+r)^t} = 0.$$

This implies that for two periods t^0 and t^{00} , the tangency condition is:

$$\frac{u'(x_{t^0})}{(1+\rho)^{t^0}} = \frac{u'(x_{t^{00}})}{(1+\rho)^{t^{00}}} \cdot \frac{1+r}{1+\rho}.$$

And, for two consecutive periods, $t^{00} = t^0 + 1$, this condition becomes:

$$\frac{u'(x_{t^0})}{(1+\rho)^{t^0}} = \frac{u'(x_{t^0+1})}{(1+\rho)^{t^0+1}} \cdot \frac{1+r}{1+\rho}. \tag{4.3}$$

Armed with this tangency condition, we are prepared to ask the question, "Under what circumstances will consumption be increasing over time?"

Intuitively, what do you think the answer is? Hint: Consumption will be increasing over time if the consumer is (more or less) impatient than the market? What does it have to do with how r and ρ compare?

To make things simple, let's consider periods 1 and 2. The same analysis holds for any other two adjacent periods. By quasiconcavity of $U()$, we know that the consumer's indifference curves will be convex in the (x_1, x_2) space, as in Figure 4.9. When $x_1 = x_2$, the slope of the utility

isoquant is given by $-\frac{u'(x_1)}{u'(x_2)} = -\frac{1+\rho}{1+r}$. When $x_1 > x_2$, this slope is less than $(1+\rho)$ in

absolute value. When $x_2 > x_1$, this slope is greater than $(1+\rho)$ in absolute value. The tangency condition (4.3) says that the absolute value of the slope of the isoquant must be the same as $(1+\rho)$. Thus if $1+r > 1+\rho$ (which is equivalent to $r > \rho$), the optimal consumption point must have $x_2 > x_1$: Consumption rises over time. If $1+r < 1+\rho$ (which is equivalent to $r < \rho$), $x_1 > x_2$, and consumption falls over time. If $1+r = 1+\rho$, consumption is constant over time.




Figure 4.9: Two-Period Consumption

What is the significance of the comparison between ρ and r ? Starting from the situation where consumption is equal in both periods, the consumer is willing to give up 1 unit of future consumption for an additional $\frac{1}{1+\rho}$ units of consumption today. By giving up one unit of future consumption, the consumer can buy an additional $\frac{1}{1+r}$ units of consumption today.

Thus if $\frac{1}{1+r} > \frac{1}{1+\rho}$, the consumer is willing to give up this unit of future consumption: Optimal consumption decreases over time. This condition will hold whenever $\rho > r$. On the other hand, if

$$\frac{1}{1+r} < \frac{1}{1+\rho}$$

$$\rho < r,$$

the consumer would rather shift consumption into the future: Optimal consumption rises over time. In words, if you are more patient than the market, consumption tends to grow over time; but if you are less patient than the market, consumption tends to shrink over time.

4.6.4 The Fisher Separation Theorem

Suppose that the consumer must choose between two careers. Career A yields endowment vector $a = (a_0, \dots, a_T)$. Career B yields endowment vector $b = (b_0, \dots, b_T)$. Which should the consumer choose?

One is tempted to think that in order to decide you have to solve the consumer's UMP for the two endowment vectors and compare the consumer's utility in the two cases. A remarkable result demonstrated by Irving Fisher, known as the Fisher Separation Theorem, shows that if the consumer has free access to credit markets, then the consumer should choose the endowment vector that has the largest present value. Put another way, the consumer's production decision (which endowment vector to choose) and her consumption decision (which consumption vector to choose)

are separate. The consumer maximizes utility by first choosing the endowment vector with the largest present value and then choosing the consumption vector that maximizes utility, subject to the budget constraint implied by that endowment vector.

First, we need to explain what is meant by free access to credit markets. Basically, this means that the consumer can borrow or lend as much wealth as she wants at interest rate r , as long as her budget balances over the entire time horizon of the model. That is, all consumption vectors such that

$$\sum_{t=0}^T \frac{x_t}{(1+r)^t} \leq \sum_{t=0}^T \frac{a_t}{(1+r)^t}$$

are available to the consumer.

The Fisher Separation theorem follows as a direct consequence of this. Let $P V_a = \sum_{t=0}^T \frac{a_t}{(1+r)^t}$ and $P V_b = \sum_{t=0}^T \frac{b_t}{(1+r)^t}$. The consumer's UMP for endowments a and b are given by:

$$\begin{aligned} & \max_x U(x_0, \dots, x_T) \\ \text{s.t.} : & \sum_{t=0}^T \frac{x_t}{(1+r)^t} \leq P V_a \end{aligned}$$

and

$$\begin{aligned} & \max_x U(x_0, \dots, x_T) \\ \text{s.t.} : & \sum_{t=0}^T \frac{x_t}{(1+r)^t} \leq P V_b. \end{aligned}$$

These problems are identical except for the right-hand side of the budget constraint. And, since we know that when utility is locally non-satiated, utility increases when the budget constraint is relaxed, so the consumer will achieve higher utility by choosing the endowment vector with the higher present value. It's that simple.

When the credit markets are not complete, the separation result will not hold. In particular, if the interest rate for saving is less than the interest rate on borrowing (as is usually the case in the real world), then the opportunities available to the consumer will depend not only on the present value of her endowment but also on when the endowment wealth is received. For example, consider Figure 4.10. Here, the interest rate on borrowing, r , is greater than the interest rate on saving, R . Because of this, beginning from initial endowment a , the budget line is flatter when the consumer saves (moves toward higher future consumption) than when she borrows (moves toward higher present consumption). Figure 4.10 depicts the budget sets for two initial endowments, a




Figure 4.10: Imperfect Credit

and b. Since neither budget set is included in the other, we cannot say whether the consumer prefers endowment a or endowment b without solving the UMP.