CHAPTER 2

Rational Choice under Certainty

2.1 Introduction

As promised, we begin by discussing the theory of rational choice. This theory forms the foundation of virtually all modern economics and is one of the first things you would learn in a graduate-level microeconomics class. As a theory of rational choice (see Section 1.1), the theory specifies what it means to make rational decisions – in short, what it means to be rational.

In this chapter, we consider **choice under certainty**. The phrase "under certainty" simply means that there is no doubt as to which outcome will result from a given act. For example, if the staff at your local gelato place is minimally competent, so that you actually get vanilla every time you order vanilla and stracciatella every time you order stracciatella, you are making a choice under certainty. (We will discuss other kinds of choice in future chapters.) Before discussing what it means to make rational choices under conditions of certainty, however, we need to talk about what preferences are and what it means to have rational preferences.

The theory of rational choice under certainty is an **axiomatic** theory. This means that the theory consists of a set of **axioms**: basic propositions that cannot be proven using the resources offered by the theory, and which will simply have to be taken for granted. When studying the theory, the first thing we want to do is examine the axioms. As we go along, we will also introduce new terms by means of definitions. Axioms and definitions have to be memorized. Having introduced the axioms and definitions, we can prove many interesting claims. Thus, much of what we will do below involves proving new propositions on the basis of axioms and definitions.

2.2 Preferences

The concept of **preference** is fundamental in modern economics, neoclassical and behavioral. Formally speaking, a preference is a **relation**. The following are examples of relations: "Alf is older than Betsy," "France is bigger than Norway," and "Bill is worried he may not do as well on the exam as Jennifer." Notice that each of these sentences expresses a relationship between two entities (things, individuals). Thus, "Alf is older than Betsy" expresses a relationship between Alf and Betsy; namely, that the former is older than the latter. Because these examples express a relation between two entities, they are called **binary** relations. The following relation is not binary: "Mom stands between Bill and Bob." This relation is **ternary**, because it involves three different entities; in this case, people.

For convenience, we often use small letters to denote entities or individuals. We may use a to denote Alf and b to denote Betsy. Similarly, we often use capital letters to denote relations. We may use R to denote the relation "is older than." If so, we can write aRb for "Alf is older than Betsy." Sometimes we write Rab. Notice that the order of the terms matters: aRb is not the same thing as bRa. The first says that Alf is older than Betsy and the second that Betsy is older than Alf. Similarly, Rab is not the same thing as Rba.

Exercise 2.1 Relations Assume that f denotes France and n denotes Norway, and that B means "is bigger than."

(a) How would you write that France is bigger than Norway?
(b) How would you write that Norway is bigger than France?
(c) How would you write that Norway is bigger than Norway?

In order to speak clearly about relations, we need to specify what sort of entities may be related to one another. When talking about who is older than whom, we may be talking about people. When talking about what is bigger than what, we may be talking about countries, houses, people, dogs, or many other things. Sometimes it matters what sort of entities we have in mind. When we want to be careful, which is most of the time, we define a **universe** U. The universe is the set of all things that can be related to one another. Suppose we are talking about Donald Duck's nephews Huey, Dewey, and Louie. If so, that is our universe. The convention is to list all members of the universe separated by commas and enclosed in curly brackets, like so: {Huey, Dewey, Louie}. Here, the order does not matter. So, the same universe can be written like this: {Louie, Dewey, Huey}. Thus: U = {Huey, Dewey, Louie} = {Louie, Dewey, Huey}.

Exercise 2.2 The universe Suppose we are talking about all countries that are members of the United Nations. How would that be written?

A universe may have infinitely many members, in which case simple enumeration is inconvenient. This is true; for instance, when you consider the time at which you entered the space where you are reading this. There are infinitely many points in time between 11:59 am and 12:01 pm, for example, as there are between 11:59:59 am and 12:00:01 pm. In such cases, we need to find another way to describe the universe.

One relation we can talk about is this one: "is at least as good as." For example, we might want to say that "coffee is at least as good as tea." The "at least as good as" relation is often expressed using this symbol: $\succcurlyeq$. If c denotes coffee and t denotes tea, we can write the sentence as $c \succcurlyeq t$. This is the **(weak) preference relation**. People may have, and often will have, their own preference relations. If we wish to specify whose preferences we are talking about, we use subscripts to denote individuals. If we want to say that for Alf coffee is at least as good as tea, and that for Betsy tea is at least as good as coffee, we say that $c \succcurlyeq_{Alf} t$ and $t \succcurlyeq_{Betsy} c$, or that $c \succcurlyeq_{A} t$ and $t \succcurlyeq_{B} c$.

Exercise 2.3 Preferences Suppose d denotes "enjoying a cool drink on a hot day" and r denotes "getting roasted over an open fire."

(**a**) How would you state your preference over these two options?

(**b**) How would you express a masochist's preference over these two options?

In economics, we are typically interested in people's preferences over **consumption bundles**, which are collections of goods. You face a choice of commodity bundles when choosing between the #1 Big Burger meal and the #2 Chicken Burger meal at your local hamburger restaurant. In order to represent commodity bundles, we think of them as collections of individual goods along the following lines: three apples and two bananas, or two units of guns and five units of butter. When talking about preference relations, the universe can also be referred to as the **set of alternatives**. If bundles contain no more than two goods, it can be convenient to represent the set of alternatives on a plane, as in Figure 2.1. When bundles contain more than two goods, it is typically more useful to write ⟨3, 2⟩ for three apples and two bananas; ⟨6, 3, 9⟩ for six apples, three bananas, and nine coconuts; and so on.

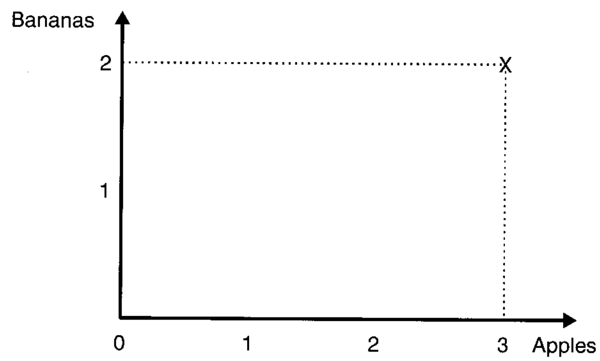

Figure 2.1 Set of alternatives

2.3 Rational preferences

We begin building our theory of rational choice by specifying what it means for a preference relation to be rational. A **rational** preference relation is a preference relation that is transitive and complete.

A relation R is **transitive** just in case the following condition holds: for all x, y, and z in the universe, if x bears relation R to y, and if y bears relation R to z, then x must bear relation R to z. Suppose the universe is the set of all the Marx brothers. If so, "is taller than" is a transitive relation: if Zeppo is taller than Groucho, and Groucho is taller than Harpo, Zeppo must be taller than Harpo (Figure 2.2).

Figure 2.2 The Marx brothers. Illustration by Cody Taylor

Example 2.4 ***30 Rock*** Consider the following exchange from the TV show *30 Rock*. Tracy, Grizz, and Dot Com are playing computer games. Tracy always beats Grizz and Dot Com. When Kenneth beats Tracy but gets beaten by Grizz, Tracy grows suspicious.

Tracy: "How were you beating Kenneth, Grizz?"
Grizz: "I don't know."
Tracy: "If Kenneth could beat me and you can beat Kenneth, then by the transitive property, you should beat me too! Have you been letting me win?"
Dot Com: "Just at some things."
Tracy: "Things? Plural?"

Now you are the first kid on the block who understands *30 Rock*. You also know that the show had a former economics or philosophy student on its staff.

If the universe consists of all people, examples of **intransitive** relations include "is in love with." Just because Sam is in love with Pat, and Pat is in love with Robin, it is not necessarily the case that Sam is in love with Robin. Sam *may* be in love with Robin. But Sam may have no particular feelings about Robin, or Sam may resent Robin for attracting Pat's attention. It may also be the case that Robin is in love with Sam. This kind of intransitivity is central to the play *No Exit*, by the French existentialist philosopher Jean-Paul Sartre. In the play, which takes place in a prison cell, a young woman craves the affection of a man who desires the respect of an older woman, who in turn is in love with the young woman. Hence the most famous line of the play is: "Hell is other people." To show that a relation is intransitive, it is sufficient to identify three members of the universe such that the first is related to the second, and the second is related to the third, but the first is not related to the third.

Formally speaking, a preference relation $\succcurlyeq$ is transitive just in case the following is true:

Axiom 2.5 Transitivity of $\succcurlyeq$ *If* $x \succcurlyeq y$ *and* $y \succcurlyeq z$*, then* $x \succcurlyeq z$ *(for all* x, y, z*).*

There are other ways of expressing the same thing. We might write: If $x \succcurlyeq y \succcurlyeq z$, then $x \succcurlyeq z$ (for all x, y, z). Using standard logic symbols, we might write: $x \succcurlyeq y$ & $y \succcurlyeq z \rightarrow x \succcurlyeq z$ (for all x, y, z). See the text box on page 18 for a useful list of logical symbols. Either way, transitivity says that if you prefer coffee to tea, and tea to root beer, you must prefer coffee to root beer; that is, you cannot prefer coffee to tea and tea to root beer while failing to prefer coffee to root beer.

A relation R is **complete** just in case the following condition holds: for any x and y in the universe, either x bears relation R to y, or y bears relation R to x (or both). If the universe consists of all people – past, present, and future – then "is at least as tall as" is a complete relation. You may not know how tall Caesar and Brutus were, but you do know this: either Caesar was at least as tall as Brutus, or Brutus was at least as tall as Caesar (or both, in case they were equally tall).

Given the universe of all people, examples of **incomplete** relations include "is in love with." For any two randomly selected people – your landlord and the current President of the US, for example – it is not necessarily the case that either one is in love with the other. Your landlord may have a crush on the President, or the other way around. But this need not be the case, and it frequently will not be. To show that a relation is incomplete, then, it is sufficient to identify two objects in the universe such that the relation does not hold either way.

Formally speaking, a preference relation $\succcurlyeq$ is complete just in case the following is true:

Axiom 2.6 Completeness of $\succcurlyeq$ *Either* $x \succcurlyeq y$ *or* $y \succcurlyeq x$ *(or both) (for all* x, y*).*

Completeness means that you must prefer tea to coffee or coffee to tea (or both); though your preference can go both ways, you cannot fail to have a preference between the two. The use of the phrase "(or both)" in the formula above is, strictly speaking, redundant: we use the "inclusive or," which is equivalent to "and/or" in everyday language. Using standard logical symbols, we might write: $x \succcurlyeq y \lor y \succcurlyeq x$ (for all x, y). If both $x \succcurlyeq y$ and $y \succcurlyeq x$, we say that there is a tie (see Section 2.4).

The following exercise serves to illustrate the concepts of transitivity and completeness.

Exercise 2.7 Assuming the universe is the set of all people – past, present, and future – are the following relations transitive? Are they complete?

(a) "is the mother of"

(b) "is an ancestor of"

(c) "is the sister of"
(d) "detests"
(e) "weighs more than"
(f) "has the same first name as"
(g) "is taller than"

When answering questions such as these, ambiguity can be a problem. A word like "sister" is ambiguous, which means that answers might depend on how it is used. As soon as the word is defined, however, the questions have determinate answers.

Exercise 2.8 The enemy of your enemy Suppose it is true, as people say, that the enemy of your enemy is your friend. What does this mean for the transitivity of "is the enemy of"? (Assume there are no true frenemies: people who are simultaneously friends and enemies.)

Exercise 2.9 Assuming the universe is the set of all natural numbers, meaning that U = {1, 2, 3, 4, ...}, are the following relations transitive? Are they complete?
(a) "is at least as great as" ($\geq$)
(b) "is equal to" ($=$)
(c) "is strictly greater than" ($>$)
(d) "is divisible by" ($|$)

Exercise 2.10 Preferences and the universe Use your understanding of transitivity and completeness to answer the following questions.
(a) If the universe is {apple, banana, starvation}, what does the transitivity of the preference relation entail?
(b) If the universe is {apple, banana}, what does the completeness of the preference relation entail?

Logical symbols

Here is a list of the most common logical symbols:

$x \& y$	x and y
$x \vee y$	x or y
$x \rightarrow y$	if x then y; x only if y
$x \leftrightarrow y$	x if and only if y; x just in case y
$\neg p$	not p

As the last exercise suggests, the completeness of the preference relation implies that it is **reflexive**, meaning that $x \succcurlyeq x$ (for all x). This result might strike you as surprising. But recall that completeness says that any time you pick two elements from the universe, the relation must hold one way or the other. The axiom does not say that the two elements must be different. If you pick the same element twice, which you may, completeness requires that the thing stands in the relation to itself.

The choice of a universe might determine whether a relation is transitive or intransitive, complete or incomplete. If the universe were U = {Romeo, Juliet}, the relation "is in love with" would be complete, since for any two members of the universe, either the one is in love with the other, or the other is in love with the one. (This assumes that Romeo and Juliet are both in love with themselves, which might perhaps not be true.) Perhaps more surprisingly, the relation would also be transitive: whenever $x \succcurlyeq y$ and $y \succcurlyeq z$, it is in fact the case that $x \succcurlyeq z$.

The assumption that the weak preference relation is rational (transitive and complete) might seem fairly modest. Yet, in combination with a couple of definitions, this assumption is in effect everything necessary to build a theory of choice under certainty. This is a wonderful illustration of how science works: based on a small number of assumptions, we will build an extensive theory, whose predictions will then be confronted with actual evidence. The rest of this chapter spells out the implications of the assumption that the weak preference relation is rational.

2.4 Indifference and strict preference

As the previous section shows, the (weak) preference relation admits ties. When two options are tied, we say that the first option is **as good as** the second or that the agent is **indifferent** between the two options. That is, a person is indifferent between two options just in case, to her, the first option is at least as good as the second and the second is at least as good as the first. We use the symbol $\sim$ to denote indifference. Formally speaking:

Definition 2.11 Definition of indifference *$x \sim y$ if and only if $x \succcurlyeq y$ and $y \succcurlyeq x$.*

Using logical symbols, we might write: $x \sim y \Leftrightarrow x \succcurlyeq y \;\&\; y \succcurlyeq x$.

Assuming that the "at least as good as" relation is rational, the indifference relation is both reflexive and transitive. It is also **symmetric**: if x is as good as y, then y is as good as x. These results are not just intuitively plausible; they can be established by means of **proofs**. (See the text box on page 23 for more about proofs.) Properties of the indifference relation are established by the following proposition.

Proposition 2.12 Properties of indifference *The following conditions hold:*

(i) $x \sim x$ *(for all x)*
(ii) $x \sim y \rightarrow y \sim x$ *(for all x, y)*
(iii) $x \sim y \;\&\; y \sim z \rightarrow x \sim z$ *(for all x, y, z)*

Proof. Each part of the proposition requires a separate proof:

(i)	1. $x \succcurlyeq x$	by Axiom 2.6
	2. $x \succcurlyeq x$ & $x \succcurlyeq x$	from (1), by logic
	$\therefore x \sim x$	from (2), by Definition 2.11
(ii)	1. $x \sim y$	by assumption
	2. $x \succcurlyeq y$ & $y \succcurlyeq x$	from (1), by Definition 2.11
	3. $y \succcurlyeq x$ & $x \succcurlyeq y$	from (2), by logic
	4. $y \sim x$	from (3), by Definition 2.11
	$\therefore x \sim y \rightarrow y \sim x$	from (1)–(4), by logic
(iii)	1. $x \sim y$ & $y \sim z$	by assumption
	2. $x \succcurlyeq y$ & $y \succcurlyeq x$	from (1), by Definition 2.11
	3. $y \succcurlyeq z$ & $z \succcurlyeq y$	from (1), by Definition 2.11
	4. $x \succcurlyeq z$	from (2) and (3), by Axiom 2.5
	5. $z \succcurlyeq x$	from (2) and (3), by Axiom 2.5
	6. $x \sim z$	from (4) and (5), by Definition 2.11
	$\therefore x \sim y$ & $y \sim z \rightarrow x \sim z$	from (1)–(6), by logic □

These are the complete proofs. In what follows, I will often outline the general shape of the proof rather than presenting the whole thing.

The indifference relation is not complete. To show this, it is enough to give a single counterexample. Any rational preference relation according to which the agent is not indifferent between all options will do (see, for instance, Figure 2.3).

Exercise 2.13 Prove the following principle: $x \succcurlyeq y$ & $y \sim z \rightarrow x \succcurlyeq z$.

In your various proofs, it is always acceptable to rely on propositions you have already established. The following exercise shows how useful this can be.

Exercise 2.14 Iterated transitivity In this exercise you will prove the following principle in two different ways: $x \sim y$ & $y \sim z$ & $z \sim p \rightarrow x \sim p$.

(a) First prove it by applying the transitivity of indifference (Proposition 2.12(iii)).

(b) Then prove it without assuming the transitivity of indifference. (You may still use the transitivity of weak preference, since it is an axiom.)

If you have difficulty completing the proofs, refer to the text box on page 23 for hints.

When a first option is at least as good as a second, but the second is not at least as good as the first, we say that the first option is **better than** the second or that the agent **strictly** or **strongly** prefers the first over the second. We use the symbol $\succ$ to denote **strict** or **strong preference**. Formally speaking:

Definition 2.15 Definition of strict preference *$x \succ y$ if and only if $x \succcurlyeq y$ and it is not the case that $y \succcurlyeq x$.*

Using logical notation, that is to say: $x \succ y \Leftrightarrow x \succcurlyeq y \,\&\, \neg y \succcurlyeq x$. For clarity, sometimes the "is at least as good as" relation will be called **weak preference.**

Assuming (still) that the weak preference relation is rational, it is possible to prove logically that the relation will have certain properties. The following proposition establishes some of them.

Proposition 2.16 Properties of strict preference *The following conditions hold:*

(i) $x \succ y \,\&\, y \succ z \rightarrow x \succ z$ *(for all x, y, z)*
(ii) $x \succ y \rightarrow \neg y \succ x$ *(for all x, y)*
(iii) *not* $x \succ x$ *(for all x)*

Proof. (i) Suppose that $x \succ y$ & $y \succ z$. In order to establish that $x \succ z$, Definition 2.15 tells us that we need to show that $x \succcurlyeq z$ and that it is not the case that $z \succcurlyeq x$. The first part is Exercise 2.17. The second part goes as follows: suppose for a **proof by contradiction** that $z \succcurlyeq x$. From the first assumption and the definition of strict preference, it follows that $x \succcurlyeq y$. From the second assumption and Axiom 2.5, it follows that $z \succcurlyeq y$. But from the first assumption and the definition of strict preference, it also follows that $\neg z \succcurlyeq y$. We have derived a contradiction, so the second assumption must be false, and therefore $\neg z \succcurlyeq x$.

(ii) Begin by assuming $x \succ y$. Then, for a proof by contradiction, assume that $y \succ x$. Given the first assumption, Definition 2.15 implies that $x \succcurlyeq y$. Given the second assumption, the same definition implies that $\neg x \succcurlyeq y$. But this is a contradiction, so the second assumption must be false, and therefore $\neg y \succ x$.

(iii) See Exercise 2.19. □

Proposition 2.16(i) says that the strict preference relation is transitive; 2.16(ii) that it is **anti-symmetric**; and 2.16(iii) that it is **irreflexive**.

Exercise 2.17 Using the definitions and propositions discussed so far, complete the first part of the proof of Proposition 2.16(i).

Notice that the proofs of Proposition 2.16(i) and (ii) involve constructing proofs by contradiction. Such proofs are also called **indirect proofs**. This mode of reasoning might look weird, but it is actually quite common in mathematics, science, and everyday thinking. For example, when mathematicians prove that $\sqrt{2}$ is an irrational number, they can proceed by assuming (for a proof by contradiction) that $\sqrt{2}$ is a rational number (meaning that $\sqrt{2}$ can be expressed as a fraction p/q of natural numbers p and q) and then use this assumption to derive a contradiction.

Exercise 2.18 The enemy of your enemy, cont. Use a proof by contradiction to establish that "is the enemy of" is not transitive, as in Exercise 2.8 above.

In future exercises, you will see just how useful proofs by contradiction can be.

Exercise 2.19 Prove Proposition 2.16(iii). Prove it by contradiction, by first assuming that there is an x such that $x \succ x$.

Exercise 2.20 Prove the following principle: $x \succ y$ & $y \succcurlyeq z \rightarrow x \succ z$ (for all x, y, z). Notice that this proof has two parts. First, prove that $x \succcurlyeq z$; second, prove that $\neg z \succcurlyeq x$.

Exercise 2.21 Establish the following important and intuitive principles. (For the record, some of them are logically equivalent.)

(a) If $x \succ y$ then $x \succcurlyeq y$
(b) If $x \succ y$ then $\neg y \succcurlyeq x$
(c) If $x \succcurlyeq y$ then $\neg y \succ x$
(d) If $x \succ y$ then $\neg x \sim y$
(e) If $x \sim y$ then $\neg x \succ y$
(f) If $\neg x \succcurlyeq y$ then $y \succcurlyeq x$
(g) If $\neg x \succcurlyeq y$ then $y \succ x$
(h) If $\neg x \succ y$ then $y \succcurlyeq x$

If you run into trouble with parts (f) and (g), note that you can always play the completeness card and throw in the expression $x \succcurlyeq y \vee y \succcurlyeq x$ any time. Also note that $p \vee q$ and $\neg p$ implies that q. If you find part (h) difficult, feel free to invoke the principle known as **de Morgan's law**, according to which $\neg(p \;\&\; q)$ is logically equivalent to $\neg p \vee \neg q$. Also note that $p \vee q$ and $p \rightarrow q$ implies that q.

For the next exercise, recall that it is acceptable to rely on propositions already established (see text box below).

Exercise 2.22 Prove that if $x \sim y$ and $y \sim z$, then $\neg x \succ z$.

Exercise 2.23 Negative transitivity Prove the following two principles. You might already have been tempted to invoke these two in your proofs. But remember that you may not do so before you have established them.

(a) If $\neg x \succcurlyeq y$ and $\neg y \succcurlyeq z$, then $\neg x \succcurlyeq z$
(b) If $\neg x \succ y$ and $\neg y \succ z$, then $\neg x \succ z$

The last two exercises illustrate some potentially problematic implications of the theory that we have studied in this chapter. Both are classics.

Exercise 2.24 Vacations Suppose that you are offered two vacation packages, one to California and one to Florida, and that you are perfectly indifferent between the two. Let us call the Florida package f and the California package c. So f $\sim$ c. Now, somebody improves the Florida package by adding an apple to it. You like apples, so the enhanced Florida package f^+ improves the original Florida package, meaning that $f^+ \succ f$. Assuming that you are rational, how do you feel about the enhanced Florida package f^+ compared to the California package c? Prove it.

How to do proofs

The aim of a **proof** of a proposition is to establish the truth of the proposition with logical or mathematical certainty (see the proofs of Proposition 2.12(i)–(iii) for examples). A proof is a sequence of propositions, presented on separate lines of the page. The last line of the proof is the proposition you intend to establish, that is, the conclusion; the lines that come before it establish its truth. The conclusion is typically preceded by the symbol ∴. All other lines are numbered using Arabic numerals. The basic rule is that each proposition in the proof must follow logically from (a) a proposition on a line above it, (b) an axiom of the theory, (c) a definition that has been properly introduced, and/or (d) a proposition that has already been established by means of another proof. Once a proof is concluded, logicians like to write "QED" – Latin for "quod erat demonstrandum," meaning "that which was to be shown" – or by the little box □.

There are some useful hints, or rules of thumb, that you may want to follow when constructing proofs. **Hint one**: if you want to establish a proposition of the form $x \rightarrow y$, you typically want to begin by assuming what is to the left of the arrow; that is, the first line will read "1. x by assumption." Then, your goal is to derive y, which would permit you to complete the proof. If you want to establish a proposition of the form $x \leftrightarrow y$, you need to do it both ways: first, prove that $x \rightarrow y$, second, that $y \rightarrow x$. **Hint two**: if you want to establish a proposition of the form $\neg p$, you typically want to begin by assuming the opposite of what you want to prove for a proof by contradiction; that is, the first line would read "1. p by assumption for a proof by contradiction." Then, your goal is to derive a contradiction, that is, a claim of the form q & $\neg q$, which would permit you to complete the proof.

Exercise 2.25 Tea cups Imagine that there are 1000 cups of tea lined up in front of you. The cups are identical except for one difference: the cup to the far left (c_1) contains one grain of sugar, the second from the left (c_2) contains two grains of sugar, the third from the left (c_3) contains three grains of sugar, and so on. Since you cannot tell the difference between any two adjacent cups, you are indifferent between c_n and c_{n+1} for all n between 1 and 999 inclusive. Assuming that your preference relation is rational, what is your preference between the cup to the far left (c_1) and the one to the far right (c_{1000})?

Your findings from Exercise 2.21 are likely to come in handy when answering these questions.

2.5 Preference orderings

The preference relation is often referred to as a **preference ordering.** This is so because a rational preference relation allows us to order all alternatives in a list, with the best at the top and the worst at the bottom. Figure 2.3 shows an example of a preference ordering.

A rational preference ordering is simple. Completeness ensures that each person will have exactly one list, because completeness entails that each element can be compared to all other elements. Transitivity ensures that the list will be linear, because transitivity entails that the strict preference relation will never have cycles, as when $x \succ y$, $y \succ z$, and $z \succ x$. Here are two helpful exercises about cycling preferences.

Exercise 2.26 Cycling preferences Using the definitions and propositions discussed so far, show that it is impossible for a rational strict preference relation to cycle. To do so, suppose (for the sake of the argument) that $x \succ y$ & $y \succ z$ & $z \succ x$ and show that this leads to a contradiction.

Exercise 2.27 Cycling preferences, cont. By contrast, it is possible for the weak preference relation to cycle. This is to say that there may well be an x, y, and z such that $x \succcurlyeq y$ & $y \succcurlyeq z$ & $z \succcurlyeq x$. If this is so, what do we know about the agent's preferences over x, y, and z? Prove it.

In cases of indifference, the preference ordering will have ties. As you may have noticed, Figure 2.3 describes a preference ordering in which two items are equally good. Assuming that the universe is {Heavenly Bliss, Coke, Pepsi, Eternal Suffering}, this preference ordering is perfectly rational.

In economics, preference orderings are frequently represented using **indifference curves**, also called **indifference maps**. See Figure 2.4 for an example of a set of indifference curves. You can think of these as analogous to contour lines on a topographic map. By convention, each bundle on one of these curves is as good as every other bundle on the same curve. When two bundles are on different curves, one of the two bundles is strictly preferred to the other. Insofar as people prefer more of each good to less, bundles on curves to the top right will be strictly preferred to bundles on curves to the bottom left.

Exercise 2.28 Indifference curves Represent the following sets of indifference curves graphically.

(**a**) Suppose that an apple for you is always as good as two bananas.

Heavenly Bliss

$\succ$

Coke $\sim$ Pepsi

$\succ$

Eternal Suffering

Figure 2.3 Preference ordering with tie

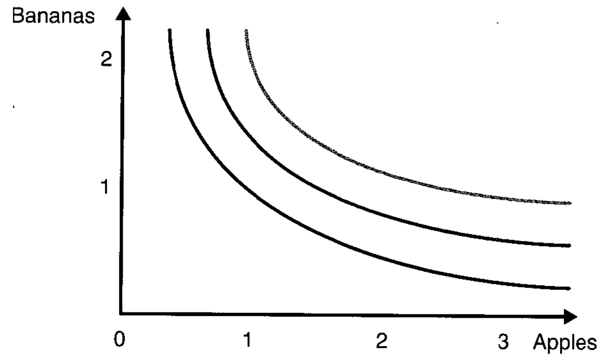

Figure 2.4 Indifference curves

(b) Suppose that one apple is always as good, as far as you are concerned, as a banana.

(c) Suppose that you do not care for tea without milk or for milk without tea. However, every time you have two units of tea and one unit of milk, you can make yourself a cup of tea with milk. You love tea with milk, and the more the better, as far as you are concerned.

2.6 Choice under certainty

To make a **choice under certainty** is to face a menu. A **menu** is a set of options such that you have to choose exactly one option from the set. This is to say that the menu has two properties. First, the items in the menu are **mutually exclusive**; that is, you can choose at most one of them at any given time. Second, the items in the menu are **exhaustive**; that is, you have to choose at least one of them.

Example 2.29 The menu If a restaurant offers two appetizers (soup and salad) and two entrées (chicken and beef) and you must choose one appetizer and one entrée, what is your set of alternatives?

Since there are four possible combinations, your set of alternatives is {soup-and-chicken, soup-and-beef, salad-and-chicken, salad-and-beef}.

Exercise 2.30 The menu, cont. If you can also choose to eat an appetizer only, or an entrée only, or nothing at all, what would the new menu be?

There is no assumption that a menu is small, or even finite, though we frequently assume that it is.

In economics, the menu is often referred to as the **budget set**. This is simply that part of the set of alternatives that you can afford given your budget, that is, your resources at hand. Suppose that you can afford at most three apples (if you buy no bananas) or two bananas (if you buy no apples). This would be the case, for instance, if you had $6 in your pocket and bananas cost $3 and apples $2. If so, your budget set – or your menu – is represented by the shaded area in Figure 2.5. Assuming that fruit is infinitely divisible, the menu is infinitely large. The line separating the items in your budget from the items outside of it is called the **budget line**.

Exercise 2.31 Budget sets Suppose that your budget is $12. Use a graph to answer the following questions:

(a) What is the budget set when apples cost $3 and bananas cost $4?
(b) What is the budget set when apples cost $6 and bananas cost $2?
(c) What is the budget set when apples always cost $2, the first banana costs $4, and every subsequent banana costs $2?

So what does it mean **to be rational**? To be rational, or **to make rational choices**, means (i) that you have a rational preference ordering, and (ii) that whenever you are faced with a menu, you choose the most preferred item, or (in the case of ties) one of the most preferred items. The second condition can also be expressed as follows: (ii′) that … you choose an item such that no other item in the menu is strictly preferred to it. Or like this: (ii″) that … you do not choose an item that is strictly less preferred to another item in the menu. *This is all we mean when we say that somebody is rational.* If you have the preferences of Figure 2.3 and are facing a menu offering Coke, Pepsi, and Eternal Suffering, the rational choice is to pick either the Coke or the Pepsi option. When there is no unique best choice, as in this case, the theory says that you have to choose one of the best options; it does not specify which one.

The rational decision can be determined if we know the agent's indifference curves and budget set. If you superimpose the former (from Figure 2.4) onto the latter (from Figure 2.5), you get a picture like Figure 2.6. The consumer will choose the bundle marked X, because it is the most highly preferred bundle in the budget set. As you can tell, there is no more highly preferred bundle in the budget set.

It is important to note what the theory of rationality does *not* say. The theory does not say why people prefer certain things to others, or why they

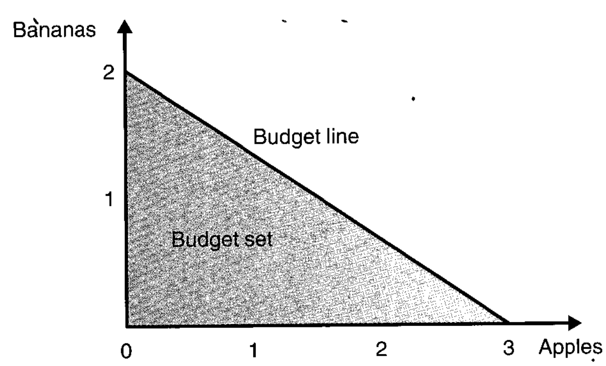

Figure 2.5 Budget set

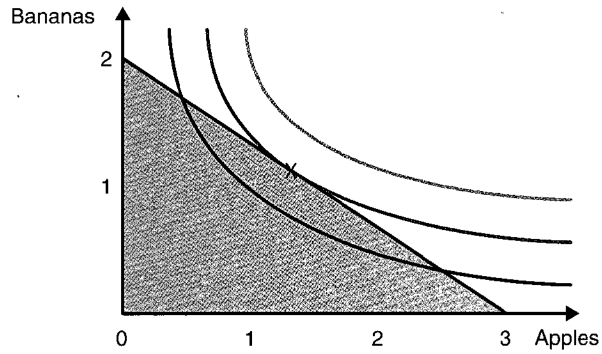

Figure 2.6 Consumer choice problem

choose so as to satisfy their preferences. It does not say that people prefer apples to bananas because they think that they will be happier, feel better, or be more satisfied, if they get apples than if they get bananas (although that may, in fact, be the case). This theory says nothing about feelings, emotions, moods, or any other subjectively experienced state. As far as this theory is concerned, the fact that you prefer a cool drink on a hot day to being roasted over an open fire is just a brute fact; it is not a fact that needs to be grounded in an account of what feels good or bad, pleasant or unpleasant, rewarding or aversive. Similarly, the theory does not say why people choose the most preferred item on the menu; as far as this theory is concerned, they just do.

Moreover, the theory does not say that people are selfish, in the sense that they care only about themselves; or that they are materialistic, in the sense that they care only about material goods; or that they are greedy, in the sense that they care only about money. The definition of rationality implies that a rational person is self-interested, in the sense that her choices reflect her own preference ordering rather than somebody else's. But this is not the same as being selfish: the rational individual may, for example, prefer dying for a just cause over getting rich by defrauding others. The theory in itself specifies only some formal properties of the preference relation; it does not say anything about the things people prefer. The theory is silent about whether or not they pursue respectable and moral ends. Rational people may be weird, evil, sadistic, selfish, and morally repugnant, or saintly, inspiring, thoughtful, selfless, and morally admirable; they can act out of compulsion, habit, feeling, or as a result of machine-like computation. This conception of rationality has a long and distinguished history. The Scottish eighteenth-century philosopher and economist David Hume wrote: "'Tis not contrary to reason to prefer the destruction of the whole world to the scratching of my finger. 'Tis not contrary to reason for me to choose my total ruin, to prevent the least uneasiness of an Indian or person wholly unknown to me. 'Tis as little contrary to reason to prefer even my own acknowledge'd lesser good to my greater, and have a more ardent affection for the former than the latter."

Rational people cannot have preferences that are intransitive or incomplete, and they cannot make choices that fail to reflect those preferences.

2.7 Utility

The notion of **utility**, which is central to modern economics, has generated a great deal of confusion. It is worth going slowly here. Suppose that you want to use numbers to express how much a person prefers something, how would you do it? One solution is obvious. Remember that a rational person's preferences allow us to arrange all alternatives in order of preference. Consider, for example, the preference ordering in Figure 2.3. The preference ordering has three "steps." In order to represent these preferences by numbers, we assign one number to each step, in such a way that higher steps are associated with higher numbers. See Figure 2.7 for an example.

A **utility function** associates a number with each member of the set of alternatives. In this case, we have associated the number 3 with Heavenly Bliss (HB). That number is called the utility of HB and is denoted u(HB). In this case, u(HB) = 3. The number associated with Eternal Suffering (ES) is called the utility of ES and is written u(ES). In this case, u(ES) = 1. If we use C to denote Coke and P to denote Pepsi, $u(C) = u(P) = 2$. Because we designed the utility function so that higher utilities correspond to more preferred items, we say that the utility function $u(\cdot)$ **represents** the preference relation $\succcurlyeq$.

As the example suggests, two conditions must hold in order for something to be a utility function. First, it must be a function (or a mapping) from the set of alternatives into the set of real numbers. This means that every alternative gets assigned exactly one number. If Figure 2.7 had empty spaces in the right-hand column, or if the figure had several numbers in the same cell, we would not have a proper utility function. While the utility function needs to assign some number to every alternative, it is acceptable (as the example shows) to assign the same number to several alternatives. Second, for something to be a utility function, it must assign larger numbers to more preferred alternatives; that is, if x is at least as good as y, the number assigned to x must be greater than or equal to the number assigned to y. To put it more formally:

Definition 2.32 Definition of u(·) *A function $u(\cdot)$ from the set of alternatives into the set of real numbers is a utility function representing preference relation $\succcurlyeq$ just in case $x \succcurlyeq y \Leftrightarrow u(x) \geq u(y)$ (for all x and y).*

A function $u(\cdot)$ that satisfies this condition can be said to be an **index** or a **measure** of preference relation $\succcurlyeq$. Historically, the word "utility" has been used to refer to many different things, including the pleasure, happiness, and

Heavenly Bliss — 3

$\curlyvee$

Coke ~ Pepsi — 2

$\curlyvee$

Eternal Suffering — 1

Figure 2.7 Preference ordering with utility function

satisfaction of receiving, owning, or consuming something. Though most people (including economics professors) find it hard to stop speaking in this way, as though utility is somehow floating around "in your head," this usage is archaic. Utility is nothing but an index or measure of preference.

Given a rational preference relation, you may ask whether it is always possible to find a utility function that represents it. When the set of alternatives is finite, the answer is yes. The question is answered by means of a so-called **representation theorem.**

Proposition 2.33 Representation theorem *If the set of alternatives is finite, then $\succcurlyeq$ is a rational preference relation just in case there exists a utility function representing $\succcurlyeq$.*

Proof. Omitted. □

When the set of alternatives is infinite, representing preference relations gets more complicated. It remains true that if a utility function represents a preference relation, then the preference relation is rational. However, even if the preference relation is rational, it is not always possible to find a utility function that represents it.

As you may suspect, a utility function will associate strictly higher numbers with strictly preferred alternatives, and equal numbers with equally preferred alternatives. That is, the following proposition is true:

Proposition 2.34 Properties of u(·) *Given a utility function $u(\cdot)$ representing preference relation $\succcurlyeq$, the following conditions hold:*

(i) $x \succ y \Leftrightarrow u(x) > u(y)$

(ii) $x \sim y \Leftrightarrow u(x) = u(y)$

Proof. (i) First, assume that $x \succ y$, so that $x \succcurlyeq y$ and $\neg y \succcurlyeq x$. Using Definition 2.32 twice, we can infer that $u(x) \geq u(y)$ and that not $u(y) \geq u(x)$. Simple math tells us that $u(x) > u(y)$. Second, assume that $u(x) > u(y)$, which implies that $u(x) \geq u(y)$ and that not $u(y) \geq u(x)$. Using Definition 2.32 twice, we can infer that $x \succcurlyeq y$ and $\neg y \succcurlyeq x$, which in turn implies that $x \succ y$.

(ii) See Exercise 2.35. □

Recall (from the text box on page 23) that if you want to prove something of the form $A \Leftrightarrow B$, your proof must have two parts.

Exercise 2.35 Prove Proposition 2.34(ii).

It is easy to confirm that the proposition is true of the utility function from Figure 2.7.

One important point to note is that utility functions are not unique. The sequence of numbers ⟨1, 2, 3⟩ in Figure 2.7 could have been chosen very differently. The sequence ⟨0, 1, 323⟩ would have done as well, as would ⟨−1000, −2, 0⟩ and ⟨$-\pi$, e, 1077⟩. All these are utility functions, in that they associate higher numbers with more preferred options. As these examples show, it is important not to ascribe any significance to absolute numbers. To know that the utility I derive from listening to Justin Bieber is 2 tells you *absolutely nothing* about my preferences. But if you know that the utility I derive from listening to Arcade Fire is 4, you know something, namely, that I strictly prefer Arcade Fire to Justin Bieber. It is equally important not to ascribe any significance to ratios of utilities. Even if the utility of Arcade Fire is twice the utility of Justin Bieber, this does not mean that I like Arcade Fire "twice as much." The same preferences could be represented by the numbers 0 and 42, in which case the ratio would not even be well defined. In brief, for every given preference relation, there are many utility functions representing it. Utility as used in this chapter is often called **ordinal utility**, because all it does is allow you to order things.

How do utilities relate to indifference curves? A utility function in effect assigns one number to each indifference curve, as in Figure 2.8. This way, two bundles that fall on the same curve will be associated with the same utility, as they should be. Two bundles that fall on different curves will be associated with different utilities, again as they should be. Of course, higher numbers will correspond to curves that are more strongly preferred. For a person who likes apples and bananas, $u_1 < u_2 < u_3$.

How does utility relate to behavior? Remember that you choose rationally insofar as you choose the most preferred item (or one of the most preferred items) on the menu. The most preferred item on the menu will also be the item with the highest utility. So to choose the most preferred item is to choose the item with the highest utility. Now, **to maximize utility** is to choose the item with the highest utility. Thus, you choose rationally insofar as you maximize utility. Hence, *to maximize utility is to choose rationally*. Notice that you can maximize utility in this sense without necessarily going through any particular calculations; that is, you do not need to be able to solve mathematical maximization problems in order to maximize utility. Similarly, you can maximize utility without maximizing feelings of pleasure, satisfaction, contentment, happiness, or whatever; utility (like preference) still has nothing to do with subjectively experienced states of any kind. This is a source of endless confusion.

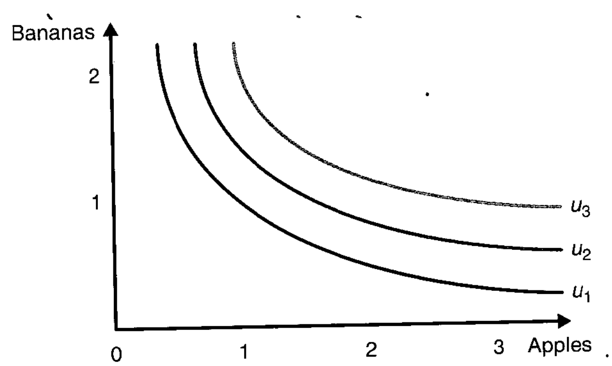

Figure 2.8 Indifference curves and utility

A final word about proofs

While the proofs discussed in this chapter may at first blush seem intimidating, notice that the basic principles are fairly simple. So far, we have introduced only two axioms, namely, the transitivity of the weak preference relation (Axiom 2.5 on page 17) and the completeness of the weak preference relation (Axiom 2.6 on page 17); three definitions, namely, the definition of indifference (Definition 2.11 on page 19), the definition of strict preference (Definition 2.15 on page 20), and the definition of utility (Definition 2.32 on page 28); and two hints (see text box on page 23). In order to complete a proof, there are only seven things that you need to know.

2.8 Discussion

The first thing to notice is how much mileage you can get out of a small number of relatively weak assumptions. Recall that we have made only two assumptions: that preferences are rational and that people choose so as to satisfy their preferences. As long as the two assumptions are true, and the set of alternatives is not too large, we can define the concept of utility and make sense of the idea of utility maximization. That is the whole theory. The second thing to notice is what the theory does not say. The theory does not say that people are selfish, materialistic, or greedy; it says nothing about why people prefer one thing over another; it does not presuppose that people solve mathematical maximization problems in their heads; and it makes no reference to things like pleasure, satisfaction, and happiness. The fact that the theory is relatively non-committal helps explain why so many economists are comfortable using it: after all, the theory is compatible with a great deal of behavior.

Though brief, this discussion sheds light on the nature of economics, as some economists see it. Nobel laureate Gary Becker defines the economic approach to behavior in terms of three features: "The combined assumptions of maximizing behavior, market equilibrium, and stable preferences, used relentlessly and unflinchingly, form the heart of the economic approach as I see it." Because this text is about individual choice, I have little to say about market equilibrium. However, what Becker has in mind when he talks about maximizing behavior and stable preferences should be eminently clear from what has already been said. In this analysis, preferences are **stable** in the sense that they are not permitted to change over time.

Exercise 2.36 Misguided criticism Many criticisms of standard economics are quite mistaken. Explain where the following critics go wrong.

(a) An otherwise illuminating article about behavioral economics in *Harvard Magazine* asserts that "the standard model of the human actor – Economic Man – that classical and neoclassical economics have used as a foundation for decades, if not centuries ... is an intelligent, analytic, selfish creature."

(b) A common line of criticism of standard economics begins with some claim of the form "the most fundamental idea in economics is that money makes people happy" and proceeds to argue that the idea is false.

Is this a plausible theory of human behavior under conditions of certainty? To answer this question we need to separate the descriptive from the normative question. The first question is whether the theory is descriptively adequate, that is, whether people's choices *do as a matter of fact* reflect a rational preference ordering. This is the same as asking whether people maximize utility. Though both transitivity and completeness may seem obviously true of people's preferences, there are many cases in which they do not seem to hold: a person's preference relation over prospective spouses, for example, is unlikely to be complete. The second question is whether the theory is normatively correct, that is, whether people's choices *should* reflect a rational preference ordering. This is the same as asking whether people *should* maximize utility. Though transitivity and completeness may seem rationally required, it can be argued that they are neither necessary nor sufficient for being rational.

Next, we explore what happens when the theory is confronted with data.

ADDITIONAL EXERCISES

Exercise 2.37 For each of the relations and properties in Table 2.1, use a check mark to identify whether or not the relation has the property.

Table 2.1 Properties of weak preference, indifference, and strong preference

	Property	Definition	$\succsim$	$\sim$	$\succ$
(a)	Transitivity	xRy & $yRz \rightarrow xRz$ (for all x, y, z)			
(b)	Completeness	$xRy \vee yRx$ (for all x, y)			
(c)	Reflexivity	xRx (for all x)			
(d)	Irreflexivity	$\neg xRx$ (for all x)			
(e)	Symmetry	$xRy \rightarrow yRx$ (for all x, y)			
(f)	Anti-symmetry	$xRy \rightarrow \neg yRx$ (for all x, y)			

Exercise 2.38 More properties of the preference relation Here are two relations: "is married to" and "is not married to." Supposing the universe is the set of all living human beings, which of these is...

(a) reflexive
(b) irreflexive
(c) symmetric
(d) asymmetric
(e) anti-symmetric

ADDITIONAL EXERCISES cont.

Exercise 2.39 As part of your answer to the following questions, make sure to specify what the universe is.

(a) Give an example of a relation that is complete but not transitive.

(b) Give an example of a relation that is transitive but not complete.

Exercise 2.40 Irrationality Explain (in words) why each of the characters below is irrational according to the theory you have learned in this chapter.

(a) In the drama *Sophie's Choice*, the title character finds herself in a Nazi concentration camp and must choose which one of her children is to be put to death. She is not indifferent and cannot form a weak preference either way.

(b) An economics professor finds that he prefers a $10 bottle of wine to a $8 bottle, a $12 bottle to a $10 bottle, and so on; yet he does not prefer a $200 bottle to a $8 bottle.

(c) Buridan's ass is as hungry as it is thirsty and finds itself exactly midway between a stack of hay and a pail of water. Unable to decide which is better, the animal starves to death.

FURTHER READING

A nontechnical introduction to decision theory is Allingham (2002). More technical accounts can be found in Mas-Colell et al. (1995, Chapters 1–2). The paragraph from David Hume comes from Hume (2000 [1739–40], p. 267). The Becker quotation is from Becker (1976, p. 5). The *Harvard Magazine* article is Lambert (2006) and the critics talking about happiness Dutt and Radcliff (2009, p. 8).