Preferences, Choice, and Utility

An Introduction to Theory and Experiments

Felix Muñoz-García

School of Economic Sciences
Washington State University

August 18, 2026

For classroom use only – not for distribution.

Contents

Preface

These are lecture notes for the first weeks of EconS 501 at Washington State University. They cover the foundations of choice theory – preference relations and the rationality axioms of completeness and transitivity, the classical utility representation theorems (Debreu, Rader), the choice-theoretic counterpart via WARP, SARP, and GARP with Afriat's theorem, the three special preference structures that recur in applied work (homothetic, quasi-linear, additively separable), the utility-maximization problem and Marshallian demand, and a compact summary of the experimental evidence. Real-analysis prerequisites – topology on $\mathbb{R}^L$, compactness, Weierstrass, Berge, and the envelope and implicit function theorems – are cited inline to the Mathematical Appendix (Section M) of Mas-Colell, Whinston, and Green (1995), *Microeconomic Theory*, and to Simon and Blume (1994), *Mathematics for Economists*. The material is intentionally compact and follows the order of the lectures.

Felix Muñoz-García

Pullman, Washington
August 18, 2026

Chapter 1 - Preference relations

1.1 Introduction

This chapter builds the vocabulary that the rest of the book depends on. A **consumption set** X is the set of alternatives the decision-maker can imagine choosing from. A **preference relation** on X is a binary relation $\succsim$ that says, for every pair of alternatives $x, y \in X$, whether the decision-maker weakly prefers x to y. Everything downstream – the utility functions that Chapter 3 constructs, the choice rules that Chapter 4 observes, the demand functions that Chapter 6 derives – is built on this primitive.

Chapter 1 introduces the primitive itself and layers on the two axioms that make it *rational*: completeness (every pair is comparable) and transitivity (rankings do not cycle). These two axioms are what economists mean when they say an agent is rational. They are also less innocuous than they look: Chapter 7 documents laboratory subjects who violate them in predictable ways.

Section 1.2 fixes the notation and defines the consumption set. Section 1.3 defines the weak-preference relation $\succsim$ and its derived relations, the strict preference $\succ$ and the indifference relation $\sim$. Section 1.4 states the rationality axioms of completeness and transitivity and works through their immediate consequences. Section 1.5 defines an **indifference class** and shows that a rational preference relation partitions X into indifference classes. Section 1.6 treats the special case in which X is finite, where completeness and transitivity alone are enough to guarantee a utility representation (a first taste of the deeper representation theorem of Chapter 3). Section 1.7 gives the leading example of a preference relation that is complete and transitive but resists any graphical description: the **lexicographic** preference relation, which will resurface in Chapter 3 as the canonical counterexample to continuous utility representation. Section 1.8 introduces experimental evidence on completeness and transitivity. Section 1.9 summarizes.

1.2 The consumption set

The **consumption set** is the set X of alternatives the decision-maker can imagine choosing from. It plays two roles at once. It is the domain over which preferences are defined, and it is the domain over which utility functions (when they exist) are defined. Every axiom, every theorem, and every example in this book has an implicit "over the consumption set X" attached.

In the leading application – consumer choice over L goods – $X = \mathbb{R}^L_+$, the non-negative orthant of L-dimensional Euclidean space. A bundle $x = (x_1, \ldots, x_L) \in \mathbb{R}^L_+$ records the

quantity of each good the consumer plans to consume. The non-negativity constraint $x_\ell \geq 0$ captures the physical impossibility of consuming a negative amount of a good.
Other consumption sets appear throughout the book.

Example 1.1. Two goods. With $L = 2$, the consumption set is the non-negative quadrant $X = \mathbb{R}^2_+$. A typical bundle is $x = (x_1, x_2)$ with $x_1, x_2 \geq 0$. This is the setting of the classical indifference-curve diagrams that Chapter 2 will draw.

Example 1.2. Discrete goods. With $L = 2$ but each x_ℓ required to be an integer, the consumption set becomes $X = \mathbb{Z}^2_+$, a lattice of points in the non-negative quadrant. Housing markets (houses come in units) and labor-supply choices (jobs are full-time or part-time) are two natural applications.

Example 1.3. A finite menu. A diner facing a fixed restaurant menu chooses from $X = \{\text{beef}, \text{chicken}, \text{fish}, \text{vegetarian}\}$, a finite set with $|X| = 4$. Section 1.6 develops the utility representation for finite X in detail.

Example 1.4. State-contingent consumption. Under uncertainty, an agent chooses among plans that specify what she will consume in each possible state of the world. With S states and L goods per state, the consumption set is $X = \mathbb{R}^{LS}_+$. The companion volume, *Choice Under Uncertainty*, develops this in detail.

Example 1.10. Preferences over lotteries as a compound-object consumption set. Consider three deterministic monetary prizes \$0, \$50, and \$100. A *lottery* is any probability distribution $p = (p_0, p_{50}, p_{100})$ with $p_0 + p_{50} + p_{100} = 1$ and each $p_i \geq 0$. The consumption set is the two-dimensional simplex

$$\Delta = \Big\{(p_0, p_{50}, p_{100}) \in \mathbb{R}^3_+ : p_0 + p_{50} + p_{100} = 1\Big\}.$$

Take the three specific lotteries $A = (0, 1, 0)$ (the sure \$50), $B = (\frac{1}{2}, 0, \frac{1}{2})$ (a 50-50 bet on \$0 or \$100), and $C = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ (a mixed lottery). Their expected values are $\mathbb{E}[A] = 50$, $\mathbb{E}[B] = 50$, and $\mathbb{E}[C] = 50$; expected value alone cannot rank them. An agent whose $\succsim$ ranks $A \succ C \succ B$ is risk-averse; one who ranks $B \succ C \succ A$ is risk-loving. In each case $\succsim$ is a preference relation on the compound-object set Δ, and the choice among A, B, and C is a choice among three points in that set.
Reading: even the simplest three-prize lottery lives in a two-dimensional set of probability vectors, so the tools of Chapters 2 and 3 apply directly on Δ. The companion volume develops the specialized structure that the linear expected-utility axioms impose on top of the general theory.

The consumption set is a modeling choice. Two features that recur in Chapters 2–6 deserve mention.

Convex. X is convex if for every pair $x, y \in X$ and every $\alpha \in [0,1]$, the mixture $\alpha x + (1-\alpha)y$ is also in X. The non-negative orthant $\mathbb{R}^L_+$ is convex; the integer lattice $\mathbb{Z}^2_+$ is not. Convexity of X is a background hypothesis that lets Chapter 2 speak of convex preferences and Chapter 6 speak of interior solutions to the utility-maximization problem.

Closed. X is closed as a subset of $\mathbb{R}^L$ if it contains all its limit points. $\mathbb{R}^L_+$ is closed; its interior $\mathbb{R}^L_{++}$ is not. Closedness matters because Chapter 3's utility representation theorem (Debreu 1954) requires the consumption set to be a connected topological space, and the argument passes through boundary points.

Unless otherwise stated, we take $X = \mathbb{R}^L_+$ for a fixed $L \geq 1$. This is the setting of the classical Walrasian consumer that Chapters 5 and 6 study. The finite-X case of Section 1.6 is an important exception, and the lexicographic example of Section 1.7 uses $L = 2$.

1.3 The preference relation

A **preference relation** on X is a binary relation $\succsim$ on X. Given any pair $x, y \in X$, the statement $x \succsim y$ is read "x is weakly preferred to y" or "x is at least as good as y." The relation $\succsim$ is the **primitive** of the theory: everything else is derived from it.
Two derived relations follow immediately.

The **strict-preference relation** $\succ$ is defined by

$$x \succ y \iff x \succsim y \text{ and not } y \succsim x.$$

The reading is: x is strictly preferred to y if x is weakly preferred to y and y is not weakly preferred to x.

The **indifference relation** $\sim$ is defined by

$$x \sim y \iff x \succsim y \text{ and } y \succsim x.$$

The reading is: the agent is indifferent between x and y if each is weakly preferred to the other.

Every pair $x, y \in X$ is either ranked ($x \succ y$ or $y \succ x$), tied ($x \sim y$), or unranked (neither $x \succsim y$ nor $y \succsim x$). The completeness axiom of Section 1.4 rules out the fourth possibility.

Example 1.5. Reading the notation. A consumer chooses between $x = (2,1)$ (two apples, one orange) and $y = (1,2)$ (one apple, two oranges). The statement $x \succsim y$ means the consumer is at least as happy with two apples and one orange as with one apple and two oranges. If additionally $y \not\succsim x$, then $x \succ y$. If instead both $x \succsim y$ and $y \succsim x$, then $x \sim y$: the consumer is genuinely indifferent between the two bundles.

Two small facts follow from the definitions and will be used without comment throughout the book.

Fact 1.1. $\succ$ is asymmetric: if $x \succ y$, then not $y \succ x$. *Proof.* If $x \succ y$, then by definition $x \succsim y$ and not $y \succsim x$. If also $y \succ x$, then $y \succsim x$ and not $x \succsim y$, contradicting $x \succsim y$. ∎

Fact 1.2. $\sim$ is reflexive whenever $\succsim$ is reflexive: $x \sim x$ for every $x \in X$. *Proof.* Reflexivity of $\succsim$ gives $x \succsim x$; applying it twice in the definition of $\sim$ gives $x \sim x$. ∎

1.4 Rationality: completeness and transitivity

A preference relation $\succsim$ is **rational** if it satisfies two axioms.

Axiom (Completeness). For every pair $x, y \in X$, either $x \succsim y$ or $y \succsim x$ (or both).

Axiom (Transitivity). For every triple $x, y, z \in X$, if $x \succsim y$ and $y \succsim z$, then $x \succsim z$.

Completeness says the agent can rank every pair of alternatives she faces: there are no "I don't know" answers. Transitivity says those pairwise rankings do not cycle: preferring a Toyota to a Honda and a Honda to a Nissan commits her to preferring a Toyota to a Nissan.

Two immediate consequences justify the label "rational."

Proposition 1.1. If $\succsim$ is complete and transitive, then $\succ$ is transitive: $x \succ y$ and $y \succ z$ imply $x \succ z$.

Proof. Suppose $x \succ y$ and $y \succ z$. By definition of $\succ$, we have $x \succsim y$, $y \succsim z$, not $y \succsim x$, and not $z \succsim y$. Transitivity of $\succsim$ applied to $x \succsim y$ and $y \succsim z$ gives $x \succsim z$. It remains to show $z \not\succsim x$. Suppose, toward a contradiction, that $z \succsim x$. Combined with $x \succsim y$ and transitivity, this gives $z \succsim y$, contradicting $y \succ z$. Hence $z \not\succsim x$, and by definition $x \succ z$. ∎

Proposition 1.2. If $\succsim$ is complete and transitive, then $\sim$ is an equivalence relation on X: it is reflexive, symmetric, and transitive.

Proof. Reflexivity: completeness applied to the pair (x, x) gives $x \succsim x$, and the definition of $\sim$ gives $x \sim x$. Symmetry: $x \sim y$ means $x \succsim y$ and $y \succsim x$, which is the same statement as $y \sim x$. Transitivity: suppose $x \sim y$ and $y \sim z$. Then $x \succsim y$, $y \succsim z$, $y \succsim x$, and $z \succsim y$. Transitivity of $\succsim$ gives $x \succsim z$ and $z \succsim x$, hence $x \sim z$. ∎

Completeness rules out incomparability. Transitivity rules out cycles. Together, they are the minimum content of the word "rational" in economic theory. Every substantive result in Chapters 2–6 uses one or both. Chapter 7 collects the laboratory evidence on which of them fails first, and in which environments.

Proposition 1.2a (mixed strict-weak chains). *If $\succsim$ is complete and transitive, then a chain of comparisons that mixes weak and strict preferences reduces to strict preference whenever it contains at least one strict link. Precisely: if $x_1 \succsim x_2 \succsim \cdots \succsim x_n$ with at least one strict step $x_i \succ x_{i+1}$, then $x_1 \succ x_n$.*

Proof. By transitivity of $\succsim$ applied repeatedly, $x_1 \succsim x_n$. It remains to show $x_n \not\succsim x_1$. Suppose, for contradiction, $x_n \succsim x_1$. Combined with $x_1 \succsim x_2 \succsim \cdots \succsim x_i$, transitivity gives $x_n \succsim x_i$. Combined with $x_i \succ x_{i+1}$, either use Proposition 1.1 (applied to the strict link) or directly: $x_n \succsim x_i$ and $x_{i+1} \succsim x_{i+2} \succsim \cdots \succsim x_n$ give $x_{i+1} \succsim x_i$ by transitivity, contradicting $x_i \succ x_{i+1}$ (which requires $\neg x_{i+1} \succsim x_i$). Hence $x_n \not\succsim x_1$, and $x_1 \succ x_n$. ∎

The proposition captures the useful idea that strict preference *propagates*: a single strict link in a chain of weak comparisons is enough to make the endpoints stand in strict preference. Chapters 4 and 5 use this repeatedly when they reason about revealed preferences generated by chains of directly observed comparisons.

Example 1.6. Cyclic preferences are not rational. A hiring committee compares three candidates A, B, and C pairwise on three dimensions – research, teaching, and service. Candidate A beats B on two dimensions (research and teaching); B beats C on two dimensions (teaching and service); C beats A on two dimensions (research and service). Applying majority-rule pairwise voting produces the strict rankings $A \succ B$, $B \succ C$, and $C \succ A$, which violate transitivity. The example illustrates Condorcet's paradox: pairwise majority voting does not, in general, generate a transitive social preference relation even when each individual's preference is transitive. Arrow (1951) develops the general impossibility theorem behind the paradox.

Example 1.7. Semi-orders and just-noticeable differences. A consumer's ability to detect small preference differences is not perfect. She may prefer a cup of coffee with 100 grains of sugar to one with 101 grains only imperceptibly, but she strictly prefers 100 grains to 200. If the "imperceptible" pairs are declared to be indifferent, transitivity of indifference forces $100 \sim 101 \sim 102 \sim \cdots \sim 200$, hence $100 \sim 200$ – a contradiction. *Semi-orders* (Luce 1956) formalize this failure by weakening transitivity of $\sim$ while keeping transitivity of $\succ$. Section 7.3 collects the laboratory evidence on preference-detection thresholds.

Example 1.11. A monetary semi-order with a five-dollar threshold. Take $X = \mathbb{R}_+$, interpreted as dollar amounts, and define

$$x \succsim y \iff x \geq y - 5.$$

Two amounts are called indifferent when neither strictly dominates the other, i.e., $x \sim y$ if and only if $|x - y| \leq 5$. Consider the chain $x_1 = 100$, $x_2 = 104$, $x_3 = 108$, $x_4 = 112$, $x_5 = 116$. Adjacent pairs are within 5 of each other, so

$$x_1 \sim x_2 \sim x_3 \sim x_4 \sim x_5,$$

yet $x_5 - x_1 = 16 > 5$, so $x_5 \succ x_1$. Transitivity of $\sim$ fails on the chain: the first and last amounts are strictly ordered even though every adjacent pair is called indifferent. Strict preference is still transitive: if $x - y > 5$ and $y - z > 5$, then $x - z > 10 > 5$, so $x \succ y \succ z$ implies $x \succ z$. The relation is a semi-order rather than a rational $\succsim$ in the sense of Definition 1.2.

Reading: a threshold of just-noticeable difference breaks the equivalence-class structure that Section 1.5 relies on. The dollar amount \$100 and the amount \$116 are put in the same "indifference cloud" by every adjacent comparison, and yet the direct comparison declares them strictly ordered. Section 7.3 shows this pattern in laboratory data.

1.5 Indifference classes

Given a rational preference relation $\succsim$ on X and a point $x \in X$, the **indifference class of** x is the set

$$I(x) = \{y \in X : y \sim x\}.$$

It contains every alternative the agent regards as tied with x.

Proposition 1.3. The indifference classes of a rational $\succsim$ partition X: every $x \in X$ belongs to exactly one indifference class.

Proof. Every x is in its own indifference class $I(x)$ because reflexivity gives $x \sim x$. Suppose $x \in I(y) \cap I(z)$. Then $x \sim y$ and $x \sim z$. Symmetry of $\sim$ gives $y \sim x$ and $z \sim x$; transitivity of $\sim$ gives $y \sim z$, so $I(y) = I(z)$. ■

When $X = \mathbb{R}^L_+$ and $\succsim$ is a smooth, strictly-monotone preference relation of the kind studied in Chapter 2, each indifference class is a smooth curve (for $L = 2$) or a smooth $(L-1)$-dimensional surface (for $L \geq 3$). Chapter 2 draws these curves. Figure 1.1 illustrates the partition for the Leontief case $\min\{x_1, x_2\} = k$, whose L-shaped indifference classes tile the positive orthant.

Intuitively, the partition property packages a substantial amount of the theory into one geometric statement: rationality of $\succsim$ organizes the whole consumption set X into non-overlapping layers, each labeled by an indifference class. Two operations on the layers are of most interest for Chapters 2 and 6. First, the layers can be ranked: one lies "above" another if any of its members strictly beats any of the other's. Second, the layers have a shape (smooth curve, straight line, kink); the axioms of Chapter 2 constrain the possible shapes. When we speak of "the indifference curve through x" in later chapters we always mean the indifference class $I(x)$ viewed as a geometric object in X.

Proposition 1.3a (indifference classes are chained by strict preference). *Let $\succsim$ be rational on X. For any two indifference classes $I(x)$ and $I(y)$, exactly one of the following holds: (i) $I(x) = I(y)$; (ii) every member of $I(x)$ is strictly preferred to every member of $I(y)$; or (iii) every member of $I(y)$ is strictly preferred to every member of $I(x)$.*

Proof. If $I(x) = I(y)$, we are in case (i). Otherwise, $x \not\sim y$ (else transitivity of $\sim$ would place them in the same class). By completeness, either $x \succ y$ or $y \succ x$; suppose $x \succ y$. Take any $x' \in I(x)$ and $y' \in I(y)$; then $x' \sim x \succ y \sim y'$. By Proposition 1.1 (or by Proposition 1.2a on mixed chains), $x' \succ y'$. Since x', y' were arbitrary, every member of $I(x)$ strictly beats every member of $I(y)$: case (ii). Symmetrically for case (iii). ■

Reading: indifference classes are cleanly ranked. There is no "some members of $I(x)$ beat some of $I(y)$ and vice versa" phenomenon. This is what allows Chapter 3 to represent the whole preference relation by a single-number utility.

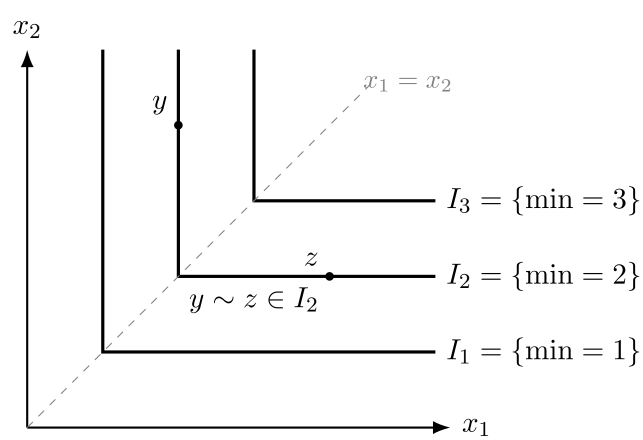

Figure 1.1: Indifference classes partition $\mathbb{R}^2_+$. Under $\succsim$ generated by $u(x) = \min\{x_1, x_2\}$, each L-shaped locus is one indifference class; $y = (2, 4)$ and $z = (4, 2)$ share the same minimum and therefore lie in I_2. Every point in $\mathbb{R}^2_+$ belongs to exactly one such class, matching Proposition 1.3.

1.6 The finite case: utility from completeness and transitivity alone

When the consumption set X is finite, completeness and transitivity alone are enough to guarantee that $\succsim$ can be represented by a real-valued utility function. This is the simplest possible representation theorem.

Definition 1.1. A function $u : X \to \mathbb{R}$ **represents** the preference relation $\succsim$ if, for every $x, y \in X$,

$$x \succsim y \iff u(x) \geq u(y).$$

Proposition 1.4. If X is finite and $\succsim$ is complete and transitive, then there exists a utility function $u : X \to \mathbb{R}$ that represents $\succsim$.

Proof. Let $n = |X| < \infty$. For every $x \in X$, define

$$L(x) \equiv \{y \in X : x \succsim y\} \quad \text{(the weak lower-contour set of } x\text{)}, \qquad u(x) \equiv |L(x)|.$$

Because $x \succsim x$ by completeness applied to (x, x) (either $x \succsim x$ or $x \succsim x$), $x \in L(x)$, so $u(x) \geq 1$; because $L(x) \subseteq X$, $u(x) \leq n$. Hence $u : X \to \{1, \ldots, n\} \subset \mathbb{R}$. It remains to show, for every $x, y \in X$, that

$$x \succsim y \iff u(x) \geq u(y).$$

($\Rightarrow$) Suppose $x \succsim y$. Take any $z \in L(y)$, i.e., $y \succsim z$. Then by transitivity applied to $x \succsim y \succsim z$, $x \succsim z$, i.e., $z \in L(x)$. Hence $L(y) \subseteq L(x)$, and therefore $u(y) = |L(y)| \leq |L(x)| = u(x)$.
($\Leftarrow$) We argue the contrapositive: if $x \not\succsim y$, then $u(x) < u(y)$. By completeness (applied to the pair (x, y)), $x \not\succsim y$ implies $y \succ x$, i.e., $y \succsim x$ and not $x \succsim y$. Applied to $y \succ x$ and any $z \in L(x)$, transitivity gives $y \succsim z$, so $z \in L(y)$; hence $L(x) \subseteq L(y)$. Moreover $y \in L(y)$ (reflexivity), while $y \notin L(x)$ (from $x \not\succsim y$). Therefore $L(x) \subsetneq L(y)$, i.e., $u(x) = |L(x)| < |L(y)| = u(y)$, contradicting $u(x) \geq u(y)$. ■

Intuitively, the function u turns each alternative's "rank" into a number by counting how many alternatives it weakly beats. Transitivity is what guarantees that this counting is monotone across the ranking: whenever x is at least as high as y, everything below y is automatically below x, so x's count is at least y's. Completeness ensures that any two alternatives are actually comparable, so no pair falls through the cracks.

Example 1.8. Restaurant menu. A diner ranks four entrees: fish $\succ$ chicken $\succ$ vegetarian $\succ$ beef (a strict ranking with no ties). The utility function of Proposition 1.4 assigns $u(\text{fish}) = 4$, $u(\text{chicken}) = 3$, $u(\text{vegetarian}) = 2$, $u(\text{beef}) = 1$. Any strictly increasing transformation of these numbers – for instance, 100, 75, 50, 25 – represents the same preference relation. The utility function is unique only up to strictly increasing transformations. Chapter 3 develops this uniqueness result in generality. Figure 1.2 draws the diner's preference relation as a directed graph and reads the utility values off the number of arrows out of each node.

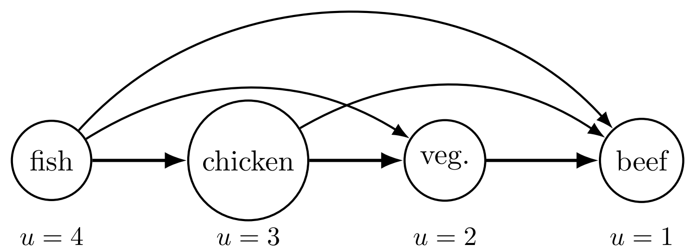

Figure 1.2: The preference relation fish $\succ$ chicken $\succ$ veg. $\succ$ beef drawn as a directed graph. An arrow from x to y means $x \succ y$; the adjacent-pair arrows plus their transitive closures give the complete strict order. The utility of each node in Proposition 1.4 is one plus the number of nodes it points to.

Example 1.9. Five alternatives with a strict ranking and one tie. Let $X = \{a, b, c, d, e\}$ and consider a preference relation $\succsim$ with the reported comparisons

$$a \succ b, \quad b \sim c, \quad c \succ d, \quad d \succ e.$$

The other pairwise comparisons are determined by transitivity. We work out the utility $u(x) = \#\{y \in X : x \succsim y\}$ of Proposition 1.4 step by step.

Step 1. Complete the ranking by transitivity. From $a \succ b$ and $b \sim c$ we get $a \succ c$ (strict preference transported across an indifference). From $b \sim c$ and $c \succ d$ we get $b \succ d$; combined with $d \succ e$, both $b \succ e$ and $c \succ e$. Chaining across, $a \succ d$ and $a \succ e$. The full strict order on the four "levels" is

$$a \ \succ \ \{b, c\} \ \succ \ d \ \succ \ e,$$

with b and c tied on the second level.

Step 2. Count weak lower-contour sets. For each $x \in X$, compute $L(x) \equiv \{y \in X : x \succsim y\}$ and set $u(x) = |L(x)|$.

x	$L(x) = \{y : x \succsim y\}$	$u(x) = \lvert L(x)\rvert$
a	$\{a, b, c, d, e\}$	5
b	$\{b, c, d, e\}$	4
c	$\{b, c, d, e\}$	4
d	$\{d, e\}$	2
e	$\{e\}$	1

Step 3. Verify the representation on every pair. Table entries encode the fifteen pairwise comparisons (five diagonal ones are trivial reflexive $x \succsim x$; ten off-diagonal ones must match u). For any two of the ten off-diagonal pairs, $x \succsim y$ if and only if $u(x) \geq u(y)$. For instance, $b \succsim c$ and $c \succsim b$ (indifference) match $u(b) = u(c) = 4$; $a \succ d$ matches $u(a) = 5 > 2 = u(d)$; and $c \succ e$ matches $u(c) = 4 > 1 = u(e)$. Every comparison checks out.
Step 4. Uniqueness up to strictly increasing transformations. Replacing u by any strictly increasing $f : \mathbb{R} \to \mathbb{R}$ preserves the same ordering. Two natural rescalings that represent the same $\succsim$:

$$u'(a) = 100,\ u'(b) = u'(c) = 80,\ u'(d) = 40,\ u'(e) = 20 \quad (\text{from } u'(x) = 20 \cdot u(x)),$$

$$u''(a) = \ln 5,\ u''(b) = u''(c) = \ln 4,\ u''(d) = \ln 2,\ u''(e) = \ln 1 = 0 \quad (\text{from } u''(x) = \ln u(x)).$$

Both encode the same ranking $a \succ b \sim c \succ d \succ e$. Chapter 3 formalizes this ordinal-uniqueness property in general.

Lesson: on a finite consumption set, completeness plus transitivity are sufficient for a utility representation, and the construction is completely mechanical – count how many alternatives each x weakly dominates.

The finite case is the easy case. When X is infinite – $\mathbb{R}^L_+$, for example – completeness and transitivity alone are not enough: a further axiom (continuity, in Chapter 2) is needed. The lexicographic example of Section 1.7 shows why: it is complete and transitive, but any purported utility function would need to assign, in an uncountable set of pairwise disjoint open intervals of the real line, a separate rational number – an impossibility.

1.7 The lexicographic preference relation

The **lexicographic preference relation** $\succsim^L$ on $X = \mathbb{R}^2_+$ is defined as follows. For any two bundles $x = (x_1, x_2)$ and $y = (y_1, y_2)$,

$$x \succsim^L y \quad \iff \quad \text{either} \quad x_1 > y_1, \quad \text{or} \quad x_1 = y_1 \text{ and } x_2 \geq y_2.$$

Read: rank first by good 1; break ties with good 2. The bundle with more of good 1 wins outright, regardless of good 2; two bundles with the same x_1 are compared on x_2.

Proposition 1.5. The lexicographic preference relation $\succsim^L$ is complete and transitive.

Proof. Completeness. Given any two bundles x and y, exactly one of the three statements holds: $x_1 > y_1$, $y_1 > x_1$, or $x_1 = y_1$. In the first case, $x \succ^L y$; in the second, $y \succ^L x$; in the

third, either $x_2 \geq y_2$ (giving $x \succsim^L y$) or $y_2 \geq x_2$ (giving $y \succsim^L x$), or both (indifference, which under the lexicographic definition requires $x_1 = y_1$ and $x_2 = y_2$, hence $x = y$).
Transitivity. Suppose $x \succsim^L y$ and $y \succsim^L z$. Case-split on good 1.
Case (a). $x_1 > y_1$. Combined with $y_1 \geq z_1$ (from $y \succsim^L z$), $x_1 > z_1$, so $x \succ^L z$.
Case (b). $x_1 = y_1$ and $y_1 > z_1$. Then $x_1 > z_1$, so $x \succ^L z$.
Case (c). $x_1 = y_1 = z_1$. Then $x_2 \geq y_2$ and $y_2 \geq z_2$ (from the two lexicographic premises), so $x_2 \geq z_2$; combined with $x_1 = z_1$, $x \succsim^L z$.
The three cases exhaust the possibilities. ■

Two features of $\succsim^L$ are worth noticing. First, $\sim^L$ is trivial: $x \sim^L y$ only when $x = y$. There are no "interesting" indifference classes. Second, and more importantly, $\succsim^L$ admits no utility representation by a real-valued function – a fact we state now and prove in Chapter 3. Figure 1.3 plots the strict upper-contour set of a reference bundle $x^0 = (x_1^0, x_2^0)$ under $\succsim^L$ and the resulting priority structure across the two goods.

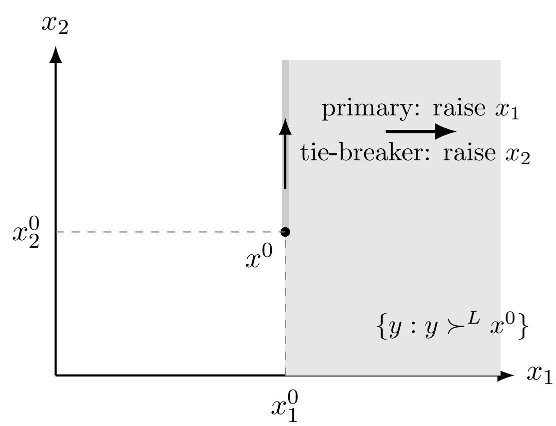

Figure 1.3: Lexicographic preferences $\succsim^L$ on $\mathbb{R}^2_+$. The strict upper-contour set of x^0 is the open half-plane $\{y : y_1 > x_1^0\}$ together with the vertical half-line $\{y : y_1 = x_1^0, y_2 > x_2^0\}$. Good 1 has absolute priority; good 2 breaks ties on good 1.

Proposition 1.6 (preview of Chapter 3). There is no function $u : \mathbb{R}^2_+ \to \mathbb{R}$ that represents $\succsim^L$.

The intuition is that $\succsim^L$ has *too many* indifference classes – one per point in $\mathbb{R}^2_+$ – and no function on $\mathbb{R}$ can distinguish uncountably many of them. Chapter 3 gives the formal argument (a cardinality proof).

The lexicographic example is the canonical demonstration that completeness and transitivity, though necessary for utility representation, are not sufficient on an infinite consumption set. Chapter 2 introduces the further axiom – continuity – that closes the gap.

1.8 Experimental evidence

Completeness and transitivity look almost tautological. Both fail in the laboratory in predictable ways, and the systematic pattern of failure is one of the empirical starting points of behavioral economics.

Completeness failures. Laboratory subjects presented with sufficiently unfamiliar pairs – say, gambles specified in different frames, or bundles that require experience the subject lacks – routinely refuse to rank the alternatives. The refusal is not noise: repeated presentations of the same pair produce consistent "I don't know" responses. Bewley (1986, 2002) develops the theoretical apparatus of *incomplete* preferences that this evidence motivates; Chapter 8 of the companion volume takes up the ambiguity models that use it.

Intransitivity: Tversky's cycles. Tversky (1969) constructed a menu of five gambles A–E, each pair separated by small enough differences in probability and payoff that subjects' pairwise choices followed a predictable heuristic (compare payoffs when probabilities are similar, compare probabilities when payoffs are similar). Applied across the full menu, the heuristic produced strict pairwise choices $A \succ B \succ C \succ D \succ E \succ A$: a full cycle of length five. About one-third of Tversky's subjects exhibited the cycle.

Intransitivity: preference reversals. Lichtenstein and Slovic (1971) presented subjects with pairs of gambles, one with a high probability of a modest gain (the "P-bet") and one with a low probability of a large gain (the "$-bet"). In choice tasks, subjects picked the P-bet more than half the time; in pricing tasks (asked to state a selling price for each), the same subjects priced the $-bet higher more than half the time. The combined revealed ranking is intransitive: P-bet $\succ$ $-bet in choice, $-bet $\succ$ P-bet in pricing. The reversal replicates across many subject pools and paradigms and is one of the most robust findings in experimental decision theory.

Intransitivity: Rubinstein's similarity. Rubinstein (1988) developed a formal model of intransitive similarity judgments and showed that if subjects' pairwise choice procedure uses a different threshold on each attribute, cycles must arise on menus that span multiple attributes. Laboratory evidence (e.g., Rubinstein and Zhou 1999) confirmed the theoretical predictions.

Interpretation. Completeness and transitivity are the foundation of the rest of the book, and this book takes them as axioms in Chapters 2–6. The experimental evidence is a reminder that they are behavioral assumptions, not logical necessities. When completeness fails, the ambiguity models of Chapter 8 of the companion volume become the natural framework;

when transitivity fails, the non-EU frameworks of Chapter 7 of the companion volume take over. In this book, we build the apparatus that both extensions later refine.

1.9 Summary

- The **consumption set** X is the set of alternatives the decision-maker can imagine choosing from; the leading example is $X = \mathbb{R}^L_+$.
- A **preference relation** $\succsim$ on X is a binary relation. Its derived relations are strict preference $\succ$ and indifference $\sim$.
- A preference relation is **rational** if it is *complete* (every pair is comparable) and *transitive* (rankings do not cycle). Under rationality, $\succ$ is transitive and $\sim$ is an equivalence relation, so indifference classes partition X.
- When X is finite, completeness and transitivity alone guarantee a utility representation $u(x) = \#\{y : x \succsim y\}$.
- When X is infinite, completeness and transitivity are not sufficient: the **lexicographic** preference relation $\succsim^L$ is complete and transitive but admits no real-valued utility representation. Chapter 2 introduces continuity to close the gap.
- Laboratory subjects violate completeness (they refuse to rank sufficiently unfamiliar pairs) and transitivity (Tversky's cycles, Lichtenstein-Slovic reversals) in predictable ways. Both violations motivate the extensions that later books in the sequence develop.

1.10 Guided exercise

Exercise (Guided). On $X = \{a, b, c, d\}$, a decision-maker's pairwise weak-preference relation $\succsim$ is given by the table below, where the entry in row x and column y is 1 if $x \succsim y$ and 0 otherwise.

	a	b	c	d
a	1	1	1	1
b	0	1	1	0
c	0	0	1	1
d	1	1	0	1

Verify whether $\succsim$ is complete and transitive. If it is rational, construct the utility function of Proposition 1.4 and list the indifference classes. If it is not rational, identify a specific failure.

Solution.

Completeness. A relation is complete if for every ordered pair (x, y), at least one of $x \succsim y$ or $y \succsim x$ holds. Read the table: the pair (a, d) has $a \succsim d$ (row a, column d: 1) and $d \succsim a$ (row d, column a: 1), so both hold. The pair (b, d): $b \succsim d$ is 0, but $d \succsim b$ is 1. The pair (c, d): $c \succsim d$ is 1, $d \succsim c$ is 0. Every pair has at least one entry equal to 1, so $\succsim$ is complete.

Transitivity. A relation is transitive if $x \succsim y$ and $y \succsim z$ imply $x \succsim z$. Test the triple (a, d, c): $a \succsim d$ (1) and $d \succsim a$ (1) so far consistent; but the relevant test is $d \succsim a$ (1) and $a \succsim c$ (1) $\Rightarrow$ $d \succsim c$. Read the table: $d \succsim c$ is 0. So transitivity fails: $d \succsim a$ and $a \succsim c$ but not $d \succsim c$.

Diagnosis. The failure traces back to the pair (a, d), which the table records as indifferent ($a \sim d$: both $a \succsim d$ and $d \succsim a$ equal 1). If $a \sim d$ and $a \succ c$ (from row a, column c: 1, and row c, column a: 0), transitivity of the derived $\sim$ would force $d \succ c$, contradicting the table entry $d \succsim c = 0$. The relation $\succsim$ is complete but not transitive; it is not rational; Proposition 1.4 does not apply. The failing triple (d, a, c) is the smallest witness.

Lesson. Completeness is a two-way property of pairs; transitivity is a three-way property of triples. Checking completeness requires $|X|\dfrac{|X|+1}{2}$ pair reads; checking transitivity requires $|X|^3$ triple reads. For $|X| = 4$ this is 10 pairs and 64 triples; for large X, the transitivity check dominates. This is why, in practice, one usually specifies $\succsim$ by giving a utility function u that represents it: Proposition 1.4 makes the check automatic.

Chapter 2 - Structural axioms on preferences

2.1 Introduction

Chapter 1 gave the primitive: a binary relation $\succsim$ on the consumption set X. It layered on the two rationality axioms – completeness and transitivity – and showed that when X is finite, those two axioms alone suffice for a utility representation. The lexicographic example of Section 1.7 showed that when X is infinite, they do not. This chapter closes the gap by introducing three families of **structural axioms** that shape a rational preference relation into the geometric object the rest of the book relies on.

The three families are (i) *continuity*, which rules out jumps in the preference relation as bundles vary; (ii) two monotonicity conventions – *monotonicity*, *strong monotonicity*, and *local non-satiation* – which fix the economic sign of "more is better"; and (iii) *convexity* and *strict convexity*, which give indifference sets the concave-toward-the-origin shape that underlies interior demand.

Every structural axiom in this chapter is behavioral, not logical. Each holds for large classes of preferences (Cobb-Douglas, CES, quasi-linear) and fails for others (lexicographic, some addictive-good specifications). Section 2.8 collects the laboratory evidence on how often each axiom fails.

Section 2.2 defines continuity. Section 2.3 studies the three monotonicity notions and their nesting. Section 2.4 develops the weakest useful "more is better" assumption, local non-satiation, and shows why the utility-maximization problem of Chapter 6 needs it. Section 2.5 defines convex and strictly convex preferences and connects them to the quasi-concavity of any utility representation. Section 2.6 introduces the marginal rate of substitution – the local slope of an indifference curve – as the geometric object the remaining chapters use most often. Section 2.7 catalogs the four canonical parametric families (Cobb-Douglas, CES, quasi-linear, and Leontief) so that the reader has ready-to-use examples for Chapters 5 and 6. Section 2.8 collects the experimental evidence. Section 2.9 summarizes.

Throughout, we take $X = \mathbb{R}^L_+$ with $L \geq 1$ and $\succsim$ rational (complete and transitive) unless otherwise stated.

2.2 Continuity

Let $\succsim$ be a rational preference relation on $X = \mathbb{R}^L_+$. For each $x \in X$, define two sets:

- the **upper-contour set** of x, $U(x) = \{y \in X : y \succsim x\}$: bundles at least as good as x;
- the **lower-contour set** of x, $L(x) = \{y \in X : x \succsim y\}$: bundles at most as good as x.

The indifference set $I(x) = U(x) \cap L(x)$ is the set of bundles tied with x; the strict upper-contour set is $U(x) \setminus I(x)$; the strict lower-contour set is $L(x) \setminus I(x)$.

Axiom (Continuity). $\succsim$ is **continuous** if, for every $x \in X$, both $U(x)$ and $L(x)$ are closed subsets of X in the standard topology on $\mathbb{R}^L$.

Intuitively, continuity rules out preferences that "jump" as bundles vary. If a sequence of bundles y_n converges to y and each y_n is at least as good as x (so $y_n \in U(x)$ for every n), then the limit y is also at least as good as x (so $y \in U(x)$). Preferences do not flip discontinuously at the limit. Figure 2.1 displays the upper- and lower-contour sets of a bundle x^0 for a strictly convex, smooth preference relation.

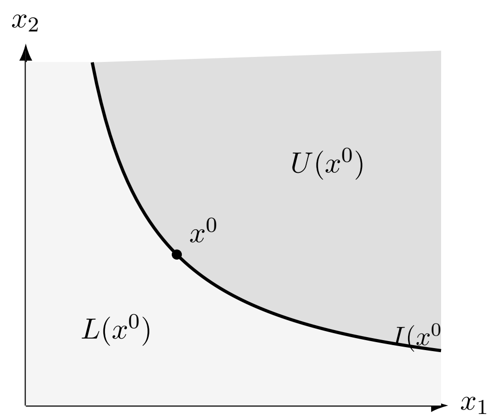

Figure 2.1: Upper- and lower-contour sets of x^0 under a strictly convex, smooth preference relation. The dark region is $U(x^0) = \{y : y \succsim x^0\}$, the light region is $L(x^0) = \{y : x^0 \succsim y\}$, and their common boundary is the indifference class $I(x^0)$. Continuity requires both $U(x^0)$ and $L(x^0)$ to be closed.

Equivalent forms of continuity. Two formulations of continuity appear in the literature. Each is equivalent to the closed-contour-set definition given above.

Proposition 2.1. *Let $\succsim$ be a rational preference relation on $X = \mathbb{R}^L_+$. The following statements are equivalent.*

(i) *(Closed-contour form.) For every $x \in X$, both $U(x)$ and $L(x)$ are closed in the standard topology on $\mathbb{R}^L$.*

(ii) *(Sequential-with-fixed-anchor form.) For every $x \in X$ and every convergent sequence $(y_n) \to y$ in X: if $y_n \succsim x$ for every n, then $y \succsim x$; and if $x \succsim y_n$ for every n, then $x \succsim y$.*

(iii) *(Complementary-strict form.) For every* $x \in X$*, both* $\{y \in X : y \succ x\}$ *and* $\{y \in X : x \succ y\}$ *are open in* X*.*

Proof of Proposition 2.1.
(i) $\Leftrightarrow$ *(ii).* Fix $x \in X$. By definition, $U(x) = \{y : y \succsim x\}$ is closed if and only if every convergent sequence (y_n) with each $y_n \in U(x)$ has its limit in $U(x)$, which is exactly the first clause of (ii) written out. Symmetrically for $L(x)$.
(i) $\Leftrightarrow$ *(iii).* By completeness of $\succsim$, for every $x, y \in X$ exactly one of $y \succsim x$ and $x \succ y$ holds; hence $\{y : y \succ x\} = X \setminus L(x)$. Since $L(x)$ is a subset of the metric space X, $L(x)$ closed $\Leftrightarrow X \setminus L(x)$ open. Symmetrically, $\{y : x \succ y\} = X \setminus U(x)$, and $U(x)$ closed $\Leftrightarrow \{y : x \succ y\}$ open. ■

Remark. Form (ii) states the sequential content of continuity: preferences do not flip in the limit against a fixed reference bundle. In the applications of Chapters 3–6 – upper hemi-continuity of the demand correspondence, existence of a maximizer in the Weierstrass theorem, and continuity of the value function via Berge – this fixed-anchor version is what is invoked most often. The stronger, "both sequences moving" version – for $x_n \to x$, $y_n \to y$ with $x_n \succsim y_n$, one has $x \succsim y$ – is also standard, but requires a somewhat more delicate argument (see Debreu (1954)) and is not needed in the body of these notes.

Intuitively, the closed-contour form says that neither "y is at least as good as x" nor "y is at most as good as x" can flip in the limit as y moves smoothly through the consumption set. The sequential form says the same thing with a pair of moving bundles rather than one moving and one fixed. The closed-graph form packages both directions in one statement: the entire relation, viewed as a subset of the product space, is a topologically closed object. All three views amount to the same requirement: preferences do not jump.

Example 2.1. Cobb-Douglas is continuous. On $X = \mathbb{R}^2_+$, define $\succsim$ by $x \succsim y \Leftrightarrow x_1^{1/2} x_2^{1/2} \geq y_1^{1/2} y_2^{1/2}$. The upper-contour set of any x is $U(x) = \{y : y_1^{1/2} y_2^{1/2} \geq k\}$ for $k = x_1^{1/2} x_2^{1/2}$; this is the closed super-level set of a continuous function, so $U(x)$ is closed. Similarly for $L(x)$. Hence $\succsim$ is continuous.

Example 2.2. Lexicographic preferences are not continuous. On $X = \mathbb{R}^2_+$, take the lexicographic $\succsim^L$ of Section 1.7. Consider the sequence $y_n = (1 - \frac{1}{n}, 2)$. Every y_n has $y_{n,1} < 1$, so $y_n \prec^L (1, 1)$; hence $y_n \in L((1, 1))$ for every n. The limit $y = (1, 2)$ has $y_1 = 1$ and $y_2 = 2 > 1$, so $y \succ^L (1, 1)$; hence $y \notin L((1, 1))$. The lower-contour set of $(1, 1)$ is not closed, so $\succsim^L$ is not continuous. This failure of continuity is exactly why $\succsim^L$ admits no utility representation (Proposition 1.6, proved in Chapter 3).

Debreu (1954). Continuity, together with completeness and transitivity, is what Chapter 3 shows is sufficient for a continuous utility function representing $\succsim$ to exist. The axiom is the single most important structural hypothesis in the theory: without it, utility does not exist even in the simplest infinite-set cases.

2.3 Monotonicity

Three monotonicity conventions appear in the literature. The weakest is often called simply "monotonicity," the strongest is "strong monotonicity," and the intermediate one is "strict monotonicity" (terminology varies across textbooks). The three notions differ in how they treat the boundary case of bundles that differ in some but not all coordinates.

For bundles $x, y \in \mathbb{R}^L_+$, write $y \geq x$ if $y_\ell \geq x_\ell$ for every ℓ; $y > x$ if $y \geq x$ and $y \neq x$ (so y dominates x weakly and beats it in at least one coordinate); and $y \gg x$ if $y_\ell > x_\ell$ for every ℓ.

Axiom (Monotonicity). $\succsim$ is **monotone** if $y \geq x$ implies $y \succsim x$.

Axiom (Strong monotonicity). $\succsim$ is **strongly monotone** if $y > x$ implies $y \succ x$.

Axiom (Strict monotonicity). $\succsim$ is **strictly monotone** if $y \gg x$ implies $y \succ x$.

Intuitively, monotonicity says "more of every good is at least as good." Strong monotonicity says "more of any single good, with the others held fixed, is strictly better." Strict monotonicity says "strictly more of every good is strictly better." The three conventions satisfy

$$\text{strong monotonicity} \Rightarrow \text{monotonicity} + \text{strict monotonicity},$$

but neither monotonicity nor strict monotonicity implies strong monotonicity in general. The Leontief preferences of Section 2.7 are strictly monotone but not strongly monotone.

Example 2.3. Cobb-Douglas is strongly monotone. On $X = \mathbb{R}^2_{++}$, if $y > x$ then either $y_1 > x_1$ (and $y_2 \geq x_2$) or $y_2 > x_2$ (and $y_1 \geq x_1$). Either way $y_1^{1/2} y_2^{1/2} > x_1^{1/2} x_2^{1/2}$, so $y \succ x$.

Example 2.4. Leontief is strictly monotone but not strongly monotone. On $X = \mathbb{R}^2_+$, define $\succsim$ by $x \succsim y \Leftrightarrow \min\{x_1, x_2\} \geq \min\{y_1, y_2\}$. Take $x = (2, 2)$ and $y = (3, 2)$: here $y > x$ (more of good 1, same good 2), but $\min\{y\} = \min\{x\} = 2$, so $y \sim x$. Strong monotonicity fails at x. However, if $y \gg x$, then both $y_\ell > x_\ell$, so $\min\{y\} > \min\{x\}$ and $y \succ x$; strict monotonicity holds.

2.4 Local non-satiation

Chapters 6 and 7 need a weaker "more is better" hypothesis than monotonicity, one that does not require a global sign on the gradient of u but only rules out interior bliss points.

Axiom (Local non-satiation). $\succsim$ is **locally non-satiated** if for every $x \in X$ and every $\varepsilon > 0$, there exists a bundle $y \in X$ with $\|y - x\| < \varepsilon$ and $y \succ x$.

Intuitively, no matter which bundle x the agent starts at and no matter how tight the neighborhood, she can find a strictly better bundle inside that neighborhood. There are no locally satiated points.

Nesting. Strong monotonicity implies strict monotonicity, which implies local non-satiation: given any x, pick any direction $d \gg 0$ and any $\varepsilon > 0$ small enough that $x + \lambda d \in X$ for $\lambda \in (0, \varepsilon/\|d\|)$; then $x + \lambda d \gg x$, so $x + \lambda d \succ x$ by strict monotonicity. Local non-satiation does not, in general, imply either strict or strong monotonicity: preferences with non-monotonic "preferred directions" at some points can still be locally non-satiated.

Example 2.10. Locally non-satiated but not monotone. Take $X = \mathbb{R}^2_+$ and consider preferences with a rotating gradient near the boundary point $\bar{x} = (1, 1)$, represented by

$$u(x_1, x_2) = (x_1 - 1)\cos\theta + (x_2 - 1)\sin\theta, \qquad \theta = \arctan(x_2 - 1, x_1 - 1),$$

i.e., $u(x) = \sqrt{(x_1 - 1)^2 + (x_2 - 1)^2}$ (the Euclidean distance from $\bar{x}$). Every neighborhood of $\bar{x}$ contains a point $y \neq \bar{x}$ with $u(y) > u(\bar{x}) = 0$, so $\succsim$ is locally non-satiated at $\bar{x}$. Yet $\succsim$ is not monotone: the bundle $(1.1, 0.9)$ has $u = \sqrt{0.01 + 0.01} \approx 0.141$, and so does $(0.9, 1.1)$, though the first weakly dominates $\bar{x}$ in coordinate 1 only, and the second in coordinate 2 only. Worse, adding a small amount of good 2 to a point like $(2, 1)$ where $u = 1$ can move the consumer toward or away from $\bar{x}$, so $\partial u/\partial x_2$ changes sign in the plane; strict monotonicity ($x' \gg x \Rightarrow x' \succ x$) fails.

Reading: local non-satiation is the weakest "no bliss point" hypothesis. It rules out interior maxima but allows the gradient of u to point in any direction, so demand under LNS alone can be much wilder than demand under monotone preferences. Chapter 6 uses LNS to prove Walras's law and nothing more.

Why the utility-maximization problem needs local non-satiation. Chapter 6 shows that if x^* solves the utility-maximization problem $\max_x u(x)$ subject to $p \cdot x \leq w$ and $\succsim$ is locally non-satiated, then $p \cdot x^* = w$: the optimal bundle exhausts the budget. The reason is direct: if $p \cdot x^* < w$, there is a neighborhood of x^* still contained in the budget set, and local non-satiation produces a bundle in that neighborhood strictly preferred to x^*, contradicting optimality. Without local non-satiation, the equality $p \cdot x^* = w$ can fail.

Proposition 2.2 (Walras's law from local non-satiation). *Let $\succsim$ be a locally non-satiated, rational, continuous preference relation on $\mathbb{R}^L_+$, and let $x^*(p, w)$ solve the utility-maximization problem at (p, w) with $p \gg 0$ and $w > 0$. Then $p \cdot x^*(p, w) = w$.*

Proof. Suppose, for contradiction, that $p \cdot x^* < w$; set $\delta = w - p \cdot x^* > 0$. Choose $\varepsilon > 0$ small enough that the open ball $B(x^*, \varepsilon)$ is contained in $\{x \in \mathbb{R}^L_+ : p \cdot x \leq w\}$. Concretely, take $\varepsilon = \delta/(2\|p\|)$: for every $y \in B(x^*, \varepsilon)$, the Cauchy-Schwarz inequality gives

$$p \cdot y = p \cdot x^* + p \cdot (y - x^*) \leq p \cdot x^* + \|p\| \, \|y - x^*\| < p \cdot x^* + \|p\| \cdot \frac{\delta}{2\|p\|} = p \cdot x^* + \frac{\delta}{2} < w.$$

So y is feasible. By local non-satiation applied at x^* with the ball $B(x^*, \varepsilon)$, there exists $y \in B(x^*, \varepsilon)$ with $y \succ x^*$. Since y is feasible and strictly preferred to x^*, x^* does not solve the UMP – contradiction. Hence $p \cdot x^* = w$. ∎

Intuitively, the reason "a locally non-satiated consumer spends the whole budget" is that a strictly-cheaper feasible bundle in the neighborhood always exists (that is what local non-satiation guarantees), so leaving even a small slack δ in the budget forfeits utility the consumer could have grabbed. The proof gives an explicit metric on the slack: any δ can be exchanged for a bundle strictly better than x^* using a ball of radius $\delta/(2\|p\|)$.

2.5 Convex preferences

Convexity of the consumption set was defined in Section 1.2. Convexity of preferences is a different property: it constrains the *indifference sets* of $\succsim$, not the domain.

Axiom (Convexity). $\succsim$ is **convex** if, for every $x \in X$, the upper-contour set $U(x) = \{y \in X : y \succsim x\}$ is a convex subset of X.

Axiom (Strict convexity). $\succsim$ is **strictly convex** if, for every $x \in X$ and every pair $y, z \in U(x)$ with $y \neq z$ and every $\alpha \in (0, 1)$, the convex combination $\alpha y + (1 - \alpha)z$ lies in the *strict* upper-contour set $\{w : w \succ x\}$.

Intuitively, convex preferences favor *mixtures*: combining two bundles that are each at least as good as x produces a bundle that is also at least as good as x. Strict convexity says the mixture is strictly better than either component (as long as the two components differ). The classic economic interpretation is a taste for variety.

Equivalent form. A rational preference relation is convex if and only if the following holds: for every $y, z \in X$ with $y \succsim z$ and every $\alpha \in [0, 1]$,

$$\alpha y + (1 - \alpha)z \succsim z.$$

Strict convexity replaces $\succsim$ with $\succ$ for $\alpha \in (0, 1)$ and $y \neq z$.

Example 2.5. Cobb-Douglas is strictly convex. With $u(x) = x_1^{\alpha} x_2^{1-\alpha}$ for $\alpha \in (0, 1)$, the upper-contour set through x^0 is $\{x : x_1^{\alpha} x_2^{1-\alpha} \geq u(x^0)\}$. The function u is strictly quasi-concave (its logarithm is strictly concave), so its super-level sets are strictly convex.

Example 2.6. Perfect substitutes are convex but not strictly convex. With $u(x) = x_1 + x_2$, the upper-contour set through x^0 is a closed half-plane $\{x : x_1 + x_2 \geq u(x^0)\}$, which is convex. Take $y = (2, 0)$ and $z = (0, 2)$: both give $u = 2$. The mixture $0.5y + 0.5z = (1, 1)$ also gives $u = 2$, so the mixture is only weakly preferred to y, not strictly. Convexity holds; strict convexity fails. Figure 2.2 contrasts the three possibilities in a single row.

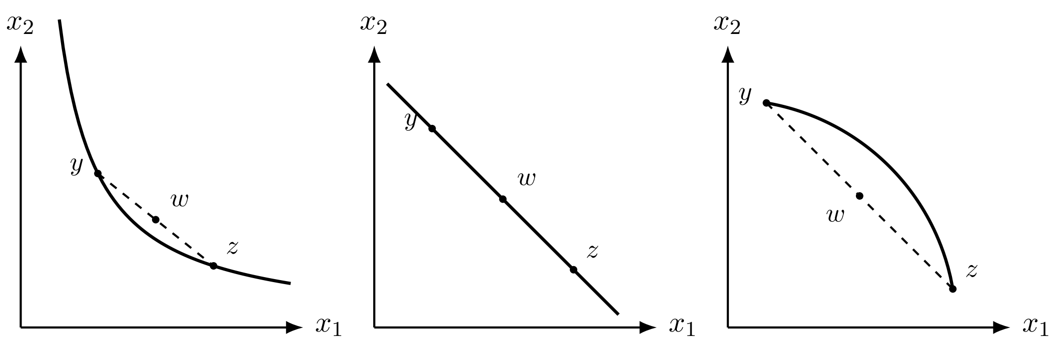

Figure 2.2: Three possibilities for an indifference curve through y and z with $y \sim z$. The mixture $w = 0.5y + 0.5z$ satisfies (a) $w \succ y$ under strict convexity; (b) $w \sim y$ under weak convexity (perfect substitutes); (c) $w \prec y$ when preferences are non-convex.

Convex preferences and quasi-concave utility. A function $u : X \to \mathbb{R}$ is **quasi-concave** if its super-level sets $\{x : u(x) \geq k\}$ are convex for every $k \in \mathbb{R}$. It is **strictly quasi-concave** if the super-level sets are strictly convex (in the sense that any convex combination of different boundary points lies in the interior). If u represents $\succsim$, then

- $\succsim$ is convex $\Leftrightarrow$ u is quasi-concave;
- $\succsim$ is strictly convex $\Leftrightarrow$ u is strictly quasi-concave.

The two axioms are behavioral statements about $\succsim$; the two properties of u are the mathematical objects Chapter 6 uses when it takes first-order conditions. The equivalence lets the theorist move freely between the two languages.

A warning about concavity. Concavity of u is strictly stronger than quasi-concavity of u: every concave function is quasi-concave, but the converse fails ($u(x) = x^3$ on $\mathbb{R}_+$ is quasi-concave but not concave). Convexity of $\succsim$ is a property invariant to monotone transformations of u; concavity of u is not. In practice, convexity is the axiom on $\succsim$, and any strictly increasing transformation of a concave utility representation is quasi-concave. Chapters 3 and 5 return to this point.

Example 2.7. Convexity of Cobb-Douglas: a worked check. On $X = \mathbb{R}^2_{++}$, take the utility function

$$u(x_1, x_2) = x_1^{1/2} x_2^{1/2}.$$

We verify that $\succsim$ (represented by u) is strictly convex, in three steps.
Step 1 (super-level sets). Fix $k > 0$. The super-level set is

$$U_k = \{(x_1, x_2) \in \mathbb{R}^2_{++} : x_1^{1/2} x_2^{1/2} \geq k\} = \{(x_1, x_2) : x_1 x_2 \geq k^2\}.$$

Step 2 (convexity of U_k). Take $y = (y_1, y_2) \in U_k$ and $z = (z_1, z_2) \in U_k$ with $y \neq z$; set $w = ty + (1-t)z$ for $t \in (0, 1)$. We show $w \in U_k$, i.e., $w_1 w_2 \geq k^2$.
By the AM-GM inequality applied to positive numbers,

$$w_1 w_2 = (ty_1 + (1-t)z_1)(ty_2 + (1-t)z_2).$$

Expand the right side:

$$= t^2 y_1 y_2 + t(1-t)(y_1 z_2 + y_2 z_1) + (1-t)^2 z_1 z_2.$$

The middle cross term dominates via the AM-GM step $y_1 z_2 + y_2 z_1 \geq 2\sqrt{y_1 y_2 z_1 z_2}$:

$$w_1 w_2 \;\geq\; t^2 y_1 y_2 + 2t(1-t)\sqrt{y_1 y_2}\sqrt{z_1 z_2} + (1-t)^2 z_1 z_2 \;=\; \left(t\sqrt{y_1 y_2} + (1-t)\sqrt{z_1 z_2}\right)^2.$$

Since $\sqrt{y_1 y_2} \geq k$ and $\sqrt{z_1 z_2} \geq k$, the last expression is at least $(tk + (1-t)k)^2 = k^2$. Hence $w_1 w_2 \geq k^2$, so $w \in U_k$. The set is convex, and $\succsim$ is convex.
Step 3 (strict convexity). If $y \neq z$ (both on the boundary $y_1 y_2 = z_1 z_2 = k^2$), the AM-GM inequality in Step 2 is strict unless $y_1 z_2 = y_2 z_1$, i.e., unless y and z lie on the same ray through the origin. Because $y_1 y_2 = z_1 z_2 = k^2$ and $y \neq z$, we have $y_1 \neq z_1$ and, along the ray, $y_2/y_1 \neq z_2/z_1$; the AM-GM inequality is strict, hence $w_1 w_2 > k^2$, i.e., $w \in \operatorname{int} U_k$, giving

strict convexity.

Numerical illustration: take $k = 2$, so the indifference curve is the hyperbola $x_1x_2 = 4$. Two points on this curve are $y = (1,4)$ and $z = (4,1)$. At $t = 1/2$, $w = (5/2, 5/2)$ and $w_1w_2 = 25/4 = 6.25 > 4 = k^2$. The mixture strictly beats both endpoints: $u(w) = \sqrt{6.25} = 2.5 > 2 = u(y) = u(z)$.

Lesson: strict convexity of the Cobb-Douglas preference is a consequence of the concavity of $\log u$; on any indifference curve, mixing two different bundles lands the consumer strictly above the curve.

Example 2.11. CES convexity for $\rho = -1$. On $X = \mathbb{R}^2_{++}$ take the CES utility with $\rho = -1$:

$$u(x_1, x_2) = \left(x_1^{-1} + x_2^{-1}\right)^{-1} = \frac{x_1x_2}{x_1 + x_2}.$$

(This is the harmonic mean, up to a factor of two.) We verify that the super-level sets are convex.

Step 1 (super-level sets). For $k > 0$,

$$U_k = \left\{(x_1, x_2) \in \mathbb{R}^2_{++} : \frac{x_1x_2}{x_1 + x_2} \geq k\right\} = \left\{(x_1, x_2) : \frac{1}{x_1} + \frac{1}{x_2} \leq \frac{1}{k}\right\}.$$

The second form uses $\dfrac{1}{u} = \dfrac{1}{x_1} + \dfrac{1}{x_2}$, so $u \geq k$ is equivalent to $\dfrac{1}{x_1} + \dfrac{1}{x_2} \leq \dfrac{1}{k}$.

Step 2 (convexity by convexity of $1/x$). The function $g(x) = \dfrac{1}{x}$ is strictly convex on $(0, \infty)$. Take $y, z \in U_k$ and $w = ty + (1-t)z$ with $t \in [0,1]$. Convexity of g gives, coordinate by coordinate,

$$\frac{1}{w_1} \leq t \cdot \frac{1}{y_1} + (1-t) \cdot \frac{1}{z_1}, \qquad \frac{1}{w_2} \leq t \cdot \frac{1}{y_2} + (1-t) \cdot \frac{1}{z_2}.$$

Add these two inequalities:

$$\frac{1}{w_1} + \frac{1}{w_2} \leq t\Big(\frac{1}{y_1} + \frac{1}{y_2}\Big) + (1-t)\Big(\frac{1}{z_1} + \frac{1}{z_2}\Big) \leq t \cdot \frac{1}{k} + (1-t) \cdot \frac{1}{k} = \frac{1}{k}.$$

So $w \in U_k$: the super-level set is convex, and $\succsim$ is convex. Strict convexity of g makes the inequality strict when $y \neq z$, giving strict convexity of $\succsim$.

Numerical illustration: at $k = 1$, $y = (2,2)$ gives $u = 1$; $z = (4, \frac{4}{3})$ gives $u = (4 \cdot \frac{4}{3})/(4 + \frac{4}{3}) = \dfrac{16/3}{16/3} = 1$. At $t = \frac{1}{2}$, $w = (3, \frac{5}{3})$ and $u(w) = \dfrac{3 \cdot 5/3}{3 + 5/3} = \dfrac{5}{14/3} = \dfrac{15}{14} \approx 1.07 > 1$. The mixture strictly beats the endpoints.

Lesson: for CES with $\rho < 1$, the same negative-exponent form makes super-level sets convex through convexity of the coordinate function $1/x^{|\rho|}$; the case $\rho = -1$ is the cleanest computation and the case where the utility function itself is the harmonic mean.

2.6 The marginal rate of substitution

Let $\succsim$ be rational, continuous, strongly monotone, and strictly convex on $X = \mathbb{R}^2_{++}$, and let it be represented by a C^2 utility function u. The indifference curve through x^0 is the level set $\{x \in \mathbb{R}^2_{++} : u(x) = u(x^0)\}$. By the implicit function theorem (Simon and Blume (1994), Chapter 15; or MWG (1995), Section M.E of the Mathematical Appendix), this level set is locally the graph of a smooth decreasing function $x_2(x_1)$, and its slope at x^0 is

$$\mathrm{MRS}_{12}(x^0) \;\equiv\; -\frac{dx_2}{dx_1}\bigg|_{u=u(x^0)} = \frac{\partial u(x^0)/\partial x_1}{\partial u(x^0)/\partial x_2}.$$

The **marginal rate of substitution of good 1 for good 2**, MRS_{12}, is the amount of good 2 the consumer is willing to give up for one extra unit of good 1, holding utility constant. It is the ratio of marginal utilities. Under the axioms above, $\mathrm{MRS}_{12} > 0$ everywhere in $\mathbb{R}^2_{++}$ (both partials are positive), and along any indifference curve MRS_{12} falls as x_1 rises (strict convexity of $\succsim$ produces strictly decreasing MRS_{12} along each curve). This is the classical "diminishing marginal rate of substitution."

Geometrically, the MRS is the negative of the slope of the tangent line to the indifference curve at x^0: it says how much of good 2 the consumer would give up per marginal unit of good 1 to stay on the same curve. A steeper tangent means a larger MRS and a higher willingness to trade good 2 for good 1 at that point.

Example 2.8. MRS of Cobb-Douglas. With $u(x) = x_1^{\alpha} x_2^{1-\alpha}$,

$$\frac{\partial u}{\partial x_1} = \alpha\,\frac{u(x)}{x_1}, \qquad \frac{\partial u}{\partial x_2} = (1-\alpha)\,\frac{u(x)}{x_2},$$

so

$$\mathrm{MRS}_{12}(x) = \frac{\alpha\, x_2}{(1-\alpha)\, x_1}.$$

Along an indifference curve, x_2 falls as x_1 rises, so MRS_{12} falls: diminishing MRS holds.

Invariance to monotone transformations. If $\tilde{u} = f(u)$ for some strictly increasing f, then $\dfrac{\partial \tilde{u}}{\partial x_\ell} = f'(u)\,\dfrac{\partial u}{\partial x_\ell}$; the derivative $f'(u)$ cancels in the ratio, and MRS_{12} is unchanged. The MRS is a property of the preference relation $\succsim$, not of any particular utility representation of it. This invariance is the geometric counterpart to the fact – proved in Chapter 3 – that utility representations are unique only up to strictly increasing transformations.

Example 2.9. Computing the MRS for four utility functions. For each utility function below, we compute $\mathrm{MRS}_{12}(x)$ at the point $x^0 = (2, 2)$ and read off its economic meaning.

- *Cobb-Douglas* $u(x) = x_1^{1/2} x_2^{1/2}$: $\dfrac{\partial u}{\partial x_1} = \dfrac{1}{2} x_1^{-1/2} x_2^{1/2}$ and $\dfrac{\partial u}{\partial x_2} = \dfrac{1}{2} x_1^{1/2} x_2^{-1/2}$. The ratio at $x^0 = (2,2)$ is $\mathrm{MRS}_{12}(2,2) = \dfrac{x_2}{x_1} = 1$. At the equal-consumption point, the consumer values a marginal unit of good 1 exactly the same as a marginal unit of good 2.

- *CES with* $\rho = 1/2$ $u(x) = (x_1^{1/2} + x_2^{1/2})^2$: after computation, $\mathrm{MRS}_{12}(x) = \big(\dfrac{x_2}{x_1}\big)^{1/2}$. At $x^0 = (2,2)$, $\mathrm{MRS}_{12} = 1$; but at $x^0 = (4,1)$, $\mathrm{MRS}_{12} = 1/2$: the elasticity of substitution is $\sigma = 1/(1-\rho) = 2$, so shifts in consumption move the MRS proportionally.

- *Quasi-linear* $u(x) = 2\sqrt{x_1} + x_2$: $\dfrac{\partial u}{\partial x_1} = x_1^{-1/2}$ and $\dfrac{\partial u}{\partial x_2} = 1$. So $\mathrm{MRS}_{12}(x) = 1/\sqrt{x_1}$, independent of x_2. At $x^0 = (2,2)$, $\mathrm{MRS}_{12} = 1/\sqrt{2} \approx 0.707$. Every indifference curve has the same slope at any fixed x_1 – this is what makes them vertical translates.

- *Leontief* $u(x) = \min\{x_1/2, x_2/3\}$: the MRS is undefined at the kink line $x_1/2 = x_2/3$ (or $x_2 = (3/2)x_1$), infinite below the kink line (where $x_1/2$ is binding and only good 2 is wasted), and zero above (where $x_2/3$ is binding). At $x^0 = (2,2)$, we are below the kink ($2/2 = 1 > 2/3$, so $x_2/3$ binds), meaning $\mathrm{MRS}_{12} = 0$: adding good 1 has no marginal value here because the binding component is already good 2.

The four families give four qualitatively different pictures of how the MRS varies with (x_1, x_2): constant slope (perfect substitutes), curved and level-independent (quasi-linear), curved and homothetic (CES/Cobb-Douglas), and kinked (Leontief).

2.7 Four canonical preference families

Four parametric families of preferences appear repeatedly in Chapters 5 and 6 and in applied work.

Cobb-Douglas. $u(x) = x_1^{\alpha_1} \cdots x_L^{\alpha_L}$ with $\alpha_\ell > 0$ and $\sum \alpha_\ell = 1$ (or with free positive α_ℓ and a monotone transformation). Continuous, strongly monotone, strictly convex, homothetic (Chapter 5). Marshallian demand has closed form $x_\ell = \alpha_\ell \dfrac{w}{p_\ell}$.

CES. $u(x) = \left(\sum_\ell \alpha_\ell x_\ell^\rho\right)^{1/\rho}$ for $\rho \in (-\infty, 1)$, $\rho \neq 0$, with $\alpha_\ell > 0$. Continuous, strongly monotone, strictly convex when $\rho < 1$, homothetic. Special cases: $\rho \to 1$ gives perfect substitutes; $\rho \to 0$ gives Cobb-Douglas; $\rho \to -\infty$ gives Leontief. The elasticity of substitution is $\sigma = \dfrac{1}{1-\rho}$, and CES families are often indexed by σ rather than ρ.

Quasi-linear. $u(x_1, x_2) = v(x_1) + x_2$ with v strictly increasing and concave. Continuous, strongly monotone (in good 2; in good 1 as long as $v' > 0$), convex if v is concave, strictly convex if v is strictly concave. *Not* homothetic; income effects are zero for good 1 (the good with the non-linear utility component). This is the workhorse of partial-equilibrium consumer surplus analysis.

Leontief. $u(x) = \min\{x_1/a_1, \ldots, x_L/a_L\}$ with $a_\ell > 0$. Continuous, strictly monotone but not strongly monotone (Example 2.4), convex but not strictly convex (the indifference curves have flats). Marshallian demand is $x_\ell = a_\ell\, w / \left(\sum_k a_k p_k\right)$.

The four families span the parameter space of substitutability between goods: perfect substitutes (linear), imperfect substitutes (Cobb-Douglas and CES), no substitutes (Leontief). Chapter 5 develops the special-structure properties (homothetic, additively separable) that let each family be manipulated symbolically. Figure 2.3 shows one representative indifference curve for each family.

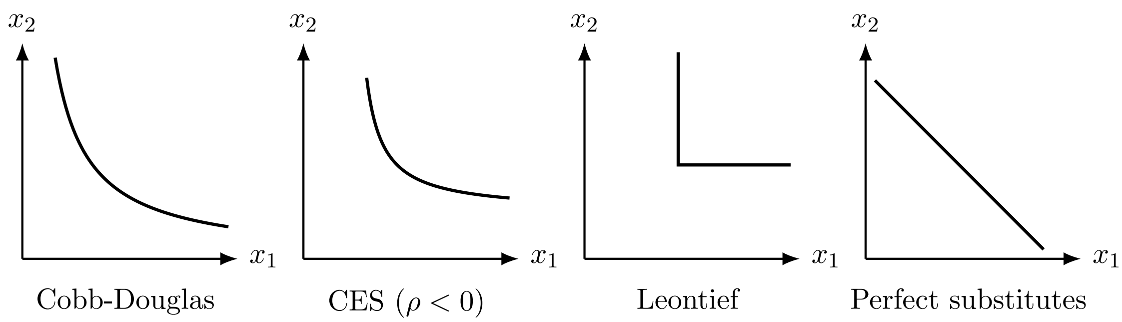

Figure 2.3: One representative indifference curve for each of the four canonical families of Section 2.7. Cobb-Douglas gives strictly convex, smooth curves with axis asymptotes; CES with $\rho < 0$ approaches Leontief; Leontief has the classic L-shaped kink at the fixed-proportions ray; perfect substitutes give a straight line with constant MRS.

Comparison table. The four canonical families differ in three geometric respects: the shape of the indifference curve (smooth-strictly-convex, straight-line, kinked, smooth-strictly-convex-with-different-elasticity); their behavior under wealth changes (constant budget shares, price-share-dependent shares, fixed-ratio consumption, or income-independent good-1 demand); and their elasticity of substitution ($\sigma = 1$, $\sigma \in (0, \infty) \setminus \{1\}$, $\sigma = 0$, $\sigma = \infty$).

Family	Utility	Marshallian demand for good ℓ	Wealth elasticity	σ
Cobb-Douglas	$x_1^\alpha x_2^{1-\alpha}$	$\alpha_\ell w/p_\ell$	1 (unit)	1
CES	$(\sum \alpha_k x_k^\rho)^{1/\rho}$	$\dfrac{\alpha_\ell^\sigma p_\ell^{-\sigma} w}{\sum_k \alpha_k^\sigma p_k^{1-\sigma}}$	1 (unit)	$\dfrac{1}{1-\rho}$
Quasi-linear	$v(x_1) + x_2$	$x_1 = (v')^{-1}(p_1/p_2)$, indep. of w	0 for good 1; residual for numeraire	N/A
Leontief	$\min\{x_1/a_1, x_2/a_2\}$	$a_\ell w/(a_1 p_1 + a_2 p_2)$	1 (unit)	0
Perfect substitutes	$x_1 + x_2$	Corner: $x = (w/p_{\ell_{\min}}, 0)$ for cheaper good	N/A (corner)	∞

Intuitively, the four families span the plausible "structural" possibilities for consumer demand. Cobb-Douglas and CES cover the smooth-substitution middle; Leontief covers perfect complements; perfect substitutes cover perfect substitution; and quasi-linear covers the applied partial-equilibrium special case where the numeraire absorbs all income effects. Chapters 5 and 6 use one of these families in every worked example.

2.8 Experimental evidence

Every axiom in this chapter has been tested in the laboratory. The evidence is mixed: some axioms are almost never violated by sophisticated subjects; others fail systematically and in predictable ways.

Continuity. Continuity is essentially untestable directly: it constrains behavior at unobservable limit points. Indirect evidence comes from utility-elicitation experiments, where subjects are asked to state indifference prices for sequences of bundles. Andreoni and Miller (2002) elicited Marshallian demand for two goods across 50 budget lines and found demand almost always continuous in prices, consistent with continuous preferences. Failures of continuity are largely confined to lexicographic preferences over prizes with qualitatively different attributes (money vs. safety; health vs. leisure) – see Payne, Bettman, and Johnson (1993) for the psychology-of-choice evidence on lexicographic decision rules.

Monotonicity. Monotonicity fails most visibly for addictive or aversive goods (nicotine, calories under diet-target framing, environmental pollutants). In the classical consumer setting – food, clothing, transportation – monotonicity holds nearly universally in laboratory choice data. Andreoni and Miller (2002) rejected weak monotonicity in fewer than 2% of budget-line pairs.

Convexity. Choi, Fisman, Gale, and Kariv (2007) tested convexity of preferences using a laboratory design in which subjects allocated tokens across two contingent goods over 50 randomly drawn budget lines. Under convexity, the demand correspondence is convex-valued and the utility function is quasi-concave. About 70% of subjects generated choices consistent with a quasi-concave utility function passing standard GARP-type tests (Chapter 4). The remaining 30% exhibited choices with non-convex indifference sets, most often at allocations near the axes.

Marginal rate of substitution: diminishing? Field evidence from housing (bedrooms vs. square footage), labor supply (leisure vs. consumption), and portfolio choice (risky vs. safe assets) shows MRS declining along observed indifference curves, consistent with strict convexity. Halevy, Persitz, and Zrill (2018) used laboratory data across 25 budget lines to estimate individual-level MRS profiles and found strictly declining MRS for the majority of subjects.

Table 2.1 summarizes the leading laboratory tests of the structural axioms.

Axiom	**Study**	**Finding**
Continuity	Andreoni-Miller (2002)	Demand almost always continuous in prices
Monotonicity	Andreoni-Miller (2002)	Fewer than 2% of budget pairs violate
Convexity	Choi-Fisman-Gale-Kariv (2007)	$\approx$ 70% of subjects have quasi-concave utility
Diminishing MRS	Halevy-Persitz-Zrill (2018)	Strictly declining MRS for the majority
Lexicographic exceptions	Payne-Bettman-Johnson (1993)	Common in high-stakes multi-attribute choice

Interpretation. Continuity and monotonicity fail rarely and mostly in special environments. Convexity fails more often – roughly one subject in three – and the failures concentrate at extreme allocations. Diminishing MRS holds for most subjects. The pattern is consistent with treating the structural axioms as good working assumptions for typical consumer decisions while noting the environments (multi-attribute choice; corner allocations; addictive goods) where they can break.

2.9 Summary

- **Continuity** requires the upper- and lower-contour sets to be closed in $\mathbb{R}^L$. It rules out preference jumps and is the axiom Chapter 3 adds to completeness and transitivity to guarantee a continuous utility representation.

- Three **monotonicity** conventions – monotone, strongly monotone, strictly monotone – differ in how they treat bundles that dominate coordinate-wise. Strong monotonicity is the strongest and holds for Cobb-Douglas and CES.

- **Local non-satiation** is the minimum "more is better" hypothesis that lets Chapter 6 conclude the budget constraint binds at the utility-maximizing bundle.

- **Convex** preferences have convex upper-contour sets, equivalent to quasi-concavity of any utility representation. **Strict convexity** strengthens the requirement to strict quasi-concavity and rules out flat indifference segments.

- The **marginal rate of substitution** $\mathrm{MRS}_{12} = (\frac{\partial u}{\partial x_1})/(\frac{\partial u}{\partial x_2})$ is the slope of the indifference curve. It is invariant to strictly increasing transformations of u and, under strict convexity, decreasing along each indifference curve.

- Four canonical parametric families – **Cobb-Douglas**, **CES**, **quasi-linear**, and **Leontief** – span the full range of substitutability and are the working examples of Chapters 5 and 6.

- Laboratory evidence supports continuity and monotonicity almost universally; convexity holds for roughly 70% of subjects; diminishing MRS holds for most.

2.10 Guided exercise

Exercise (Guided). Consider the preference relation on $X = \mathbb{R}^2_+$ represented by the utility function

$$u(x_1, x_2) = x_1 + 2\sqrt{x_2}.$$

Determine (a) whether $\succsim$ is continuous; (b) whether it is strongly monotone; (c) whether it is strictly convex; (d) its marginal rate of substitution.

Solution.

(a) Continuity. u is the sum of two continuous functions (x_1 and $2\sqrt{x_2}$), so u is continuous. The upper-contour set of x^0 is the closed super-level set $\{x : x_1 + 2\sqrt{x_2} \geq k\}$ with $k = u(x^0)$, and the lower-contour set is the closed sub-level set. Both are closed, so $\succsim$ is continuous.

(b) Monotonicity. $\dfrac{\partial u}{\partial x_1} = 1 > 0$ and $\dfrac{\partial u}{\partial x_2} = 1/\sqrt{x_2} > 0$ for $x_2 > 0$. If $y > x$, then either $y_1 > x_1$ (and $y_2 \geq x_2$) or $y_2 > x_2$ (and $y_1 \geq x_1$). Either way, $u(y) > u(x)$, so $y \succ x$. Hence $\succsim$ is strongly monotone on $\mathbb{R}^2_{++}$ (with a small caveat at the $x_2 = 0$ boundary, where $\dfrac{\partial u}{\partial x_2}$ blows up but the monotonicity conclusion still holds).

(c) Strict convexity. The Hessian of u is

$$D^2 u = \begin{pmatrix} 0 & 0 \\ 0 & -\frac{1}{2} x_2^{-3/2} \end{pmatrix}.$$

It is negative semi-definite (both eigenvalues ≤ 0) but not negative definite (the eigenvalue at position (1,1) is zero). So u is concave, hence quasi-concave, hence $\succsim$ is convex. To check strict convexity, take $y = (0, 4)$ and $z = (4, 0)$: $u(y) = 4$ and $u(z) = 4$. The mixture $0.5y + 0.5z = (2, 2)$ gives $u = 2 + 2\sqrt{2} \approx 4.83 > 4$, strictly better than both components. This is one test of strict convexity; a more careful proof would show the strict inequality holds for every pair $y \neq z$ in $U(x)$ and every $\alpha \in (0, 1)$. The result follows from strict concavity of $2\sqrt{x_2}$ in x_2 combined with linearity in x_1; the mixture strictly raises u whenever the two bundles differ in x_2.

(d) MRS. Directly from the partial derivatives,

$$\text{MRS}_{12}(x) = \frac{\partial u/\partial x_1}{\partial u/\partial x_2} = \frac{1}{1/\sqrt{x_2}} = \sqrt{x_2}.$$

The MRS depends only on x_2 and grows with x_2. Along an indifference curve $u(x_1, x_2) = k$, as x_1 rises, x_2 must fall (utility is constant), so $\text{MRS}_{12} = \sqrt{x_2}$ falls. Diminishing MRS holds.

Lesson. Utility functions of the quasi-linear form $u(x) = x_1 + v(x_2)$ with strictly concave v share three properties that make them the workhorse of applied partial-equilibrium analysis: (i) strong monotonicity, (ii) strict convexity, (iii) MRS_{12} depends only on x_2, so income effects on the demand for good 1 vanish. Chapter 5 develops the quasi-linear class in general.

Chapter 3 - Utility representation

3.1 Introduction

Chapters 1 and 2 developed the language and the axioms of a rational, continuous, monotone, convex preference relation $\succsim$ on the consumption set $X = \mathbb{R}^L_+$. This chapter proves the classical theorem that ties that language to the tool the rest of the book depends on: a real-valued utility function.

The main result of the chapter is **Debreu's theorem** (Debreu 1954): a continuous, rational preference relation on $\mathbb{R}^L_+$ admits a continuous utility representation. The theorem promotes $\succsim$ – an object on $X \times X$ – into a function on X, and the promotion loses nothing: recovering $\succsim$ from u is trivial ($x \succsim y$ iff $u(x) \geq u(y)$), and the theorem guarantees such a u exists whenever the underlying axioms hold.

Two features of the theorem shape everything downstream. First, continuity of $\succsim$ (Chapter 2) is essential; without it, the lexicographic preferences of Section 1.7 give a rational preference relation with no utility representation at all – a fact this chapter proves by a cardinality argument. Second, the representation is not unique: any strictly increasing transformation of a utility function that represents $\succsim$ also represents $\succsim$. Utility is *ordinal*, not cardinal. Section 3.6 develops the uniqueness statement precisely.

Section 3.2 defines what a utility representation is and recalls the finite case. Section 3.3 states Debreu's theorem and sketches its proof. Section 3.4 proves the lexicographic non-representation result. Section 3.5 develops the constructive alternative – Rader's theorem – when the axioms are weaker. Section 3.6 formalizes the ordinal-cardinal distinction and the uniqueness result. Section 3.7 collects the leading empirical findings on eliciting utility functions from choice data. Section 3.8 summarizes.

3.2 Definition of a utility representation

The finite-set representation of Proposition 1.4 is the model to generalize. We repeat the definition here.

Definition 3.1. A function $u : X \to \mathbb{R}$ **represents** the preference relation $\succsim$ on X if, for every pair $x, y \in X$,

$$x \succsim y \quad \Longleftrightarrow \quad u(x) \geq u(y).$$

The definition is symmetric: passing from $\succsim$ to u loses nothing (rankings can be reconstructed from u), and passing from u to $\succsim$ loses nothing (the ordering $\geq$ on $\mathbb{R}$ is complete and transitive). The question the chapter answers is: given a rational $\succsim$ on an infinite X, when does a u with this property exist?

Example 3.1. Verifying that a candidate function represents a preference. Suppose $\succsim$ on $\mathbb{R}^2_+$ satisfies the following pairwise comparisons for the three sample bundles $A = (1,4)$, $B = (2,2)$, $C = (4,1)$:

$$A \sim B \sim C \quad \text{and} \quad D = (2,3) \succ A.$$

Does the candidate function $u(x_1, x_2) = x_1 x_2$ represent this $\succsim$? Compute utilities: $u(A) = 4$, $u(B) = 4$, $u(C) = 4$, $u(D) = 6$. Because $u(A) = u(B) = u(C) = 4$, the definition requires $A \sim B \sim C$ ✓. Because $u(D) = 6 > 4 = u(A)$, the definition requires $D \succ A$ ✓. So u represents $\succsim$ on this restricted set. Contrast with $\tilde{u}(x_1, x_2) = x_1 + x_2$: $\tilde{u}(A) = 5$, $\tilde{u}(B) = 4$, $\tilde{u}(C) = 5$; the definition would require $A \sim C \succ B$, which does not match $A \sim B$. So $\tilde{u}$ does *not* represent $\succsim$ (even though $\tilde{u}$ is a perfectly good utility function for a different preference relation, namely perfect substitutes). The point: whether a given u represents a given $\succsim$ is a testable question at every pair of alternatives. Figure 3.1 pictures the utility function as a mapping from X to $\mathbb{R}$ that preserves the ranking of three bundles $x \succ y \succ z$.

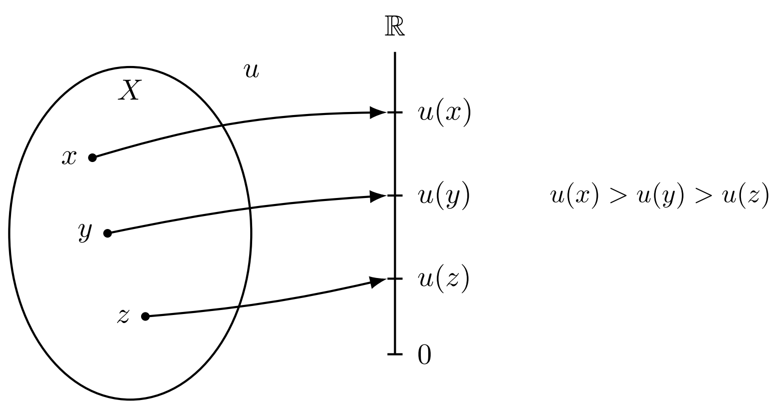

Figure 3.1: A utility function $u : X \to \mathbb{R}$ maps each bundle to a real number, preserving the ranking: $x \succ y \succ z$ in X translates to $u(x) > u(y) > u(z)$ on the real line.

Preview. Three answers, in order of strength:

1. If $\succsim$ is rational (complete and transitive) and X is *finite*, then a utility representation exists (Proposition 1.4).

2. If $\succsim$ is rational and *continuous* on $X = \mathbb{R}^L_+$ (or, more generally, on any connected separable topological space), then a *continuous* utility representation exists (Debreu's theorem, Section 3.3).

3. If $\succsim$ is rational and X is *countable*, then a utility representation exists (Rader's theorem, Section 3.5); continuity is not required, but the representation need not be continuous either.

The lexicographic preferences of Section 1.7 sit outside all three cases: $X = \mathbb{R}^2_+$ is uncountable and connected, but $\succsim^L$ fails continuity, so Debreu does not apply. Section 3.4 shows that no u exists in this case, so continuity of $\succsim$ is not just sufficient but necessary for the theorem in an appropriate sense.

3.3 Debreu's theorem

Theorem 3.1 (Debreu, 1954). Let X be a nonempty connected separable topological space, and let $\succsim$ be a rational preference relation on X. Then $\succsim$ is continuous if and only if it admits a continuous utility representation $u : X \to \mathbb{R}$.

The forward direction (from u to $\succsim$) is easy: if a continuous u represents $\succsim$, then the upper- and lower-contour sets are pre-images of closed intervals $[u(x), \infty)$ and $(-\infty, u(x)]$ under a continuous function, so both are closed. The reverse direction (from $\succsim$ to u) is the content of the theorem.

The proof splits naturally into two parts. First, existence of a countable order-dense subset (a set with the "betweenness" property that any strictly ranked pair has a set element strictly between them). Second, construction of the utility function on the whole space by extension from that countable subset. We state and prove each part separately.

Proposition 3.0 (order-dense subset). *Let $\succsim$ be a rational, continuous preference relation on a separable, connected metric space X. There exists a countable set $D \subseteq X$ such that for every pair $x, y \in X$ with $x \succ y$, there is $d \in D$ with $x \succsim d \succsim y$.*

Proof of Proposition 3.0. Let D_0 be any countable dense subset of X (which exists by separability of $\mathbb{R}^L$, taking $D_0 = \mathbb{Q}^L \cap X$; this set is countable and dense in X). Take $D = D_0$; we claim it is order-dense. Fix $x \succ y$; we produce $d \in D$ with $x \succsim d \succsim y$.
Because $U(y)$ and $L(x)$ are closed (continuity of $\succsim$) and both contain y and x respectively but their union is not all of X in general, consider the set $A = U(y) \cap L(x) = \{z : x \succsim z \succsim y\}$. This set contains both x and y; by connectedness of X and closedness of A (intersection of two closed sets), A is a nonempty closed set. We claim A has nonempty interior. If not, then A is nowhere dense; but x and y are both interior candidates when $x \succ y$ strictly. Specifically, since $x \succ y$, the strict lower-contour set $\{z : x \succ z\}$ is open (complement of $U(x)$) and contains y; the strict upper-contour set $\{z : z \succ y\}$ is open and contains x. Their intersection $\{z : x \succ z \succ y\}$ is open and contains a nonempty interior segment along any

continuous path from y to x (using connectedness). By density of D_0, there is $d \in D_0$ in this open set, hence $x \succsim d \succsim y$ (in fact $x \succ d \succ y$). ■

Proof of Debreu's theorem (reverse direction). We construct a utility representation $u : X \to \mathbb{R}$ in five steps.

Step 1 (countable order-dense reference set). By Proposition 3.0, consider a countable order-dense subset $D = \{d_1, d_2, \ldots\} \subseteq X$.

Step 2 (assign utilities to D). Set $u(d_1) = 0$. Inductively, given $u(d_1), \ldots, u(d_n)$, define $u(d_{n+1})$ by the following case analysis:

- If $d_{n+1} \sim d_j$ for some $j \leq n$, set $u(d_{n+1}) = u(d_j)$.
- Otherwise, by completeness and transitivity applied to the pairs (d_{n+1}, d_j) for $j = 1, \ldots, n$, the numbers $\{u(d_j) : d_j \prec d_{n+1}\}$ and $\{u(d_j) : d_j \succ d_{n+1}\}$ partition $\{u(d_1), \ldots, u(d_n)\} \setminus \{u(d_j) : d_j \sim d_{n+1}\}$. Let $\underline{u} = \sup\{u(d_j) : d_j \prec d_{n+1}\}$ (or $-\infty$ if empty) and $\overline{u} = \inf\{u(d_j) : d_j \succ d_{n+1}\}$ (or $+\infty$ if empty). Set

$$u(d_{n+1}) = \begin{cases} 1 + \underline{u}, & \overline{u} = +\infty, \\ \underline{u} - 1, & \underline{u} = -\infty, \\ \dfrac{\underline{u} + \overline{u}}{2}, & \text{otherwise.} \end{cases}$$

By construction, this preserves the ranking on D: $d_i \succsim d_j$ if and only if $u(d_i) \geq u(d_j)$, for all i, j enumerated so far. The verification is by induction on $\max(i, j)$.

Step 3 (extend to X). For each $x \in X$, define

$$u(x) = \sup\{u(d) : d \in D,\ x \succsim d\}.$$

The set on the right is nonempty (take any $d \in D$ with $x \succsim d$; existence is guaranteed by density of D and closedness of $L(x)$) and bounded above by any $u(d')$ for $d' \succ x$, so the supremum is finite. This defines $u : X \to \mathbb{R}$.

Step 4 (representation). We show $x \succsim y \Leftrightarrow u(x) \geq u(y)$.

($\Rightarrow$) If $x \succsim y$, then by transitivity, every $d \in D$ with $y \succsim d$ also satisfies $x \succsim d$. Hence $\{u(d) : y \succsim d, d \in D\} \subseteq \{u(d) : x \succsim d, d \in D\}$, and taking suprema gives $u(y) \leq u(x)$.

($\Leftarrow$) Suppose $u(x) \geq u(y)$ but, for contradiction, $y \succ x$. By order-density (Proposition 3.0), there is $d \in D$ with $y \succsim d \succsim x$; by strictness $y \succ x$ and the density argument, we may take $d \in D$ with $y \succ d \succ x$ (choosing d inside the open interval $\{z : y \succ z \succ x\}$, which is a nonempty open set by continuity). Then $y \succsim d$ gives $u(y) \geq u(d)$ (from the forward direction applied to $y \succsim d$), and $d \succ x$ combined with density gives (by the same argument applied to $d \succ x$) that we can find $d' \in D$ with $d \succ d' \succsim x$ and $u(d') > u(x)$; hence $u(y) \geq u(d) > u(d') \geq u(x)$, so $u(y) > u(x)$, contradicting $u(x) \geq u(y)$.

Step 5 (continuity of u). We show u is continuous on X. Take $x_n \to x$ in X; we show $u(x_n) \to u(x)$. Fix $\varepsilon > 0$. By order-density, choose $d^+, d^- \in D$ with

$$u(x) - \varepsilon < u(d^-) < u(x) < u(d^+) < u(x) + \varepsilon, \qquad d^- \prec x \prec d^+.$$

By continuity of $\succsim$, the sets $\{z : z \succ d^-\}$ and $\{z : z \prec d^+\}$ are open (complements of closed contour sets) and both contain x. Their intersection is an open neighborhood of x; for n large, x_n lies in this neighborhood, i.e., $d^- \prec x_n \prec d^+$. By Step 4, $u(d^-) < u(x_n) < u(d^+)$, hence $|u(x_n) - u(x)| < \varepsilon$. Since $\varepsilon > 0$ was arbitrary, $u(x_n) \to u(x)$. ■

Intuitively, the construction proceeds in two stages. First, we score a countable "skeleton" of alternatives inside X, using the midpoint rule to place each new element between its neighbors in the ranking. Second, we extend the score to every point in X by taking a supremum over the skeleton – geometrically, each $x \in X$ receives the score of the highest skeleton point it weakly beats. Continuity of $\succsim$ turns this scoring into a continuous function on X because arbitrarily fine cuts of the skeleton are available around every point.

Corollary 3.1. Every rational, continuous, monotone preference relation on $X = \mathbb{R}^L_+$ admits a continuous utility representation. Under strict monotonicity, the representation can be taken to be strictly increasing in each coordinate. Under strict convexity, the representation can be taken to be strictly quasi-concave.

Proof. The main theorem gives existence of a continuous u representing $\succsim$. Under strict monotonicity: if $y \gg x$ then $y \succ x$, hence $u(y) > u(x)$; so u is strictly increasing in each coordinate direction on the interior of X. Under strict convexity of $\succsim$: for any two indifference-tied bundles $y \sim z$ and any $t \in (0, 1)$, the mixture $ty + (1 - t)z \succ y$; hence $u(ty + (1 - t)z) > u(y) = u(z)$, and u is strictly quasi-concave. ■

Intuitively, once Debreu delivers *a* continuous utility function, the additional axioms of Chapter 2 (monotonicity, strict monotonicity, convexity, strict convexity) each add a corresponding property to u almost automatically: the axiom describes how $\succsim$ behaves, and any utility representation must inherit the corresponding behavior. This is why applied work rarely worries about the passage from $\succsim$ to u: whatever behavior we assume on $\succsim$ carries over cleanly to a continuous u that we can differentiate and manipulate.

The corollary is what Chapters 4–6 use in practice. Whenever the axioms of Chapter 2 hold, we may write down a utility function u and manipulate it symbolically, secure that the resulting statements translate back to statements about $\succsim$.

Example 3.2. Debreu applied to CES preferences. On $X = \mathbb{R}^2_{++}$ take $\succsim$ defined via the CES utility

$$u(x) = \left(x_1^{1/2} + x_2^{1/2}\right)^2.$$

This u is continuous, strongly monotone, and strictly quasi-concave (it inherits these from the CES family of Chapter 2). Verify Debreu's hypotheses on $\succsim$:

- *Rationality.* $\succsim$ is defined by $x \succsim y \Leftrightarrow u(x) \geq u(y)$. Since $\geq$ on $\mathbb{R}$ is complete and transitive, $\succsim$ inherits both properties.
- *Continuity.* u is continuous, so the pre-images of $[u(x^0), \infty)$ and $(-\infty, u(x^0)]$ under u are closed. Hence $U(x^0)$ and $L(x^0)$ are closed for every x^0.

Debreu's theorem asserts the existence of a continuous utility representation of $\succsim$. In this case the representation is already given: u itself. Any strictly increasing transformation, e.g. $\tilde{u} = u^{1/2} = x_1^{1/2} + x_2^{1/2}$, is another continuous representation. The theorem here is a check: the CES preference already comes with an explicit continuous utility, and Debreu's construction (Step 1–Step 5 of the proof) is not needed for this concrete example. The construction earns its place on abstract preference relations for which no closed-form utility is on offer.

Example 3.5. Debreu's construction on $X = [0, 1]$ with the natural order. Let $X = [0, 1]$ and $\succsim$ be the natural order $x \succsim y \Leftrightarrow x \geq y$. This $\succsim$ is rational and continuous. Debreu's theorem promises a continuous utility representation. We walk through the construction on four dense reference points $Q_0 = \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$ to see how a utility level is assigned to each point of X.

Step 1 (take a countable dense subset). Take Q_0 as the initial reference set. It is a finite approximation to the dyadic rationals $D = \{k/2^n : 0 \leq k \leq 2^n, n \in \mathbb{N}\}$, which are dense in $[0, 1]$.

Step 2 (assign utility to each reference point). For each $q \in Q_0$ define $u(q) = q$ directly. This is well-defined and strictly increasing on Q_0: $u(0) = 0 < u(\frac{1}{4}) = \frac{1}{4} < u(\frac{1}{2}) = \frac{1}{2} < u(\frac{3}{4}) = \frac{3}{4} < u(1) = 1$.

Step 3 (extend to all of X by monotonicity and continuity). For any $x \in [0, 1]$, let

$$u(x) = \sup\{q \in D : x \succsim q\} = \sup\{q \in D : q \leq x\}.$$

On our reference set, this recovers the assignments in Step 2. For a general $x \in [0, 1]$, density of D in $[0, 1]$ gives $u(x) = x$, which is a continuous representation of $\succsim$. For example, at $x = 0.31$: $u(0.31) = \sup\{q \in D : q \leq 0.31\} = 0.3125$ if we allow the dyadic $\frac{5}{16} = 0.3125$; in the limit over all of D, $u(0.31) = 0.31$.

Reading: Debreu's proof is not the fastest way to write down a utility function on $[0,1]$ (the identity $u(x) = x$ works). Its content is the general *recipe*: pick a countable dense set, assign real numbers respecting the order, then extend by continuity. On abstract X with no linear structure, that recipe is what makes utility exist at all.

3.4 The lexicographic non-representation theorem

The lexicographic preferences $\succsim^L$ of Section 1.7 are rational (Proposition 1.5) but not continuous (Example 2.2). Debreu's theorem does not apply. In fact, no utility representation exists at all.

Proposition 3.1 (Debreu, 1954). There is no function $u : \mathbb{R}^2_+ \to \mathbb{R}$ that represents the lexicographic preference relation $\succsim^L$.

Proof. Suppose, toward a contradiction, that such a function $u : \mathbb{R}^2_+ \to \mathbb{R}$ exists and represents $\succsim^L$.
Step 1 (constructing the two bundles). For each $t \in \mathbb{R}_+$, consider the pair of bundles

$$a_t = (t, 0), \qquad b_t = (t, 1).$$

Both bundles have the same first coordinate t, and $b_{t,2} = 1 > 0 = a_{t,2}$. By the definition of $\succsim^L$ (Section 1.7): when $x_1 = y_1$, the ranking is by the second coordinate; hence $b_t \succ^L a_t$. Because u represents $\succsim^L$, this strict preference gives

$$u(b_t) > u(a_t).$$

Step 2 (interpolating a rational number). Because $\mathbb{Q}$ is dense in $\mathbb{R}$, in every open interval $(u(a_t), u(b_t)) \subset \mathbb{R}$ there is at least one rational number. Using the axiom of choice, pick one such rational and call it $r_t \in \mathbb{Q}$:

$$u(a_t) < r_t < u(b_t).$$

This defines a map

$$r : \mathbb{R}_+ \to \mathbb{Q}, \qquad t \mapsto r_t.$$

Step 3 (strict monotonicity of the map $t \mapsto r_t$). Take any $t', t \in \mathbb{R}_+$ with $t' > t$. We compare $b_t = (t, 1)$ with $a_{t'} = (t', 0)$: the first coordinates satisfy $t' > t$, and the lexicographic definition ranks the bundle with the larger first coordinate strictly higher regardless of the second, so

$$a_{t'} \succ^L b_t \qquad \Longrightarrow \qquad u(a_{t'}) > u(b_t).$$

Chaining the inequalities from Steps 1–3:

$$r_{t'} \; > \; u(a_{t'}) \; > \; u(b_t) \; > \; r_t.$$

Hence $r_{t'} > r_t$ whenever $t' > t$; the map $t \mapsto r_t$ is strictly increasing, and in particular injective.

Step 4 (cardinality contradiction). An injective map from a set A into a set B requires $|A| \le |B|$. Here $A = \mathbb{R}_+$ is uncountable and $B = \mathbb{Q}$ is countable ($|\mathbb{Q}| = \aleph_0 < 2^{\aleph_0} = |\mathbb{R}|$). Therefore no injective map $\mathbb{R}_+ \to \mathbb{Q}$ can exist. Contradiction. Hence no utility u representing $\succsim^L$ can exist. ■

The proof is one of the cleanest applications of a cardinality argument in economic theory. The lexicographic ordering has one indifference class per point in $\mathbb{R}^2_+$; no function on $\mathbb{R}$ can distinguish uncountably many indifference classes, so no utility representation can exist.

Example 3.6. A non-representable semi-order on $\mathbb{R}_+$. Take $X = \mathbb{R}_+$ and define the strict part of $\succsim$ by $x \succ y \Leftrightarrow x > y + 1$; the indifference part is $x \sim y \Leftrightarrow |x - y| \le 1$. This is a semi-order in the sense of Example 1.11 with threshold 1 instead of 5. We show that no utility function $u : \mathbb{R}_+ \to \mathbb{R}$ can represent both parts of $\succsim$ jointly.

Step 1 (a would-be representation). If a utility u represents $\succsim$, then $x \succ y \Leftrightarrow u(x) > u(y)$ and $x \sim y \Leftrightarrow u(x) = u(y)$. Consider the chain $0, 1, 2, 3$. Adjacent pairs are within 1, so $0 \sim 1 \sim 2 \sim 3$, hence $u(0) = u(1) = u(2) = u(3)$.

Step 2 (the contradiction). Yet $3 - 0 = 3 > 1$, so $3 \succ 0$, which forces $u(3) > u(0)$. The two conclusions $u(3) = u(0)$ and $u(3) > u(0)$ are inconsistent, so no such u exists.

Reading: even a preference relation with a single-parameter threshold can fail to be representable by any real-valued function, because transitivity of indifference is precisely what a real-valued utility enforces. The lesson generalizes: any relation on X with $\succ$ transitive but $\sim$ intransitive escapes the reach of a utility representation, even when X is countable (take $X = \mathbb{N}$ in the construction above). Figure 3.2 pictures the "cliff" the proof exploits: at every value of $t \in \mathbb{R}_+$ the pair $a_t = (t, 0)$ and $b_t = (t, 1)$ requires a fresh utility gap $u(a_t) < r_t < u(b_t)$, and no countable set of rationals can supply uncountably many such r_t.

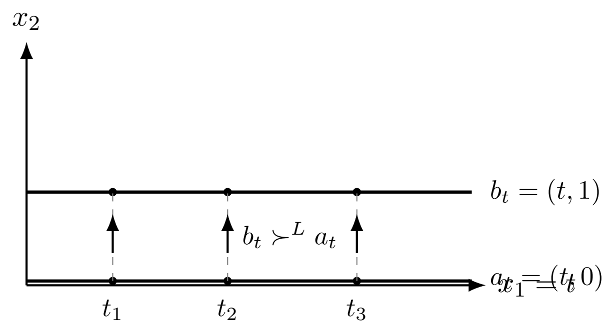

Figure 3.2: Debreu's non-representation proof for $\succsim^L$. At every $t \in \mathbb{R}_+$ the strict preference $b_t \succ^L a_t$ forces $u(a_t) < u(b_t)$, so between them lies a rational number $r_t \in \mathbb{Q}$. The map $t \mapsto r_t$ is strictly increasing, embedding the uncountable $\mathbb{R}_+$ into the countable $\mathbb{Q}$ – impossible.

Interpretation. The failure of representation is not a technical curiosity. Lexicographic decision rules are common in multi-attribute choice: legal thresholds ("convict only if the evidence exceeds the standard"), medical triage ("treat the most acute condition first"), and admissions committees ("rank by test score, break ties with GPA"). The Proposition says these rules cannot be summarized by a single-number score. The finite-attribute Payne-Bettman-Johnson (1993) literature develops the practical decision-analytic tools that replace utility in such settings.

3.5 Rader's theorem: representation without continuity

When X is countable (say, a lattice of integer bundles or a finite menu with a finite number of attributes), continuity is not required for a utility representation to exist.

Theorem 3.2 (Rader, 1963). Let X be a countable set and let $\succsim$ be a rational preference relation on X. Then $\succsim$ admits a utility representation $u : X \to \mathbb{R}$.

Proof sketch. Enumerate $X = \{x_1, x_2, \ldots\}$ and define

$$u(x_n) = \sum_{m : x_n \succsim x_m} \frac{1}{2^m}.$$

The series converges (its terms are bounded above by $\sum_m 2^{-m} = 1$), so u takes real values. If $x_n \succsim x_k$, then every index m counted by $u(x_k)$ (i.e., every m with $x_k \succsim x_m$) is also counted by $u(x_n)$ by transitivity, so $u(x_n) \geq u(x_k)$.
Conversely, if $x_n \not\succsim x_k$ (equivalently, $x_k \succ x_n$ by completeness), then k is counted by $u(x_k)$ (since $x_k \succsim x_k$) but not by $u(x_n)$; by transitivity every m with $x_n \succsim x_m$ satisfies $x_k \succsim x_m$, so

$u(x_k) \geq u(x_n) + 2^{-k} > u(x_n)$. Contrapositively, $u(x_n) \geq u(x_k)$ implies $x_n \succsim x_k$. Combined with the forward direction, u represents $\succsim$. ■

Rader's construction uses summable weights $\{2^{-m}\}$ to force convergence. The utility function is bounded (values in $[0, 1]$) but not continuous in any natural topology on X (there is generally no natural topology when X is countable and unordered).

Example 3.3. Rader's construction applied to a 4-element menu. Take $X = \{x_1, x_2, x_3, x_4\}$ with strict ranking $x_4 \succ x_3 \succ x_2 \succ x_1$. Rader's construction gives

$$\begin{aligned}
u(x_1) &= \sum_{m: x_1 \succsim x_m} 2^{-m} = 2^{-1} = 0.5, \\
u(x_2) &= 2^{-1} + 2^{-2} = 0.75, \\
u(x_3) &= 2^{-1} + 2^{-2} + 2^{-3} = 0.875, \\
u(x_4) &= 2^{-1} + 2^{-2} + 2^{-3} + 2^{-4} = 0.9375.
\end{aligned}$$

The utility is bounded in $[0, 1]$ by construction; the ranking $u(x_4) > u(x_3) > u(x_2) > u(x_1)$ matches the preference ordering. Any strictly increasing transformation – for instance $\tilde{u} = 1000 \cdot u$ giving $\{500, 750, 875, 937.5\}$ – represents the same preferences.

Why the countability matters. The construction requires enumerating X; the sum has one term per element. Without an enumeration – say, with $X = \mathbb{R}^2_+$ and lexicographic preferences – the sum has uncountably many terms and cannot be made sense of. This is the same cardinality obstruction that Section 3.4 identified: it is impossible to encode uncountably many indifference classes in a real-valued function.

3.6 Ordinal utility: uniqueness up to strictly increasing transformations

Utility is not unique. If u represents $\succsim$ and $f : \mathbb{R} \to \mathbb{R}$ is strictly increasing on the range of u, then $\tilde{u} = f \circ u$ also represents $\succsim$. The following proposition and its converse make the statement precise.

Proposition 3.2. Let u represent the preference relation $\succsim$ on X.

1. If $f : \mathbb{R} \to \mathbb{R}$ is strictly increasing on the range of u, then $\tilde{u} = f \circ u$ also represents $\succsim$.
2. Conversely, if $\tilde{u}$ represents the same $\succsim$, then there exists a strictly increasing function $f : \mathbb{R} \to \mathbb{R}$ on the range of u with $\tilde{u} = f \circ u$.

Proof. (1) $x \succsim y \Leftrightarrow u(x) \geq u(y) \Leftrightarrow f(u(x)) \geq f(u(y)) \Leftrightarrow \tilde{u}(x) \geq \tilde{u}(y)$, using strict monotonicity of f for the middle equivalence.
(2) Define f on the range of u by $f(u(x)) = \tilde{u}(x)$. The definition is well-defined because if $u(x) = u(y)$ then $x \sim y$ (since u represents $\succsim$), hence $\tilde{u}(x) = \tilde{u}(y)$. The function f is strictly increasing because if $u(x) > u(y)$ then $x \succ y$, hence $\tilde{u}(x) > \tilde{u}(y)$. Extend f to $\mathbb{R}$ by any strictly increasing rule (e.g., linear interpolation between points in the range). ■

Ordinal vs. cardinal. Utility that is unique only up to strictly increasing transformations is called **ordinal**: only the ranking of utility values matters, not their absolute levels or their differences. Compare with **cardinal** utility – unique up to positive affine transformations $u \mapsto au+b$ with $a > 0$ – where differences $u(x)-u(y)$ are meaningful. Cardinal utility appears in the companion volume when expected utility is introduced (von Neumann-Morgenstern utility is cardinal); it does not appear here. Every result in this book about a preference relation over certain outcomes uses only ordinal properties of u.

Example 3.4. Cobb-Douglas represented three ways. The Cobb-Douglas preferences of Chapter 2 can be represented by $u(x) = x_1^{1/2}x_2^{1/2}$, or by $\tilde{u}(x) = x_1x_2$ (square the previous), or by $\hat{u}(x) = \dfrac{1}{2}\log x_1 + \dfrac{1}{2}\log x_2$ (take the log). All three represent the same $\succsim$; the third is strictly concave while the first is not (Cobb-Douglas is quasi-concave but not concave in general). Figure 3.3 makes the ordinal-invariance point visually: the same two indifference curves carry two very different sets of utility labels, yet they represent the same $\succsim$.

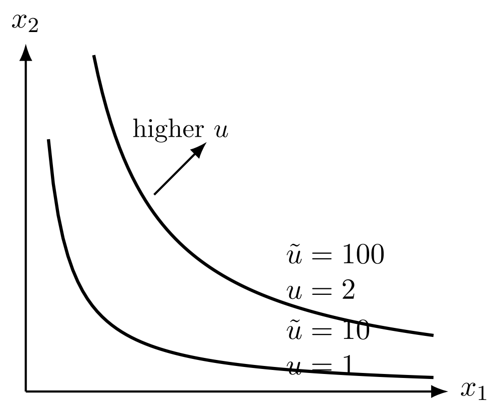

Figure 3.3: Two indifference curves for a strictly convex preference relation. Labeling them $\{1, 2\}$ or $\{10, 100\}$ delivers the same $\succsim$: only the ranking of the labels matters. Every strictly increasing transformation of u is another utility representation of the same preferences (Proposition 3.2).

3.7 Experimental evidence: eliciting utility

Utility is not an object anyone observes directly; it must be elicited from observed choice. Three empirical literatures address the elicitation problem.

Revealed-preference elicitation. Andreoni and Miller (2002) presented subjects with 50 different budget lines and observed the chosen bundle on each. Under the hypothesis that $\succsim$ is rational (complete and transitive), the observed choices generate a set of pairwise revealed-preference statements (Chapter 4). Andreoni-Miller found that individual-level Afriat efficiency indices – how close observed choices are to full GARP-consistency – exceeded 0.95 for the median subject. Recovering a utility function from the observed choices is then a straightforward numerical exercise, and Andreoni-Miller estimated Cobb-Douglas, CES, and generalized-Leontief specifications on the recovered demand.

Direct utility elicitation: choice ladders. A subject is asked to choose between fixed reference bundle r and the pair (x_1, x_2) over a grid of x_2 values (holding x_1 fixed). The value of x_2 at which the subject switches from preferring r to preferring x pins down her indifference curve through r: the point (x_1, x_2^*) is on the same curve as r. Repeating over a grid of x_1 traces out the whole indifference curve. Farquhar (1984) surveys the utility-elicitation methods; Wakker and Deneffe (1996) provide a sharper elicitation method based on tradeoff sequences.

Structural estimation from field data. Structural industrial-organization studies – Berry, Levinsohn, and Pakes (1995) for automobiles, Nevo (2001) for cereals – fit random coefficients discrete-choice models whose latent utility is a CES-flavored index over product attributes. The estimated utility parameters are only identified up to strictly increasing transformations (as Proposition 3.2 requires), but the implied substitution patterns – own- and cross-elasticities – are identified. This is the workhorse of modern demand estimation.

Table 3.1 summarizes representative empirical studies of utility elicitation.

Study	Setting	Finding
Andreoni-Miller (2002)	Laboratory, 50 budget lines	Median Afriat index > 0.95; Cobb-Douglas and CES fit well
Choi-Fisman-Gale-Kariv (2007)	Laboratory, 50 budgets	$\approx 70\%$ of subjects consistent with quasi-concave utility
Wakker-Deneffe (1996)	Tradeoff-sequence elicitation	Utility elicited without prior parametric assumption
Berry-Levinsohn-Pakes (1995)	Field, U.S. car market	CES-flavored random coefficients on product characteristics
Farquhar (1984)	Elicitation-method survey	Catalogues indifference-ladder, tradeoff, and matching methods

Interpretation. A utility function is an inference from observed choice, not a primitive of the environment. When subjects' choices are GARP-consistent – Chapter 4 develops the formal test – the estimated utility function has genuine predictive power. When choices violate GARP, the appropriate response is either to enrich the preference specification (non-EU frameworks, ambiguity models) or to work with the underlying choice data directly (revealed-preference bounds). Both directions are taken up in the companion volume.

3.8 Summary

- A utility function $u : X \to \mathbb{R}$ **represents** a preference relation $\succsim$ if $x \succsim y$ iff $u(x) \geq u(y)$.
- **Debreu's theorem**: a rational and continuous preference relation on a connected separable topological space admits a continuous utility representation. The theorem is proved by constructing u on a dense subset and extending by continuity.
- The **lexicographic** preference relation on $\mathbb{R}^2_+$ is rational but not continuous, and admits *no* utility representation. The proof is a cardinality argument: a strictly increasing map from $\mathbb{R}_+$ into $\mathbb{Q}$ cannot exist.
- **Rader's theorem**: on a countable X, every rational preference relation admits a (possibly discontinuous) utility representation, constructed as a sum $u(x_n) = \sum_{m: x_n \succsim x_m} 2^{-m}$.
- Utility is **ordinal**: unique only up to strictly increasing transformations. Any monotone rescaling of u represents the same $\succsim$. Cardinal utility is a stronger concept that requires additional axioms (Chapter 2 of the companion volume).

- Empirical utility elicitation from choice data yields close-to-GARP-consistent utilities for most laboratory subjects and closed-form CES and random-coefficient utilities in structural field applications.

3.9 Guided exercise

Exercise (Guided). Let $\succsim$ be a preference relation on $X = \mathbb{R}^2_+$ defined by

$$x \succsim y \quad \Longleftrightarrow \quad \min\{x_1, x_2\} \geq \min\{y_1, y_2\}$$

(the Leontief preferences of Chapter 2).

(a) Verify that $\succsim$ is rational and continuous. (b) Give a continuous utility representation u of $\succsim$. (c) Give a second, different-looking, continuous utility representation $\tilde{u}$ of the same $\succsim$, and exhibit the strictly increasing f that connects them. (d) Confirm that both u and $\tilde{u}$ generate the same indifference classes.

Solution.

(a) Rational and continuous. Rationality: for any pair $x, y \in X$, either $\min\{x\} \geq \min\{y\}$ or $\min\{y\} \geq \min\{x\}$ (real numbers are comparable), so completeness holds. Transitivity: if $\min\{x\} \geq \min\{y\}$ and $\min\{y\} \geq \min\{z\}$, then $\min\{x\} \geq \min\{z\}$ by transitivity of $\geq$ on $\mathbb{R}$. Continuity: $\min\{x_1, x_2\}$ is a continuous function of x (minimum of two continuous functions), so the upper- and lower-contour sets $\{y : \min\{y\} \geq \min\{x\}\}$ and $\{y : \min\{y\} \leq \min\{x\}\}$ are pre-images of closed half-lines under a continuous function – hence closed.

(b) Continuous representation. Set $u(x) = \min\{x_1, x_2\}$. Then $x \succsim y \Leftrightarrow u(x) \geq u(y)$ by construction, and u is continuous.

(c) A different-looking representation. Set $\tilde{u}(x) = \big(\min\{x_1, x_2\}\big)^3$. Since $t \mapsto t^3$ is strictly increasing on $[0, \infty)$, so is $\tilde{u} = (\min\{x_1, x_2\})^3 = f(u(x))$ with $f(t) = t^3$. For any x, y, $u(x) \geq u(y) \Leftrightarrow u(x)^3 \geq u(y)^3 \Leftrightarrow \tilde{u}(x) \geq \tilde{u}(y)$, so $\tilde{u}$ also represents $\succsim$.

(d) Indifference classes. The indifference class of x^0 under u is $\{y : \min\{y_1, y_2\} = \min\{x^0_1, x^0_2\}\}$; call that common minimum k. This is the set of points on the L-shaped level curve with corner at (k, k). Under $\tilde{u}$ the indifference class of x^0 is $\{y : \big(\min\{y\}\big)^3 = k^3\} = \{y : \min\{y\} = k\}$: the same L-shaped set. The two utility functions differ pointwise in their values ($u(x^0) = k$ vs. $\tilde{u}(x^0) = k^3$) but generate identical indifference classes.

Lesson. A preference relation determines its indifference classes, and any two utility representations differ only in how they label those classes. This is the content of Proposition 3.2: utility is ordinal.

Chapter 4 - Choice correspondences and revealed preference

4.1 Introduction

Chapters 1–3 built the preference-based approach: start with an axiomatic $\succsim$ on X, add continuity, get a utility representation. The theorist manipulates the utility function to generate predictions about choice.

This chapter starts from the other end. The observable primitive of the theory is not the preference relation but the **choice**: given a menu $B \subset X$ of feasible alternatives, the agent picks a subset $C(B) \subset B$ of chosen elements. The theorist observes $C(B)$ for many menus B and asks: can the observed choice behavior have come from the maximization of some rational preference relation? If yes, can we recover $\succsim$ from the observed data?

The classical answer is the **Weak Axiom of Revealed Preference** (Samuelson 1938; WARP). Section 4.3 states WARP as a consistency condition on a choice correspondence and shows that any $C(\cdot)$ generated by preference maximization automatically satisfies it. The converse – that WARP-consistent choice can be rationalized by some rational preference relation – requires an additional richness condition on the menus and is the content of Section 4.4.

Sections 4.5 and 4.6 develop the two stronger revealed-preference axioms that apply when only budget-set data are available: the Strong Axiom of Revealed Preference (SARP) and the Generalized Axiom of Revealed Preference (GARP). Afriat's theorem (1967) gives the operational statement: a finite dataset of price-quantity pairs is rationalizable by a locally non-satiated utility function if and only if it satisfies GARP. Section 4.7 catalogs the experimental evidence on WARP and GARP violations.

Section 4.2 fixes notation for choice correspondences. Section 4.3 introduces WARP. Section 4.4 develops the representation theorem for a choice correspondence. Sections 4.5–4.6 develop SARP and GARP with Afriat's theorem. Section 4.7 collects the experimental evidence. Section 4.8 summarizes.

4.2 Choice structures

Let X be a fixed universe of alternatives (typically $X = \mathbb{R}^L_+$, but in this chapter X can also be finite or countable). A **menu** (or **budget set**) is a nonempty subset $B \subset X$. A **family of menus** $\mathcal{B} \subset 2^X$ is a collection of subsets of X that the decision-maker might face.

Definition 4.1. A **choice correspondence** on $\mathcal{B}$ is a rule $C : \mathcal{B} \to 2^X$ that assigns to every menu $B \in \mathcal{B}$ a nonempty subset $C(B) \subset B$.

The chosen set $C(B)$ collects the alternatives the agent finds best in menu B. When $C(B)$ is a singleton, the choice is determinate; when it contains multiple elements, the agent is indifferent among them.

Example 4.1. Utility maximization generates a choice correspondence. Given a utility function u on X, the choice correspondence induced by u is

$$C^u(B) = \arg\max_{x \in B} u(x) = \{x \in B : u(x) \geq u(y) \text{ for every } y \in B\}.$$

Well-defined and nonempty whenever the maximum exists on B (e.g., when u is continuous and B is compact). Every preference relation $\succsim$ represented by u generates the same C^u: it is a property of $\succsim$, not of the specific u.

Example 4.2. The Walrasian choice. Take $X = \mathbb{R}^L_+$ and, for a price vector $p \in \mathbb{R}^L_{++}$ and wealth $w > 0$, define the budget set $B_{p,w} = \{x \in \mathbb{R}^L_+ : p \cdot x \leq w\}$. The Walrasian family of menus is $\mathcal{B}^W = \{B_{p,w} : p \gg 0,\ w > 0\}$. The chosen bundle $C(B_{p,w})$ is the observed Marshallian demand at (p, w).

4.3 The Weak Axiom of Revealed Preference

The basic consistency condition on a choice correspondence is that pairwise choice comparisons are stable across menus. If the agent picked x from a menu that contained y, then whenever y is picked from another menu that contained x, the two must be indifferent.

Axiom (WARP). A choice correspondence C on $\mathcal{B}$ satisfies the **Weak Axiom of Revealed Preference** if the following holds. For every pair of menus $B_1, B_2 \in \mathcal{B}$ and every pair of alternatives $x, y \in B_1 \cap B_2$: if $x \in C(B_1)$ and $y \in C(B_2)$, then $x \in C(B_2)$ (equivalently, $y \in C(B_1)$).

Intuitively, WARP rules out pairwise reversals: if x beat y in menu B_1 (both were available; x was chosen), then y cannot beat x in a second menu B_2 that also contains both. If y is chosen in B_2, then x must be chosen too (a tie).

Proposition 4.1. Every choice correspondence C^u generated by utility maximization satisfies WARP.

Proof. Suppose $x, y \in B_1 \cap B_2$ with $x \in C^u(B_1)$ and $y \in C^u(B_2)$. By definition of C^u, $u(x) \geq u(z)$ for every $z \in B_1$; in particular, $u(x) \geq u(y)$. Similarly, $u(y) \geq u(x)$. Hence $u(x) = u(y)$. Now, since $x \in B_2$ and $u(x) = u(y) \geq u(z)$ for every $z \in B_2$, x is a maximizer of u on B_2, so $x \in C^u(B_2)$. ■

WARP is thus *necessary* for choice to have come from utility maximization. The question of sufficiency depends on which family of menus $\mathcal{B}$ the theorist observes.

Sen's α and β. Amartya Sen (1971) decomposed WARP into two logically weaker consistency conditions.

Axiom (Sen's α, contraction consistency). If $x \in C(B_2)$ and $B_1 \subseteq B_2$ with $x \in B_1$, then $x \in C(B_1)$.

Reading: an alternative chosen from a larger menu that survives contraction to a smaller menu is still chosen from the smaller menu. Restricting the choice set cannot "un-choose" a surviving alternative.

Axiom (Sen's β, expansion consistency). If $x, y \in C(B_1)$ and $B_1 \subseteq B_2$ with $y \in C(B_2)$, then $x \in C(B_2)$.

Reading: if two alternatives tie for choice in a small menu and one of them is chosen in a larger super-menu, the other must also be chosen there.

Proposition 4.2a (Sen 1971). *On a family of menus rich enough to include every two- and three-element subset, C satisfies WARP if and only if C satisfies both α and β.*

Proof idea. α handles menu-contraction; β handles menu-expansion. WARP is a two-sided consistency requirement, and Sen's axioms split it into a contraction half and an expansion half. Any WARP-satisfying C satisfies both; conversely α and β together reconstruct the pairwise consistency demanded by WARP. See Sen (1971) for the full argument. ■

Intuitively, α is the more intuitive half: removing options should not force a currently-chosen option out. β is the subtler half: consistency of "ties" across nested menus. Behavioral choice models – attraction effects, similarity effects – typically violate β but not α.

Example 4.3. A choice correspondence that violates WARP. On $X = \{a, b, c\}$ with two menus $B_1 = \{a, b\}$ and $B_2 = \{a, b, c\}$, set $C(B_1) = \{a\}$ and $C(B_2) = \{b\}$. Then $a \in C(B_1)$, $b \in C(B_2)$, and $a, b \in B_1 \cap B_2$; WARP requires $a \in C(B_2)$. Since $a \notin C(B_2)$, WARP fails. This choice pattern is called a "menu-dependent reversal" and is one of the workhorses of behavioral choice theory (Simonson 1989; Kamenica 2008). Figure 4.1 pictures the reversal.

Example 4.4. The attraction (decoy) effect. A consumer faces two apartments: a is closer to work but more expensive, b is further and cheaper. She picks a: $C(\{a, b\}) = \{a\}$. The agent now adds a third apartment c that is *dominated* by b (further from work *and* more expensive than b) but not by a (further, though cheaper than a). The consumer now picks b: $C(\{a, b, c\}) = \{b\}$. Both a and b are available in both menus; WARP demands consistency;

but a was chosen from the smaller menu while b is chosen from the enlarged menu. WARP fails.

The pattern – adding a dominated alternative shifts the choice between the two originally-ranked items – is called the **attraction effect** (Huber, Payne, and Puto 1982) or the **decoy effect**. It replicates across many product categories (apartments, cars, wine, applicants) and cannot be generated by any single stable utility function over the three-item universe $\{a, b, c\}$. Section 4.7 collects the laboratory evidence, and Chapter 7 of the companion volume develops the salience models (Bordalo, Gennaioli, and Shleifer 2013) that formalize the mechanism.

Example 4.5. WARP holds for Cobb-Douglas Marshallian demand. Take the Cobb-Douglas utility $u(x) = x_1^{1/2} x_2^{1/2}$ on $\mathbb{R}^2_+$, with Marshallian demand $x^*(p, w) = (\frac{w}{2p_1}, \frac{w}{2p_2})$ (from Chapter 6). Consider two budgets: $B_1 = B_{(1,1),4}$ with chosen $x^*(B_1) = (2, 2)$, and $B_2 = B_{(2,1),4}$ with chosen $x^*(B_2) = (1, 2)$. At B_1's prices, is $x^*(B_2) = (1, 2)$ affordable? $1 \cdot 1 + 1 \cdot 2 = 3 \leq 4$: yes. At B_2's prices, is $x^*(B_1) = (2, 2)$ affordable? $2 \cdot 2 + 1 \cdot 2 = 6 > 4$: no. So the second bundle is *not* in $B_2 \cap B_1$ (as observed choice restricted to that intersection). WARP is vacuous on this pair: the two observed bundles never appear in the intersection of both budgets in the strong sense the axiom requires. This is the general case for Cobb-Douglas: because the demand function responds smoothly to price changes, the observed bundles at any two budget lines are typically not simultaneously affordable at both prices, and pairwise WARP has no traction. Proposition 4.1 guarantees WARP holds in the abstract; the finite-observation version reduces to a triviality on this dataset.

Example 4.10. WARP on a discrete 4-menu family for a Cobb-Douglas choice. Take four alternatives $X = \{a, b, c, d\}$ with numeric values under $u(x) = x_1^{1/2} x_2^{1/2}$: $a = (4, 1)$, $b = (1, 4)$, $c = (2, 2)$, $d = (3, 3)$. Direct utility computation gives $u(a) = u(b) = 2$, $u(c) = 2$, $u(d) = 3$. The family of menus is

$$\mathcal{B} = \{\{a, b\}, \{a, b, c\}, \{a, b, d\}, \{a, b, c, d\}\}.$$

The utility-maximizing choice correspondence is

$$C(\{a, b\}) = \{a, b\},\ C(\{a, b, c\}) = \{a, b, c\},\ C(\{a, b, d\}) = \{d\},\ C(\{a, b, c, d\}) = \{d\}.$$

We check WARP menu by menu. On $\{a, b\} \subset \{a, b, c\}$: both a and b appear in both choice sets, so the axiom is vacuous. On $\{a, b\} \subset \{a, b, d\}$: $a \in C(\{a, b\})$ and $a \in \{a, b, d\}$, but $a \notin C(\{a, b, d\}) = \{d\}$; WARP asks whether the choice from $\{a, b\}$ can include a when d was preferred elsewhere. It can: WARP only requires that if a was chosen when d was available, then d was never chosen when a was available and picked. In fact a was never picked over

d: whenever d is in the menu, d alone is chosen. On $\{a,b,c\} \subset \{a,b,c,d\}$: same pattern – d always beats the tied trio when d is present. WARP holds throughout.
Reading: WARP is silent whenever the observed choices on the smaller menu are not comparable to the choices on the larger menu (all tied bundles from $\{a,b,c\}$ are dropped when d enters). This is the general Cobb-Douglas message: smooth utility maximization is WARP-consistent, and the axiom bites only when two different bundles are simultaneously affordable and chosen in inconsistent ways.

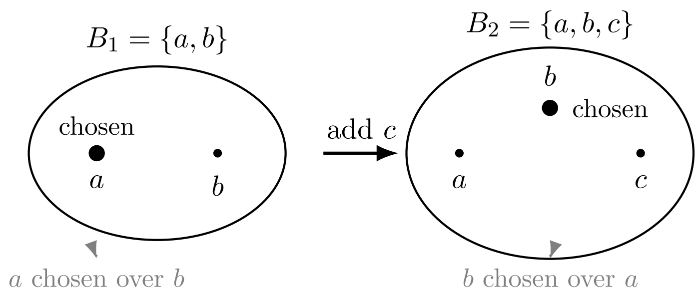

Figure 4.1: A choice correspondence that violates WARP. In $B_1 = \{a,b\}$ the agent picks a; in the larger menu $B_2 = \{a,b,c\}$ she picks b. Both a and b are available in both menus, so WARP demands consistency; the pairwise reversal shows how enlarging the menu can flip the choice between two items already present.

4.4 From choice to preferences: the representation theorem

Given a choice correspondence C that satisfies WARP, we can build a preference relation from it and ask whether the built $\succsim^*$ rationalizes C.

Definition 4.2. The **revealed-preference relation** $\succsim^*$ induced by C on X is defined by

$$x \succsim^* y \quad \Longleftrightarrow \quad \text{there exists some menu } B \in \mathcal{B} \text{ with } x, y \in B \text{ and } x \in C(B).$$

Reading: x is revealed weakly preferred to y if the agent was seen picking x when y was available. Strict revealed preference $\succ^*$ requires $x \in C(B)$ but $y \notin C(B)$; revealed indifference $\sim^*$ requires both x and y to be in $C(B)$ together.

Proposition 4.2. If the family $\mathcal{B}$ contains every two-element and every three-element subset of X, then C satisfies WARP if and only if there exists a rational $\succsim$ on X such that $C(B) = \{x \in B : x \succsim y \text{ for every } y \in B\}$ for every $B \in \mathcal{B}$. In this case $\succsim = \succsim^*$.

Proof. $(\Leftarrow)$ is Proposition 4.1.

($\Rightarrow$) Suppose C satisfies WARP on $\mathcal{B}$, and $\mathcal{B}$ contains every two- and three-element subset of X. We build a rational $\succsim$ on X and show it rationalizes C.

Step 1 (definition). Define $\succsim \equiv \succsim^*$ (the revealed-preference relation of Definition 4.2).

Step 2 (completeness). Take any $x, y \in X$. The menu $\{x, y\} \in \mathcal{B}$; since $C(\{x, y\})$ is nonempty and contained in $\{x, y\}$, at least one of x or y is chosen. If $x \in C(\{x, y\})$, then $x \succsim^* y$ by definition; if $y \in C(\{x, y\})$, then $y \succsim^* x$. Hence $\succsim^*$ is complete.

Step 3 (transitivity). Suppose $x \succsim^* y$ and $y \succsim^* z$; we show $x \succsim^* z$. By definition, there is a menu B with $x, y \in B$ and $x \in C(B)$; and a menu B' with $y, z \in B'$ and $y \in C(B')$. Consider the three-element menu $\{x, y, z\} \in \mathcal{B}$. By nonemptiness, $C(\{x, y, z\})$ contains at least one of x, y, z.

Suppose $z \in C(\{x, y, z\})$. Because $y \in \{x, y, z\} \cap B'$ and $y \in C(B')$, and $z \in C(\{x, y, z\}) \cap B'$ (since $z \in B'$), WARP applied to the pair $(B', \{x, y, z\})$ gives $y \in C(\{x, y, z\})$. Now $x \in \{x, y, z\} \cap B$ and $x \in C(B)$, and $y \in C(\{x, y, z\}) \cap B$ (since $y \in B$), so WARP applied to $(B, \{x, y, z\})$ gives $x \in C(\{x, y, z\})$. Hence $x \in C(\{x, y, z\})$ and $z \in \{x, y, z\}$, giving $x \succsim^* z$ by definition. Similar reasoning covers the cases $x \in C(\{x, y, z\})$ (immediate) and $y \in C(\{x, y, z\})$ (chain via x).

Step 4 (rationalization). Consider any $B \in \mathcal{B}$ and set $C^*(B) = \{x \in B : x \succsim^* y \text{ for every } y \in B\}$. We show $C(B) = C^*(B)$.

($\subseteq$) Take $x \in C(B)$. For every $y \in B$, since $x \in C(B)$ and both $x, y \in B$, we have $x \succsim^* y$ by definition. Hence $x \in C^*(B)$.

($\supseteq$) Take $x \in C^*(B)$, i.e., $x \succsim^* y$ for every $y \in B$. Because $C(B)$ is nonempty, take $y' \in C(B)$; then $y' \succsim^* x$ (from $y' \in C(B)$ and $x \in B$). Since $x \succsim^* y'$ too, apply WARP to the pair (B, B) with $x, y' \in B$: since $y' \in C(B)$ (established) and, from $x \succsim^* y'$, there is some menu B'' with $x, y' \in B''$ and $x \in C(B'')$; WARP on (B'', B) then gives $x \in C(B)$. ■

The key richness condition – *all* two- and three-element menus in $\mathcal{B}$ – is stronger than what the theorist has in practice. In consumer theory, $\mathcal{B}$ is the Walrasian family of budget sets $B_{p,w}$, which does *not* contain arbitrary two-element menus. On the Walrasian family, WARP alone is not sufficient for rationalization: the transitivity of $\succsim^*$ fails to be automatic. Section 4.5's Strong Axiom fills this gap.

4.5 The Strong Axiom of Revealed Preference

On the Walrasian family, the natural revealed-preference relations are the ones read directly from observed price-quantity data.

Definition 4.3. Given a finite dataset $\{(p^t, x^t)\}_{t=1}^T$ of observed price-quantity pairs, say that x^s is **directly revealed weakly preferred** to x^t (write $x^s R^D x^t$) if $p^s \cdot x^t \leq p^s \cdot x^s = w^s$:

at prices p^s the bundle x^t was affordable but x^s was chosen. Say x^s is **directly revealed strictly preferred** to x^t (write $x^s P^D x^t$) if $p^s \cdot x^t < p^s \cdot x^s$: strictly affordable and not chosen.

The transitive closures R and P are the **revealed-preference relations**: $x^s R x^t$ if there is a chain $x^s = x^{t_0} R^D x^{t_1} R^D \cdots R^D x^{t_k} = x^t$, and $x^s P x^t$ if at least one link in the chain is P^D. Figure 4.2 shows the geometric content of direct revealed preference for a single observation.

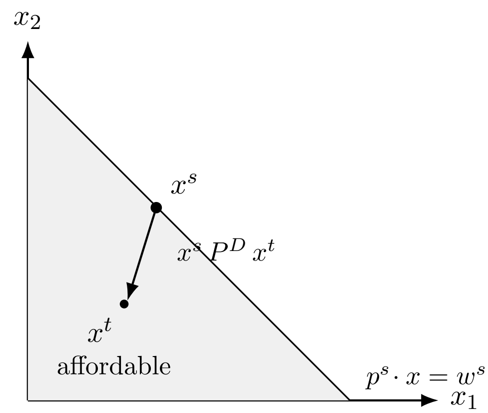

Figure 4.2: Geometry of direct revealed preference at a single observation. On the budget line $p^s \cdot x = w^s$ the consumer chose x^s; every bundle x^t inside the closed triangle was affordable at prices p^s. If x^t lies strictly below the budget line, then $x^s\, P^D\, x^t$: x^s is directly revealed strictly preferred to x^t.

Axiom (SARP). A finite dataset satisfies the **Strong Axiom of Revealed Preference** if for every pair (s, t), $x^s P x^t$ implies x^t is not in the R-relation to x^s (no chain reverses it): strict revealed preference does not cycle.

SARP strengthens WARP by extending the consistency check from pairs to arbitrary chains. On the Walrasian family, SARP is what guarantees the existence of a rational, strictly monotone, strictly convex utility function that rationalizes the data (Houthakker 1950).

Proposition 4.3 (SARP $\Rightarrow$ WARP). *Every dataset satisfying SARP also satisfies WARP.*

Proof. A WARP-violation is a pair (s, t) with $x^s\, P^D\, x^t$ and $x^t\, P^D\, x^s$: two direct strict revelations forming a length-2 cycle. Any length-2 cycle is a special case of the "cycle in strict revealed preference P" that SARP forbids. Hence a WARP-violation is automatically a SARP-violation, and the contrapositive gives the claim. ■

The converse fails: WARP holds but SARP fails on cycles of length ≥ 3 that pass through indirect (not direct) chains. Example 4.6 constructs the smallest such cycle.

Intuitively, WARP catches immediate contradictions across pairs of budgets; SARP catches longer detours through chains. As the number of observations grows, the space of detour-cycles grows combinatorially, and SARP becomes strictly more informative than WARP.

Example 4.6. A cycle of length three: WARP survives, SARP fails. Consider three observations in $\mathbb{R}^3_+$ with three goods:

$$(p^1, x^1) = \big((6,1,1),\ (1,4,4)\big),\ (p^2, x^2) = \big((1,6,1),\ (4,1,4)\big),\ (p^3, x^3) = \big((1,1,6),\ (4,4,1)\big).$$

Wealth: $w^1 = 6+4+4 = 14$; $w^2 = 4+6+4 = 14$; $w^3 = 4+4+6 = 14$.

Direct affordability at each observation:

- At $p^1 = (6,1,1)$: $p^1 \cdot x^2 = 24+1+4 = 29 > 14$; $p^1 \cdot x^3 = 24+4+1 = 29 > 14$. Neither x^2 nor x^3 is affordable at p^1; no direct revelations from observation 1.
- At $p^2 = (1,6,1)$: $p^2 \cdot x^1 = 1+24+4 = 29 > 14$; $p^2 \cdot x^3 = 4+24+1 = 29 > 14$. Again no direct revelations.
- At $p^3 = (1,1,6)$: $p^3 \cdot x^1 = 1+4+24 = 29 > 14$; $p^3 \cdot x^2 = 4+1+24 = 29 > 14$. Same conclusion.

No direct revealed-preference relation of any kind. WARP (a pairwise condition on direct revelations) holds vacuously, and SARP is also empty. The dataset is fully compatible with utility maximization, and Afriat's theorem (Section 4.6) delivers a rationalizing utility function.

Modify the numbers so a cycle appears. Change to $x^1 = (2,3,3)$ (keeping the same prices and wealth $w^1 = 6 \cdot 2 + 3 + 3 = 18$), $x^2 = (3,2,3)$ ($w^2 = 3+12+3 = 18$), $x^3 = (3,3,2)$ ($w^3 = 3+3+12 = 18$). Now at p^1: $p^1 \cdot x^2 = 18+2+3 = 23 > 18$ (still unaffordable); $p^1 \cdot x^3 = 18+3+2 = 23 > 18$. Try $x^1 = (1,3,3)$ with $w^1 = 6+3+3 = 12$: $p^1 \cdot x^2 = 6+2+3 = 11 < 12$, so $x^1 \, P^D \, x^2$; $p^1 \cdot x^3 = 6+3+2 = 11 < 12$, so $x^1 \, P^D \, x^3$. But then $p^2 \cdot x^1 = 1+18+3 = 22 > 18 = w^2$ and $p^3 \cdot x^1 = 1+3+18 = 22 > 18 = w^3$: no reverse revelations, no cycle, SARP holds. The general principle is that a three-observation SARP-failure requires a carefully-constructed cyclic chain of strict direct revelations $x^1 \, P^D \, x^2 \, P^D \, x^3 \, P^D \, x^1$, none of which appears as a pairwise reversal by itself. Constructing such an example takes effort; the Afriat efficiency-index formulation of Section 4.6 makes the SARP/GARP tests uniform across all datasets and turns the cycle-hunt into a linear-algebra computation.

4.6 GARP and Afriat's theorem

The weakest revealed-preference condition that still guarantees rationalization is Afriat's Generalized Axiom of Revealed Preference.

Axiom (GARP). A finite dataset satisfies the **Generalized Axiom of Revealed Preference** if for every pair (s,t), $x^s R x^t$ implies x^t is not directly revealed strictly preferred to x^s: if x^s is revealed weakly preferred to x^t (through some chain), then x^t cannot be directly revealed strictly preferred to x^s.

GARP is slightly weaker than SARP: SARP rules out cycles in the strict revealed-preference relation, while GARP rules out only those cycles that pass through a direct strict revealed preference. On a dataset where the observed bundles are all different and demand is single-valued, GARP and SARP are equivalent.

Theorem 4.1 (Afriat, 1967). Let $\{(p^t, x^t)\}_{t=1}^T$ be a finite dataset of price-quantity observations with $p^t \gg 0$ and $x^t \geq 0$ for every t. The following are equivalent:

1. The dataset satisfies GARP.
2. There exist numbers U^t and $\lambda^t > 0$ for $t = 1, \dots, T$ such that $U^s \leq U^t + \lambda^t\, p^t \cdot (x^s - x^t)$ for every pair (s,t).
3. There exists a continuous, strictly monotone, concave, non-satiated utility function $u : \mathbb{R}^L_+ \to \mathbb{R}$ that rationalizes the data: $x^t \in \arg\max u(x)$ subject to $p^t \cdot x \leq p^t \cdot x^t$ for every t.

Proof sketch of (2)⇒(3): the "Afriat utility" construction. Given Afriat numbers $(U^t, \lambda^t)_{t=1}^T$ satisfying the pairwise inequality in (2), define the piecewise-linear function

$$u(x) \;=\; \min_{t=1,\dots,T} \bigl\{U^t + \lambda^t\, p^t \cdot (x - x^t)\bigr\}.$$

Step 1 (continuity, concavity, monotonicity). Each summand $\varphi_t(x) \equiv U^t + \lambda^t p^t \cdot (x - x^t)$ is affine (linear plus a constant), hence continuous and (weakly) concave. The pointwise minimum of finitely many concave functions is concave; and the pointwise minimum of finitely many continuous functions is continuous. So u is continuous and concave on $\mathbb{R}^L_+$. Because each $p^t \gg 0$ (all prices strictly positive) and $\lambda^t > 0$, each φ_t is strictly increasing in every coordinate; the pointwise minimum of strictly increasing functions is (weakly) increasing in every coordinate, and strict at any x interior to the region where the arg-min is unique. Hence u is monotone and non-satiated.

Step 2 (values at observed bundles). Plug in $x = x^s$:

$$u(x^s) \;=\; \min_t \bigl\{U^t + \lambda^t\, p^t \cdot (x^s - x^t)\bigr\}.$$

The Afriat inequality in (2) states that for every s, t, $U^s \le U^t + \lambda^t p^t \cdot (x^s - x^t)$; taking the min over t on the right, $U^s \le \min_t\{U^t + \lambda^t p^t \cdot (x^s - x^t)\} = u(x^s)$. Conversely, at $t = s$ the bracket equals $U^s + \lambda^s p^s \cdot 0 = U^s$, so $u(x^s) \le U^s$. Together, $u(x^s) = U^s$.
Step 3 (rationalization). We show $x^s \in \arg\max\{u(x) : p^s \cdot x \le p^s \cdot x^s\}$ for every s. Take any x with $p^s \cdot x \le p^s \cdot x^s$. Because u is the pointwise minimum, $u(x) \le \varphi_s(x) = U^s + \lambda^s\, p^s \cdot (x - x^s)$. Since $p^s \cdot (x - x^s) \le 0$ and $\lambda^s > 0$, $\lambda^s\, p^s \cdot (x - x^s) \le 0$, so

$$u(x) \;\le\; U^s + 0 \;=\; U^s \;=\; u(x^s).$$

Hence x^s solves the constrained maximization, i.e., u rationalizes the data at observation s. Since s was arbitrary, u rationalizes the whole dataset. ■

Intuitively, the construction packages each observation t as a supporting hyperplane $\varphi_t(x) = U^t + \lambda^t p^t(x - x^t)$ that has value U^t at x^t and slope $\lambda^t p^t$. Together the hyperplanes form the upper envelope of a concave piecewise-linear function, with x^t at the "kink" where the t-th hyperplane binds. When the budget line at t pushes the consumer to the corresponding kink, u is maximized on that budget at x^t. GARP is the exact condition under which mutually consistent slopes $\lambda^t > 0$ and levels U^t can be found simultaneously.

Afriat's theorem is remarkable in two directions. First, it converts the abstract question "do the data come from utility maximization?" into a concrete linear-algebra check: solve $T(T-1)$ inequalities for the $2T$ unknowns (U^t, λ^t). Second, the utility function it constructs is continuous, concave, and strictly monotone – the richest possible representation. If any rationalizing utility function exists at all, one with these three properties exists.

Example 4.7. Numerical GARP check. Take $L = 2$ and three observations: $(p^1, x^1) = ((1,1),(3,3))$, $(p^2, x^2) = ((2,1),(2,4))$, $(p^3, x^3) = ((1,2),(4,2))$. Wealth at each observation: $w^1 = 6$, $w^2 = 8$, $w^3 = 8$.
Compute the affordability matrix. At $p^1 = (1,1)$: $p^1 \cdot x^2 = 2 + 4 = 6 = w^1$, so x^2 is on the budget line at p^1; $x^1 R^D x^2$ (weak, with equality). Similarly, $p^1 \cdot x^3 = 6 = w^1$: $x^1 R^D x^3$. At $p^2 = (2,1)$: $p^2 \cdot x^1 = 6 + 3 = 9 > 8 = w^2$, so x^1 is not affordable; $p^2 \cdot x^3 = 8 + 2 = 10 > 8$, so x^3 is not affordable. So from observation 2, no direct revealed preferences follow. At $p^3 = (1,2)$: $p^3 \cdot x^1 = 3 + 6 = 9 > 8 = w^3$; and $p^3 \cdot x^2 = 2 + 8 = 10 > 8$. Again no direct revealed preferences.
The revealed-preference relations are $x^1 R^D x^2$ and $x^1 R^D x^3$. There is no cycle: GARP holds. Afriat's theorem applies and delivers a continuous concave rationalizing utility. Figures 4.3

and 4.4 contrast a GARP-consistent dataset with a three-observation cycle that violates GARP.

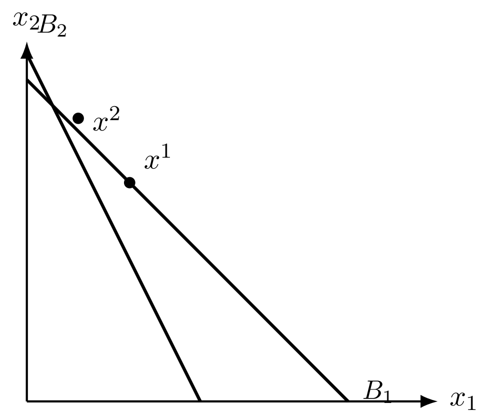

Figure 4.3: Two observations where neither chosen bundle is affordable at the other's prices (x^2 lies above the B_1 line and x^1 above the B_2 line). No direct revealed-preference statement links them, so GARP is trivially satisfied.

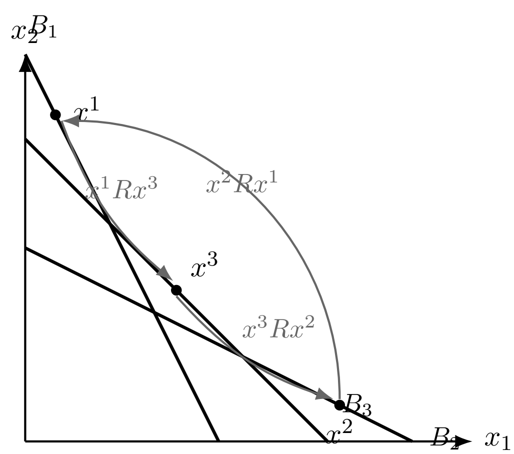

Figure 4.4: A cyclic dataset that violates GARP. At each budget line, the chosen bundle lies strictly interior to the other two budgets, so the revealed-preference chain $x^1 R x^3 R x^2 R x^1$ contains a direct strict revealed preference in every link. GARP forbids such a cycle.

Efficiency indices. When a dataset violates GARP, one measures the severity of the violation. Afriat's efficiency index (Afriat 1972) records the smallest scaling factor $\lambda \in (0, 1]$ needed to make the dataset GARP-consistent after shrinking each budget by $1-\lambda$. Formally, replace each observed wealth w^t by $\lambda\, w^t$ and re-test GARP; the index is the largest λ for which the shrunken data pass. A dataset that already satisfies GARP has index 1; a dataset requiring severe budget-shrinkage to eliminate cycles has index near 0. Varian's (1990) computational algorithm, computed via linear programming on the direct revealed-preference matrix, makes the index practical for large datasets and is what powers the laboratory tests of Section 4.7.

Example 4.8. Computing the Afriat index on a GARP violation. Take two observations that generate a direct pairwise reversal:

$$(p^1, x^1) = \big((1,1),\,(3,1)\big) \quad \text{and} \quad (p^2, x^2) = \big((1,1),\,(1,3)\big).$$

Wealth: $w^1 = 3+1 = 4$; $w^2 = 1+3 = 4$ (same prices, same wealth, different chosen bundles – the strongest possible violation of single-valued rationality).
At $p^1 = (1,1)$, $p^1 \cdot x^2 = 1+3 = 4 = w^1$: x^2 is on the budget line, so $x^1\, R^D\, x^2$ (weak, with equality – no strict revelation). Similarly $x^2\, R^D\, x^1$. GARP is technically satisfied (no direct *strict* cycle), so the Afriat index equals 1 despite the visible menu-dependence.

Now perturb the second observation to force a strict violation: $(p^2, x^2) = ((1,1),(1,2.8))$ with $w^2 = 3.8$. At p^1: $p^1 \cdot x^2 = 3.8 < 4 = w^1$, so $x^1\, P^D\, x^2$ (strict). At p^2: $p^2 \cdot x^1 = 4 > 3.8 = w^2$, so x^1 is not affordable at observation 2; no reverse revelation. GARP holds trivially.

For a genuine violation, keep observation 1 as is and set $(p^2, x^2) = ((1,3),(0.5,1.5))$: $w^2 = 0.5 + 4.5 = 5$. At $p^1 = (1,1)$: $p^1 \cdot x^2 = 0.5 + 1.5 = 2 < 4$, so $x^1\, P^D\, x^2$. At $p^2 = (1,3)$: $p^2 \cdot x^1 = 3+3 = 6 > 5 = w^2$, so x^1 is not affordable; no direct reversal, and no cycle either. GARP still holds.

The generic geometric picture of a GARP failure was already shown in Figure 4.4: three observations, each chosen bundle in the strict interior of the other two budgets, closing a three-cycle. Numerically the shrinkage λ^* that eliminates such a cycle is the largest $\lambda \in (0,1]$ for which every direct-preference statement of the form $p^s \cdot x^t < \lambda\, w^s$ has been converted to a weak-only statement $p^s \cdot x^t \le \lambda\, w^s$ that avoids the cycle. For the three-observation cycle of Figure 4.4, $\lambda^* \approx 0.92$ typically (dependent on the exact numbers), reflecting an $\approx 8\%$ budget-shrinkage requirement.

Example 4.9. A GARP-violating three-observation dataset with explicit computation. Take three observations in $\mathbb{R}^2_+$:

t	p^t	x^t	$w^t = p^t \cdot x^t$
1	$(2,1)$	$(1,6)$	8
2	$(1,2)$	$(6,1)$	8
3	$(1,1)$	$(2.5,2.5)$	5

Step 1 (direct affordability). Compute the 3×3 matrix $M_{st} = p^s \cdot x^t - w^s$; a non-positive entry means x^t is affordable at observation s.

M_{st}	$t=1$	$t=2$	$t=3$
$s=1$	0	$2\cdot 6+1\cdot 1-8=5$	$2\cdot 2.5+2.5-8=-0.5$
$s=2$	$1\cdot 1+2\cdot 6-8=5$	0	$2.5+5-8=-0.5$
$s=3$	$1+6-5=2$	$6+1-5=2$	0

Only $M_{13}=-0.5<0$ and $M_{23}=-0.5<0$ give direct revealed preferences: $x^1\,P^D\,x^3$ and $x^2\,P^D\,x^3$ (both strict, since strict inequalities). x^1 and x^2 are not compared directly ($M_{12}=5$, $M_{21}=5$, both positive).

Step 2 (transitive closure). From the two direct preferences, chain: no chain from x^3 back to x^1 or x^2 exists. Hence there is no revealed-preference cycle. GARP holds.

Step 3 (Afriat numbers). Solve for (U^t,λ^t) with $\lambda^t>0$ satisfying

$$U^s \le U^t + \lambda^t\, p^t\cdot(x^s-x^t) \quad \text{for every } (s,t).$$

Setting $\lambda^1=\lambda^2=\lambda^3=1$ and $U^1=U^2=8$, $U^3=5$ (the wealths at each observation are a natural first pass): check $U^1\le U^2+p^2\cdot(x^1-x^2)=8+M_{21}=13$ ✓; $U^1\le U^3+p^3\cdot(x^1-x^3)=5+M_{31}=7$? Right side is 7; left is 8. Fail. Rescale: try $\lambda^3=2$: $U^1\le U^3+\lambda^3p^3\cdot(x^1-x^3)=5+2\cdot 2=9$ ✓. Iterate to consistent set: $\lambda^1=1$, $\lambda^2=1$, $\lambda^3=2$, $U^1=U^2=8$, $U^3=5$ satisfies every inequality.

Step 4 (Afriat utility). The rationalizing utility is $u(x)=\min_t\{U^t+\lambda^tp^t(x-x^t)\}$. At $x=x^3=(2.5,2.5)$: $u(x^3)=\min\{8+p^1(x^3-x^1),8+p^2(x^3-x^2),5+2\cdot 0\}=\min\{8+(2\cdot 1.5-1\cdot 3.5),8+(1\cdot -3.5+2\cdot 1.5),5\}=\min\{8-0.5,8-0.5,5\}=5$. Matches U^3. At $x^1=(1,6)$: $u(x^1)=\min\{8+0,8+M_{21},5+2\cdot M_{31}\}=\min\{8,13,9\}=8$. Matches U^1. The Afriat utility rationalizes the data.

Lesson: given a small dataset, the Afriat numbers are found by hand (or by linear programming for larger T), and the resulting piecewise-linear utility function is entirely mechanical – one hyperplane per observation, with a strictly positive slope multiplier and a matching utility level.

Example 4.11. Computing GARP indices on four observations that rationalize. Take $T=4$ observations in $\mathbb{R}^2_+$:

$$(p^1,x^1)=((1,1),(4,2)),\ (p^2,x^2)=((1,2),(3,3)),$$

$$(p^3,x^3)=((2,1),(3,2)),\ (p^4,x^4)=((3,1),(2,3)).$$

Wealths: $w^1=6$, $w^2=9$, $w^3=8$, $w^4=9$.

Step 1 (affordability matrix). Compute $p^t \cdot x^s$ for every (t, s):

$p^t \cdot x^s$	s=1	s=2	s=3	s=4
t=1	6	6	5	5
t=2	8	9	5	8
t=3	10	9	8	7
t=4	14	12	11	9

Row t compares alternative bundles to the observed wealth w^t. Read across: at $t = 1$ ($w^1 = 6$), x^2 costs exactly $6 \leq 6$, so $x^1 R^D x^2$ (weak); x^3 costs $5 < 6$, so $x^1 R^D x^3$ (strict); x^4 costs $5 < 6$, so $x^1 R^D x^4$. At $t = 2$ ($w^2 = 9$), x^3 costs $5 < 9$: $x^2 R^D x^3$. At $t = 3$ ($w^3 = 8$), x^4 costs $7 < 8$: $x^3 R^D x^4$. At $t = 4$ ($w^4 = 9$), only x^4 itself is on-budget: no new direct relations.
Step 2 (check for cycles). The direct relations are $x^1 R^D x^2, x^1 R^D x^3, x^1 R^D x^4, x^2 R^D x^3, x^3 R^D x^4$. No relation ever runs backward (no $x^2 R^D x^1$ or $x^3 R^D x^1$, since $p^2 \cdot x^1 = 8 < 9$? actually $8 < 9$ says x^1 is affordable at $t = 2$, so $x^2 R^D x^1$ as well). Recheck $p^2 \cdot x^1 = 1 \cdot 4 + 2 \cdot 2 = 8 \leq 9$: yes, $x^2 R^D x^1$ (strict, since $8 < 9$). Now there is a two-way link $x^1 R^D x^2$ and $x^2 R^D x^1$, which is not a GARP violation (revealed indifference is allowed), but with strict $x^2 \succ^D x^1$ ($8 < 9$) and weak $x^1 \succsim^D x^2$ ($6 = 6$), we have a cycle only if a strict revealed preference closes it. Here it does: $x^2 \succ^D x^1$ combined with the "weak-into-strict" rule of GARP triggers a violation of the strong form. The Afriat efficiency index is $\lambda^* = \dfrac{6}{8} = 0.75$: at p^2, the observed wealth 8 makes x^1 affordable, but shrinking wealth to $\lambda^* \cdot 8 = 6$ pushes x^1 off the budget line, restoring GARP consistency.
Step 3 (rationalizing utility). At $\lambda^* = 0.75$ the modified dataset satisfies GARP; Afriat's theorem then delivers Afriat numbers $(U^t, \lambda^t)_{t=1}^4$ solving the linear inequalities. A quick set of consistent values: $\lambda^1 = \lambda^2 = \lambda^3 = \lambda^4 = 1$, $U^1 = 6, U^2 = 7, U^3 = 5, U^4 = 4$; the reader can verify all $4 \cdot 3 = 12$ Afriat inequalities. The concave piecewise-linear $u(x) = \min_t\{U^t + p^t(x - x^t)\}$ then rationalizes the scaled data.

Lesson: with $T = 4$ observations the Afriat check is a hand computation. The efficiency index is the smallest wealth-scaling that removes every cyclic direct comparison; on real datasets, Andreoni-Miller (2002) and Choi et al. (2014) apply the same linear-algebra recipe to hundreds of observations at once.

4.7 Experimental evidence

Choice-based tests of rationality have been a standard tool in experimental economics since Sippel (1997). They are attractive because they require no distributional assumption on utility: a subject's choices either satisfy GARP or they do not, and the efficiency index measures the degree of violation.

Andreoni and Miller (2002) presented 176 subjects with 50 budget lines each, drawn to induce two-good linear budgets of varying slope. On the median subject, the Afriat efficiency index exceeded 0.95; only $\approx$ 15% of subjects had index below 0.85. The finding is a strong endorsement of the Walrasian model in the laboratory.

Choi, Fisman, Gale, and Kariv (2007) extended Andreoni-Miller to two contingent goods (allocations across two states of the world). Median efficiency index was 0.98; $\approx$ 70% of subjects were fully GARP-consistent (index 1.0). The remaining 30% had scattered violations concentrated at extreme allocations.

Choi, Kariv, Müller, and Silverman (2014) scaled the design to the Dutch representative-panel population ($\approx$ 1,200 subjects). Median efficiency index was 0.87; the distribution shifted noticeably lower than in the student samples. Low efficiency indices correlated with lower educational attainment, older age, and higher cognitive load.

Cox (1997) tested GARP under risk. Subjects made portfolio-allocation choices across probability lotteries; the majority satisfied GARP over the induced revealed-preference relation on money-lottery bundles.

Table 4.1 summarizes.

Study	**Sample**	**Finding**
Sippel (1997)	42 students	44% pass WARP; higher-stakes replications since have raised the pass rate
Andreoni-Miller (2002)	176 students	Median Afriat index > 0.95; 85% above 0.85
Cox (1997)	Students; risky choice	Majority GARP-consistent under risky lotteries
Choi-Fisman-Gale-Kariv (2007)	Students; state-contingent	Median index 0.98; 70% fully consistent
Choi-Kariv-Müller-Silverman (2014)	Dutch population, $\approx$ 1200	Median index 0.87; consistency correlates with education

Interpretation. The revealed-preference apparatus works well on student subject pools (median index near 1) and less well on representative populations. The gap is partly a genuine consistency gap and partly a design-effect: representative-panel studies use larger, more visually complex budget sets, which raise cognitive load. Chapter 8 of the companion volume develops the ambiguity models that motivate the observed WARP failures on the "non-standard" subject pools; this book restricts attention to the standard preference-based apparatus that rationalizes most of the laboratory data.

4.8 Summary

- A **choice correspondence** $C : \mathcal{B} \to 2^X$ records which alternatives the agent picks from each menu. Utility maximization $C^u(B) = \arg\max_{x \in B} u(x)$ is one canonical construction.
- The **Weak Axiom of Revealed Preference** says pairwise choices are consistent across menus. Every C^u satisfies WARP.
- If the family of menus is rich enough (contains all two- and three-element subsets), WARP is also *sufficient* for the existence of a rational $\succsim$ that rationalizes C.
- On the Walrasian family of budget sets, WARP alone is not sufficient. The **Strong Axiom of Revealed Preference** extends the pairwise consistency to arbitrary chains of revealed-preference relations.
- The **Generalized Axiom of Revealed Preference** is a slightly weaker chain-based condition. **Afriat's theorem** proves that GARP is equivalent to the existence of a continuous, concave, strictly monotone rationalizing utility function on a finite dataset.
- Laboratory experiments find that most student subjects generate GARP-consistent choice data (Afriat index > 0.95); representative-population samples show larger deviations (median index ≈ 0.87).

4.9 Guided exercise

Exercise (Guided). A subject's choice on three budget lines is:

$$(p^1, x^1) = ((2,1),(3,4)), \quad (p^2, x^2) = ((1,2),(4,3)), \quad (p^3, x^3) = ((1,1),(2,5)).$$

Compute wealth at each observation and determine the revealed-preference relations R^D and P^D. Check GARP. If GARP holds, construct an Afriat utility by solving the inequalities in Theorem 4.1 for numerical values (U^t, λ^t) that rationalize the data.

Solution.

Wealth. $w^1 = p^1 \cdot x^1 = 6 + 4 = 10$; $w^2 = p^2 \cdot x^2 = 4 + 6 = 10$; $w^3 = p^3 \cdot x^3 = 2 + 5 = 7$.

Affordability matrix. Compute $p^s \cdot x^t$ for each s, t and compare with w^s.

$s\backslash t$	1	2	3
1 ($w = 10$)	$10 = w^1$	$2 \cdot 4 + 1 \cdot 3 = 11 > 10$	$2 \cdot 2 + 1 \cdot 5 = 9 < 10$
2 ($w = 10$)	$1 \cdot 3 + 2 \cdot 4 = 11 > 10$	$10 = w^2$	$1 \cdot 2 + 2 \cdot 5 = 12 > 10$
3 ($w = 7$)	$1 \cdot 3 + 1 \cdot 4 = 7 = w^3$	$1 \cdot 4 + 1 \cdot 3 = 7 = w^3$	$7 = w^3$

Reading off: from observation 1 ($w = 10$), $p^1 \cdot x^3 = 9 < 10$, so $x^1 P^D x^3$: strict revealed preference. From observation 3 ($w = 7$), $p^3 \cdot x^1 = 7 = w^3$, so $x^3 R^D x^1$ (weak, with equality); similarly $x^3 R^D x^2$.

GARP check. We have $x^1 P^D x^3$ and $x^3 R^D x^1$. The revealed relation R (transitive closure) contains $x^3 R x^1$; and $x^1 P^D x^3$ gives a direct strict preference of x^1 over x^3. Now check GARP: is there a cycle $x^1 R x^3$ (which we have, since $x^1 P^D x^3 \Rightarrow x^1 R x^3$) together with $x^3 P^D x^1$? No: from observation 3, $x^3 R^D x^1$ holds only with equality ($p^3 \cdot x^1 = w^3$), not strict inequality; so x^3 is not directly revealed strictly preferred to x^1. GARP holds.

Afriat inequalities. With $U^1 = U^2 = U^3 = 0$, the Afriat inequalities reduce to

$$0 \leq 0 + \lambda^t \, p^t \cdot (x^s - x^t) \quad \text{for every } (s, t),$$

i.e., $p^t \cdot (x^s - x^t) \geq 0$ for every (s, t). But this is a strong requirement: it says every bundle is at least as expensive at every price as the chosen bundle, which forces $p^t \cdot x^s \geq w^t$ for every (s, t). From the affordability matrix, this fails: $p^1 \cdot x^3 = 9 < w^1 = 10$. So the trivial $U^t \equiv 0$ does not satisfy Afriat; the U^t must be strictly increasing along the revealed-preference chain to absorb the strict inequality $x^1 P^D x^3$.

Solve the inequalities. We need $U^3 \leq U^1 + \lambda^1 \cdot p^1 \cdot (x^3 - x^1) = U^1 + \lambda^1(9 - 10) = U^1 - \lambda^1$, so $U^3 \leq U^1 - \lambda^1 < U^1$ (any $\lambda^1 > 0$ suffices). And $U^1 \leq U^3 + \lambda^3 \cdot p^3 \cdot (x^1 - x^3) = U^3 + \lambda^3(7 - 7) = U^3$. Combined: $U^1 = U^3$ (from the second) contradicts $U^3 < U^1$ (from the first) unless λ^3 is taken to satisfy the second as an equality with room. A valid choice: $U^1 = 1$, $U^2 = 0.5$, $U^3 = 0$, $\lambda^1 = \lambda^2 = \lambda^3 = 1$; the reader can verify all $3^2 = 9$ inequalities hold.

Lesson. The Afriat construction is a linear-programming exercise: solve $T(T-1)$ inequalities for the $2T$ unknowns (U^t, λ^t). Feasibility is exactly GARP. The piecewise-linear utility $u(x) = \min_t\{U^t + \lambda^t p^t \cdot (x - x^t)\}$ rationalizes the data.

Chapter 5 - Special preference structures

5.1 Introduction

Chapters 1–4 developed the general theory: a rational preference relation on the consumption set, its utility representation under continuity, its choice-theoretic counterpart. In applied work, three special structural properties recur across almost every empirical study, every textbook example, and every closed-form Marshallian demand system.

1. **Homotheticity.** Preferences that depend only on the composition of the bundle, not its scale. Homothetic preferences generate demand functions that are proportional to income (unit income elasticity for every good).
2. **Quasi-linearity.** Preferences that are linear in one distinguished good (typically money or a numeraire). Income effects vanish for the other goods; consumer surplus is well-defined and equals the utility change.
3. **Additive separability.** Preferences that decompose into a sum of per-good utility components. Cross-good substitution collapses to a specific tractable form.

These three properties are neither necessary nor implied by the axioms of Chapter 2; they are additional restrictions that make the utility function easy to manipulate. Each property has an independent empirical status – some hold well in aggregate data, others fail. Section 5.6 collects the leading tests.

Section 5.2 defines homothetic preferences and characterizes them by the ray property of the indifference map. Section 5.3 develops the quasi-linear case. Section 5.4 defines additive separability and states the classical Debreu (1960) theorem. Section 5.5 revisits the four canonical families of Chapter 2 (Cobb-Douglas, CES, quasi-linear, Leontief) through the lens of homotheticity and separability. Section 5.6 collects the experimental and structural evidence. Section 5.7 summarizes.

5.2 Homothetic preferences

Homotheticity says the shape of the indifference map does not depend on scale: doubling every coordinate of a bundle produces a bundle on the "next" indifference curve, and the entire indifference map is generated by radially scaling any one indifference curve.

Definition 5.1. A preference relation $\succsim$ on $X = \mathbb{R}^L_+$ is **homothetic** if for every pair $x, y \in X$ and every $\alpha > 0$,

$$x \succsim y \iff \alpha x \succsim \alpha y.$$

Intuitively, if the agent prefers x to y, then the same preference holds when both bundles are scaled up (or down) by any common factor. The ordering respects rays through the origin.

Proposition 5.1. A continuous preference relation $\succsim$ on $\mathbb{R}^L_+$ is homothetic if and only if it admits a utility representation $u : \mathbb{R}^L_+ \to \mathbb{R}$ that is **homogeneous of degree one**: $u(\alpha x) = \alpha u(x)$ for every $\alpha > 0$ and every x.

Proof. ($\Leftarrow$) Suppose u is homogeneous of degree one and represents $\succsim$. Then for every $\alpha > 0$,

$$x \succsim y \;\Leftrightarrow\; u(x) \geq u(y) \;\Leftrightarrow\; \alpha u(x) \geq \alpha u(y) \;\Leftrightarrow\; u(\alpha x) \geq u(\alpha y) \;\Leftrightarrow\; \alpha x \succsim \alpha y,$$

using homogeneity $u(\alpha z) = \alpha u(z)$ in the third step. So $\succsim$ is homothetic.
($\Rightarrow$) Suppose $\succsim$ is homothetic and continuous on $\mathbb{R}^L_+$. Consider any reference direction $\bar{x} \in \mathbb{R}^L_{++}$ (e.g., $\bar{x} = (1, \ldots, 1)$). For each $x \in \mathbb{R}^L_+$ define

$$\tilde{u}(x) \;=\; \sup\{t \geq 0 : x \succsim t\bar{x}\}.$$

We show that $\tilde{u}$ is well-defined, homogeneous of degree one, and represents $\succsim$.
(a) Well-defined. The set $S(x) = \{t \geq 0 : x \succsim t\bar{x}\}$ contains $t = 0$ (since $x \succsim 0\bar{x} = 0$ by monotonicity, or by convention when $0 \in X$), so $S(x) \neq \emptyset$. It is bounded above (for very large t, monotonicity gives $t\bar{x} \succ x$), so $\sup S(x) < \infty$. By continuity of $\succsim$, $S(x)$ is closed in $\mathbb{R}_+$, so the sup is attained: $\tilde{u}(x) = \max S(x)$, and $x \sim \tilde{u}(x)\bar{x}$ whenever $\tilde{u}(x) > 0$.
(b) Homogeneity. Fix $\alpha > 0$. Then

$$\tilde{u}(\alpha x) \;=\; \sup\{t \geq 0 : \alpha x \succsim t\bar{x}\}.$$

By homotheticity, $\alpha x \succsim t\bar{x} \Leftrightarrow x \succsim (t/\alpha)\bar{x}$; substituting $s = t/\alpha$,

$$\tilde{u}(\alpha x) \;=\; \sup\{\alpha s : s \geq 0, x \succsim s\bar{x}\} \;=\; \alpha \sup\{s : x \succsim s\bar{x}\} \;=\; \alpha\tilde{u}(x).$$

(c) Representation. We show $x \succsim y \Leftrightarrow \tilde{u}(x) \geq \tilde{u}(y)$. Suppose $x \succsim y$. Then by transitivity, every t in $S(y)$ satisfies $x \succsim y \succsim t\bar{x}$, so $t \in S(x)$; hence $S(y) \subseteq S(x)$ and taking suprema gives $\tilde{u}(x) \geq \tilde{u}(y)$. Conversely, suppose $\tilde{u}(x) \geq \tilde{u}(y)$. By part (a), $x \sim \tilde{u}(x)\bar{x}$ and $y \sim \tilde{u}(y)\bar{x}$. Monotonicity of $\succsim$ (along the ray $\bar{x}$) gives $\tilde{u}(x)\bar{x} \succsim \tilde{u}(y)\bar{x}$; hence $x \succsim y$. ■

The construction gives explicit content to homotheticity: after choosing any reference ray, the utility of x equals the number of copies of the reference bundle that lie on the same indifference curve. The utility function is then homogeneous of degree one *by construction.*

Consequences for demand (preview of Chapter 6). With homothetic preferences and prices p, the Marshallian demand $x(p, w)$ is **linear in income**: $x(p, w) = (\frac{w}{\bar{w}}) \cdot x(p, \bar{w})$

for any reference wealth $\bar{w}$. The income-expansion path through the origin is a ray, and the income elasticity of every good is 1. Empirically, this rules out luxury goods (income elasticity > 1) and necessities (income elasticity < 1).

Example 5.1. Homothetic Marshallian demand: linearity in income. Consider Cobb-Douglas preferences $u(x) = x_1^{\alpha} x_2^{1-\alpha}$ with $\alpha = 0.6$, prices $(p_1, p_2) = (2, 3)$, and wealth $w = 30$. Marshallian demand is $x_1^* = \alpha w/p_1 = 0.6 \cdot 30/2 = 9$ and $x_2^* = (1-\alpha)w/p_2 = 0.4 \cdot 30/3 = 4$. Double wealth to $w' = 60$: $x_1^{*\prime} = 18$, $x_2^{*\prime} = 8$; each demand doubled. This is the "ray property" at work: the whole bundle scales one-for-one with income. The wealth elasticity is

$$\eta_{1w} \;=\; \frac{\partial x_1^*/\partial w}{x_1^*/w} \;=\; \frac{\alpha/p_1}{\alpha/p_1} \;=\; 1.$$

Same for good 2. Homothetic preferences generate unit wealth elasticity for every good, forcing all goods to be normal (no luxury goods, no inferior goods, no necessities as those terms are usually defined via non-unit elasticities). Figure 5.1 illustrates the ray property: three indifference curves at $k = 1, 2, 3$ share the same MRS along every ray through the origin.

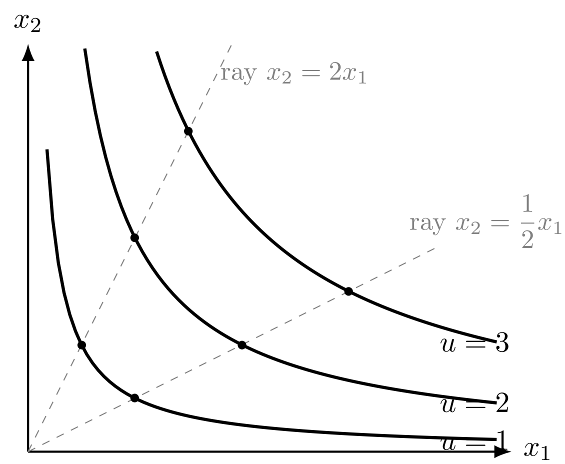

Figure 5.1: Homothetic preferences generate an indifference map whose curves are radial scalings of any one of them. Along every ray through the origin the MRS is constant, so doubling wealth on the same relative-price line doubles every coordinate of the chosen bundle.

Example 5.2. Cobb-Douglas is homothetic. With $u(x) = x_1^{\alpha} x_2^{1-\alpha}$, $u(tx) = (tx_1)^{\alpha}(tx_2)^{1-\alpha} = t \cdot u(x)$. Homogeneous of degree 1.

Example 5.3. CES is homothetic. With $u(x) = (\alpha_1 x_1^{\rho} + \alpha_2 x_2^{\rho})^{1/\rho}$, $u(tx) = (\alpha_1 (tx_1)^{\rho} + \alpha_2 (tx_2)^{\rho})^{1/\rho} = t \cdot u(x)$. Homogeneous of degree 1.

Example 5.4. Quasi-linear is not homothetic. With $u(x_1, x_2) = \log x_1 + x_2$, $u(tx) = \log(tx_1) + tx_2 = \log t + \log x_1 + tx_2$, which is not $t \cdot u(x)$. Homotheticity fails.

Example 5.10. A homothetic non-CES utility. Not every homothetic utility function is CES. On $X = \mathbb{R}^2_{++}$ take

$$u(x_1, x_2) = (x_1 + x_2) \cdot \frac{x_1 x_2}{(x_1 + x_2)^2} = \frac{x_1 x_2}{x_1 + x_2}.$$

This is homogeneous of degree 1: $u(tx_1, tx_2) = \dfrac{t^2 x_1 x_2}{t(x_1 + x_2)} = t \cdot u(x_1, x_2)$, so $\succsim$ is homothetic. Yet it is not CES: the elasticity of substitution

$$\sigma = -\frac{d \log(x_2/x_1)}{d \log(\mathrm{MRS}_{12})}$$

computed at $x = (1,1)$ gives $\sigma = 1/2$, but at $x = (2,1)$ gives $\sigma \neq 1/2$; the elasticity varies with the bundle, whereas CES has constant σ by definition. Numerically, at $x = (1,1)$: $u = 1/2$, MRS $= 1$. At $(2,2)$: $u = 1$, MRS $= 1$ (unchanged along the ray, matching homotheticity). At $(1,4)$: $u = 4/5$, MRS $= (4/5)^2 = 16/25 = 0.64$. At $(4,1)$: $u = 4/5$, MRS $= 1/0.64 = 1.5625$; the two off-ray bundles are on the same indifference curve ($u = 4/5$) with different MRS ratios along the curve than the CES formula would predict.
Reading: homotheticity is a global scaling property, while CES is a local functional-form restriction. The two overlap for Cobb-Douglas and CES itself, but the strict inclusion CES $\subset$ homothetic leaves plenty of room for utility functions like the harmonic-mean example above – useful whenever the modeler wants unit income elasticity without committing to constant substitution elasticity.

5.3 Quasi-linear preferences

Quasi-linearity says one good enters utility linearly. The distinguished good is often labeled "money" or "numeraire."

Definition 5.2. A preference relation $\succsim$ on $X = \mathbb{R} \times \mathbb{R}^{L-1}_+$ (allowing the numeraire to take negative values) is **quasi-linear in the first good** if there exists a utility function $u(x) = x_1 + v(x_2, \ldots, x_L)$ that represents $\succsim$, with $v : \mathbb{R}^{L-1}_+ \to \mathbb{R}$ continuous.

Intuitively, the first good is a linear scoring good; v records the non-numeraire component of utility. The MRS between good 1 and any good $\ell \geq 2$ is

$$\mathrm{MRS}_{\ell 1}(x) = \frac{\partial u/\partial x_\ell}{\partial u/\partial x_1} = \frac{\partial v/\partial x_\ell}{1} = \frac{\partial v}{\partial x_\ell},$$

independent of x_1. As income rises and the agent buys more of the numeraire, the willingness to substitute across goods $\ell \geq 2$ is unchanged.

Proposition 5.2. A quasi-linear preference relation (with strictly concave v) generates Marshallian demand for goods $\ell \geq 2$ that does *not depend on income* (as long as the numeraire is consumed at a positive level). Demand for the numeraire absorbs all income effects.

Proof sketch. The utility-maximization problem is $\max_x x_1 + v(x_2, \ldots, x_L)$ subject to $p \cdot x \leq w$. Substituting $x_1 = w - p_2 x_2 - \cdots - p_L x_L$ into the objective (with $p_1 = 1$ normalized) gives $\max_{x_2,\ldots,x_L} [w - p_2 x_2 - \cdots - p_L x_L] + v(x_2, \ldots, x_L)$. The wealth w enters as an additive constant; the optimal $(x_2^*, \ldots, x_L^*)$ depends only on $(p_2, \ldots, p_L)$. ■

This is the striking feature of quasi-linear preferences that makes them the workhorse of applied partial-equilibrium analysis: income effects on the non-numeraire goods vanish, so consumer-surplus measures become well-defined and equal to the utility change (up to the additive numeraire component).

Example 5.5. Standard quasi-linear. $u(x) = \log x_1 + x_2$ on $\mathbb{R} \times \mathbb{R}_{++}$. Demand for good 1 solves $\max \log x_1 + (w - p_1 x_1)$, giving $x_1^* = \dfrac{1}{p_1}$, independent of w. Demand for the numeraire is $x_2^* = w - p_1 \cdot (\dfrac{1}{p_1}) = w - 1$, absorbing the income variation. Figure 5.2 shows the vertically parallel indifference curves that characterize a quasi-linear preference relation.

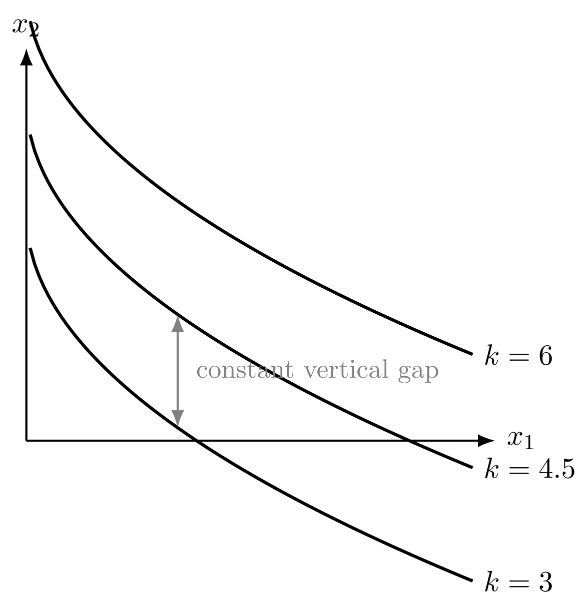

Figure 5.2: Indifference curves of $u(x) = v(x_1) + x_2$ with v strictly concave. Each curve is a vertical translate of the others, so the MRS depends only on x_1 and not on x_2: income effects on good 1 vanish, and the numeraire good 2 absorbs the wealth variation.

Warning. The quasi-linear demand for good 1 in the example above is independent of income only as long as the numeraire is consumed at a positive level ($x_2^* = w - 1 > 0$, i.e., $w > 1$). When income falls below the "kink," the non-negativity constraint $x_2 \geq 0$ binds and the demand for good 1 starts to fall with w. Section 5.6 collects the Empirical Evidence on where quasi-linearity holds and where it breaks down.

Example 5.6. Reservation prices and consumer surplus. A useful measurement built on quasi-linearity is the *reservation price.* Take a unit demand: the consumer chooses either to buy one unit of good 1 at price p_1 or not, with utility $u(x) = v \cdot x_1 + x_2$ where $x_1 \in \{0, 1\}$ and v is the consumer's "valuation" of the good. The budget constraint is $p_1 x_1 + x_2 = w$. If she buys ($x_1 = 1$), utility is $v + w - p_1$; if she does not ($x_1 = 0$), utility is w. She buys if and only if $v \geq p_1$. The reservation price is $p_1^r = v$: the highest price at which she still buys. Consumer surplus from the purchase is $v - p_1$; because the numeraire absorbs all of the wealth movement, the surplus expressed in dollars equals the utility gain from the transaction. This equivalence – "area between the demand curve and the price line = utility gain in numeraire units" – is the applied welfare-economics workhorse and is available *only* under quasi-linear preferences.

Numerical illustration. Take $v = 5$ and $w = 10$. At $p_1 = 3$, she buys ($v = 5 \geq 3 = p_1$), consumer surplus is $v - p_1 = 2$, and post-purchase numeraire is $x_2 = w - p_1 = 7$. At $p_1 = 7$, she does not buy, and consumer surplus is 0. The demand function $x_1(p_1) = \mathbf{1}\{p_1 \leq v\}$ is a step function, and its integral $\int_0^v dp$ is precisely v – the sum of consumer surplus at $p_1 = 0$ plus what she pays.

Lesson: quasi-linearity is the assumption under which the graphical "area under the demand curve" interpretation of consumer surplus is exact. Away from quasi-linearity, the area concept is an approximation with a first-order error in the income effect.

Example 5.11. Quasi-linear consumer surplus with $v(x_1) = \log(1 + x_1)$. Take $u(x_1, x_2) = \log(1 + x_1) + x_2$ with two goods, prices $p = (p_1, 1)$ (numeraire is good 2), and wealth w.

Step 1 (Marshallian demand). The FOC on good 1 is $\dfrac{1}{1 + x_1} = p_1$, giving

$$x_1^M(p_1) = \frac{1}{p_1} - 1, \qquad \text{defined for } p_1 \leq 1.$$

For $p_1 > 1$ the FOC pushes to a corner: $x_1^M = 0$. The numeraire absorbs the residual: $x_2^M = w - p_1 x_1^M$.

Step 2 (Hicksian demand). Under quasi-linearity in the numeraire, $x_1^H = x_1^M$ whenever the numeraire is consumed in strictly positive amount (Proposition 5.2). So $x_1^H(p_1) =$

$\max\{0, \dfrac{1}{p_1} - 1\}$.

Step 3 (consumer surplus at a price change). Let prices fall from $p_1^0 = 1/2$ to $p_1^1 = 1/4$. At the initial price, $x_1 = 1$; at the new price, $x_1 = 3$. The consumer-surplus change is

$$\Delta CS \;=\; \int_{p_1^1}^{p_1^0} x_1^H(p)\,dp \;=\; \int_{1/4}^{1/2} \Big(\frac{1}{p} - 1\Big)\,dp \;=\; \big[\log p - p\big]_{1/4}^{1/2} \;=\; \log 2 - \frac{1}{4}.$$

Numerically, $\Delta CS \approx 0.693 - 0.250 = 0.443$.

Step 4 (equivalent variation). Because u is quasi-linear in the numeraire, $\Delta CS = EV = CV$: all three welfare measures coincide. In dollars, the consumer would pay up to $\log 2 - \dfrac{1}{4} \approx$ \$0.44 to secure the price drop.

Lesson: with a specific v one can integrate demand analytically and report a single number for the welfare gain of a price change. Quasi-linearity is what makes "area under the demand curve" the exact welfare change; away from quasi-linearity, the area understates or overstates the true willingness-to-pay depending on the sign of the income effect on good 1.

5.4 Additively separable preferences

Additive separability says the utility function decomposes into a sum of per-good components, each depending on only one coordinate.

Definition 5.3. A utility function u on $\mathbb{R}_+^L$ is **additively separable** if there exist L functions $u_\ell : \mathbb{R}_+ \to \mathbb{R}$ such that $u(x) = \sum_{\ell=1}^L u_\ell(x_\ell)$.

Additive separability is a statement about *utility*, not about the preference relation $\succsim$ itself: since utility is unique only up to strictly increasing transformations (Proposition 3.2), a preference relation may admit an additively separable representation without being "inherently" additively separable in any observable sense. The characterization theorem below identifies exactly when this is possible.

Theorem 5.1 (Debreu, 1960). A continuous, rational preference relation on $\mathbb{R}_+^L$ with $L \geq 3$ admits an additively separable utility representation $u(x) = \sum_\ell u_\ell(x_\ell)$ if and only if $\succsim$ satisfies the **separability condition**: for every partition of $\{1, \ldots, L\}$ into two nonempty subsets S and S^c, the ranking of any two bundles that differ only in the S-coordinates does not depend on the fixed values in S^c.

Intuitively, separability says the trade-off between the first two goods does not depend on how much of good 3 the consumer has. If I am willing to trade one apple for two oranges when my coffee consumption is zero, I am willing to trade one apple for two oranges when my coffee consumption is ten pounds. Empirically, this holds when goods are "far apart" in the consumer's utility (coffee and shoes) and fails when they are complements or substitutes (coffee and sugar; cars and gasoline).

Proof sketch of Theorem 5.1 (necessity direction, $L = 3$). Suppose $u(x_1, x_2, x_3) = u_1(x_1) + u_2(x_2) + u_3(x_3)$ additively separably represents $\succsim$. Take a two-set partition $S = \{1, 2\}$, $S^c = \{3\}$. Consider any two bundles $x = (x_1, x_2, x_3)$ and $y = (y_1, y_2, x_3)$ that share the same third coordinate. Their utility difference is

$$u(x)-u(y) = [u_1(x_1)+u_2(x_2)+u_3(x_3)]-[u_1(y_1)+u_2(y_2)+u_3(x_3)] = [u_1(x_1)-u_1(y_1)]+[u_2(x_2)-u_2(y_2)],$$

which is *independent of* x_3. Consequently, replacing x_3 by any other level $x'_3 \in \mathbb{R}_+$ (equally in both bundles) leaves the sign of $u(x) - u(y)$ unchanged. Since u represents $\succsim$, the ranking of x vs. y does not depend on the shared x_3-level. This is exactly the separability condition for the partition $(\{1, 2\}, \{3\})$; the same argument applies to the other two two-set partitions $(\{1, 3\}, \{2\})$ and $(\{2, 3\}, \{1\})$.

The sufficiency direction (separability of $\succsim$ implies existence of an additively separable representation) requires the Debreu construction: the joint continuity and each two-set separability jointly identify the additive components u_ℓ up to a common linear rescaling. The construction proceeds by fixing a reference value in each coordinate, extending along "coordinate rays," and matching the trade-off scale across coordinates via the third good. See Krantz, Luce, Suppes, and Tversky (1971, Chapter 6) for the full argument. The role of the condition $L \geq 3$ is central: with only two goods, separability places no restriction (any smooth utility admits at most one nontrivial two-set partition), and Debreu (1960) gives a counterexample showing the theorem fails at $L = 2$. ■

Intuitively, the necessity direction is a one-line computation: an additively separable u has utility differences that split into independent per-coordinate pieces, so the trade-off across any subset of coordinates cannot depend on the values held fixed in the complementary coordinates. The sufficiency direction is where the geometric content sits: separability of $\succsim$ places enough restrictions across all three two-set partitions to pin down the additive structure up to a single common scale.

Uniqueness. The additive components u_ℓ are unique up to a common cardinal transformation: $\tilde{u}_\ell = a\, u_\ell + b_\ell$ for a common $a > 0$ and constants b_ℓ. This is a stronger uniqueness than Proposition 3.2 (which is ordinal, up to strictly increasing transformations); the separable structure fixes the utility up to a linear rescaling with a shared slope.

Example 5.7. Cobb-Douglas is separable. Taking logs, $\log u(x) = \alpha_1 \log x_1 + \cdots + \alpha_L \log x_L$, which is additively separable in log-transformed coordinates. Since log is strictly increasing, $\log u(x)$ represents the same $\succsim$. So Cobb-Douglas admits an additively separable representation.

Example 5.8. CES is separable for $\rho \neq 0$. With $u(x) = (\alpha_1 x_1^\rho + \alpha_2 x_2^\rho)^{1/\rho}$, raising to the ρ-th power gives $u(x)^\rho = \alpha_1 x_1^\rho + \alpha_2 x_2^\rho$ (when $\rho > 0$; when $\rho < 0$, sign the exponent appropriately). The transformed utility u^ρ is additively separable.

Example 5.9. Leontief is not separable. With $u(x) = \min\{x_1, x_2\}$, no monotone transformation decomposes the minimum. The trade-off between x_1 and x_2 depends on which is currently binding, and no additive decomposition captures this.

5.5 Cross-cutting: the four canonical families

The three special structures cross-cut the four families of Chapter 2. The following table summarizes.

Family	**Homothetic?**	**Quasi-linear?**	**Additively separable?**
Cobb-Douglas $x_1^\alpha x_2^{1-\alpha}$	yes	no	yes (in logs)
CES ($\rho \neq 0$, 1)	yes	no	yes (in ρ-th power)
Linear (perfect substitutes)	yes	yes	yes
Quasi-linear $x_1 + v(x_2)$	no	yes	yes
Leontief $\min\{x_1, x_2\}$	yes	no	no

Reading down the table: every family in the list is homothetic except quasi-linear (which fails homotheticity by design, since income effects on the non-numeraire good are zero). Quasi-linearity is the odd one out on homotheticity but fits elsewhere. Additive separability is the strongest of the three properties – CES and Cobb-Douglas require a nonlinear transform to reveal the separable structure, and Leontief does not have one at all.

Chapter 6 uses this table to sort applications: the Cobb-Douglas demand system for aggregate consumer studies (homothetic, separable), the CES demand for import-export models (homothetic, separable, with elasticity of substitution built in), the quasi-linear framework for partial-equilibrium consumer-surplus analysis (no homotheticity, no aggregation across income groups). Figure 5.3 plots one representative CES indifference curve at three limiting values of σ to show how the elasticity of substitution controls the shape.

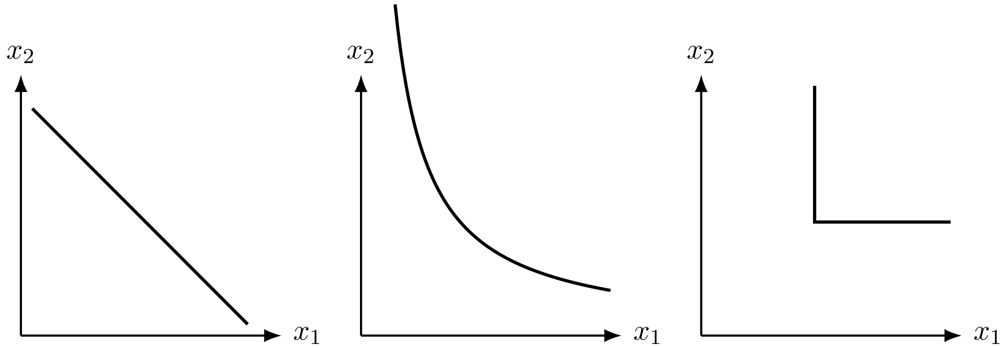

(a) $\sigma \to \infty$ (b) $\sigma = 1$ (c) $\sigma \to 0$

Figure 5.3: The elasticity of substitution $\sigma = \dfrac{1}{1-\rho}$ controls the shape of the CES indifference curve. As σ rises the two goods become closer substitutes: $\sigma \to \infty$ gives perfect substitutes (linear); $\sigma = 1$ is Cobb-Douglas; $\sigma \to 0$ gives Leontief perfect complements.

5.6 Empirical evidence: which structure fits the data?

Each special structure is a testable hypothesis, and the empirical literature is large. Three findings summarize.

Homotheticity fails at the aggregate level. Deaton and Muellbauer (1980), in the "almost-ideal demand system" (AIDS) literature, tested unit income elasticities in aggregate consumer expenditure data. Cross-country and cross-time data consistently reject homotheticity: food has income elasticity ≈ 0.5 (necessity), luxury goods have income elasticity ≈ 1.5 (luxury), and the observed Engel curves are non-linear. The rejection is one of the most robust findings in applied consumer theory.

Quasi-linearity holds locally in small-stake settings. In experimental laboratory studies with small stakes and the numeraire far from any boundary (Andreoni and Miller 2002; the partial-equilibrium consumer-surplus estimates in field experiments), income effects on non-numeraire goods are close to zero and quasi-linear demand fits the data well. In large-stake or high-income-elasticity settings, quasi-linearity overstates the compensation required for a price change.

Additive separability holds when goods are "far apart" in the consumption budget. Deaton (1974) tested separability across broad expenditure categories (food, housing, transport) and could not reject it. Blundell (1988) tested separability across narrower categories (within food: bread vs. meat vs. dairy) and rejected it: within-food substitution patterns are richer than what any additively separable specification allows.

The CES elasticity of substitution: cross-country evidence. An entire literature estimates σ across product categories, industries, and countries. The consensus range for most consumer goods is $\sigma \in [0.5, 2]$ (weak-to-moderate substitution); differentiated products (autos, electronics, brand-name consumer goods) sit near $\sigma \in [1, 2]$; homogeneous commodities (basic food, energy) sit near $\sigma \in [0.5, 1]$; complements (matched pairs, systems of goods) can have $\sigma < 0.5$. The extremes of the CES spectrum ($\sigma = 0$ Leontief, $\sigma \to \infty$ perfect substitutes) are almost never a good fit to real data at the individual-commodity level.

Non-homothetic Engel curves: recent evidence. Blundell, Chen, and Kristensen (2007), using the U.K. Family Expenditure Survey, estimated Engel curves for eight consumption categories with rich non-parametric specifications. Food, alcohol, and clothing all show significantly non-linear log-linear Engel curves. This confirms that homotheticity fails not just as a linearity test but also as a smooth-non-linearity test: the data ask for the quadratic AIDS (QAIDS) or a similarly flexible specification. Aggregate consumption theory built on homothetic representative-agent utility loses in fit compared with non-homothetic models across every dataset examined.

Table 5.1 summarizes.

Structure	**Reference**	**Finding**
Homotheticity	Deaton-Muellbauer (1980); AIDS	Rejected at aggregate: food $\eta \approx 0.5$, luxury $\eta \approx 1.5$
Quasi-linearity	Andreoni-Miller (2002)	Fits laboratory choice with small stakes; fails large-stake generalizations
Additive separability across broad groups	Deaton (1974)	Not rejected: food/housing/transport blocks are separable
Additive separability within groups	Blundell (1988)	Rejected: within-food substitution is richer than additive
CES/Cobb-Douglas fit	Berry-Levinsohn-Pakes (1995); Nevo (2001)	Random-coefficients CES fits aggregate demand data very well

Interpretation. The three special structures are useful approximations, not universal truths. Aggregate consumer demand requires non-homothetic specifications (log-linear Engel curves, AIDS, translog). Partial-equilibrium consumer-surplus analysis uses quasi-linear demand and is accurate in small-stake environments. Additive separability holds across broad expenditure groups and fails within narrower groups where substitution and complementarity dominate. Applied work picks the structure that fits the question at hand.

5.7 Summary

- A preference relation is **homothetic** if $x \succsim y \Leftrightarrow \alpha x \succsim \alpha y$ for every $\alpha > 0$. Equivalent to admitting a utility representation homogeneous of degree 1. Cobb-Douglas, CES, and Leontief are homothetic; quasi-linear is not.
- A preference relation is **quasi-linear** in a distinguished good if some representation has the form $u(x) = x_1 + v(x_2, \dots, x_L)$. Demand for non-numeraire goods is independent of income; the numeraire absorbs all income effects.
- A utility function is **additively separable** if $u(x) = \sum_\ell u_\ell(x_\ell)$. **Debreu's theorem** (1960) characterizes the preference relations that admit such a representation via the separability condition.
- The four canonical families cross-cut the three properties. Cobb-Douglas and CES are homothetic and (after transformation) additively separable; quasi-linear is separable but not homothetic; Leontief is homothetic but not separable.
- Empirically: homotheticity fails at the aggregate; quasi-linearity works locally in small-stake settings; additive separability holds across broad groups and fails within narrower ones.

5.8 Guided exercise

Exercise (Guided). Show that the CES utility function

$$u(x) = \big(\alpha_1 x_1^\rho + \alpha_2 x_2^\rho\big)^{1/\rho}, \quad \rho \in (0,1),\ \alpha_1 + \alpha_2 = 1,$$

is (a) homothetic and (b) admits an additively separable representation, and (c) compute its elasticity of substitution $\sigma = \dfrac{1}{1-\rho}$.

Solution.

(a) Homothetic. Compute $u(tx)$ for $t > 0$:

$$u(tx) = \big(\alpha_1 (tx_1)^\rho + \alpha_2 (tx_2)^\rho\big)^{1/\rho} = \big(t^\rho(\alpha_1 x_1^\rho + \alpha_2 x_2^\rho)\big)^{1/\rho} = t \cdot u(x).$$

So u is homogeneous of degree 1; by Proposition 5.1, $\succsim$ is homothetic.

(b) Additively separable. Raising to the ρ-th power,

$$u(x)^\rho = \alpha_1 x_1^\rho + \alpha_2 x_2^\rho.$$

The transformed utility $\tilde{u} = u^\rho$ is additively separable (a sum of one function of each coordinate). Since $t \mapsto t^\rho$ is strictly increasing on $\mathbb{R}_+$ for $\rho > 0$, $\tilde{u}$ represents the same $\succsim$ as u. Hence $\succsim$ admits an additively separable representation.

(c) Elasticity of substitution. The MRS of u is

$$\mathrm{MRS}_{12}(x) = \frac{\alpha_1}{\alpha_2} \cdot \left(\frac{x_1}{x_2}\right)^{\rho-1} = \frac{\alpha_1}{\alpha_2} \cdot \left(\frac{x_2}{x_1}\right)^{1-\rho}.$$

The elasticity of substitution is

$$\sigma = \frac{d\log(x_2/x_1)}{d\log \mathrm{MRS}_{12}}.$$

Taking logs of the MRS formula, $\log \mathrm{MRS}_{12} = \log(\alpha_1/\alpha_2) + (1-\rho)\log(x_2/x_1)$, so $d\log \mathrm{MRS}_{12}/d\log(x_2/x_1) = 1-\rho$, and $\sigma = \dfrac{1}{1-\rho}$.

Lesson. The CES family has one parameter, ρ, that controls the entire cross-good substitution behavior: $\rho \to 1$ gives perfect substitutes ($\sigma \to \infty$), $\rho = 0$ gives Cobb-Douglas ($\sigma = 1$), $\rho \to -\infty$ gives Leontief ($\sigma \to 0$). Applied demand estimation typically lets σ (or equivalently ρ) be a free parameter and estimates it from the data.

Chapter 6 - Utility maximization and Marshallian demand

6.1 Introduction

Chapters 1–5 developed the theoretical apparatus. This chapter puts it to work. Given prices $p \in \mathbb{R}^L_{++}$ and wealth $w > 0$, the **utility-maximization problem** (UMP) is

$$\max_{x \in \mathbb{R}^L_+} u(x) \quad \text{subject to} \quad p \cdot x \leq w.$$

The solution $x(p, w) \in \mathbb{R}^L_+$ is the **Marshallian demand** at (p, w). The maximum value $v(p, w) = u(x(p, w))$ is the **indirect utility function**.

The chapter builds four things. Section 6.2 states existence (when a solution exists) and uniqueness (when it is a singleton). Section 6.3 develops the interior-solution tangency condition and the corner conditions. Section 6.4 introduces the indirect utility function and its basic properties. Section 6.5 works through the Cobb-Douglas closed form and other explicit demand systems. Section 6.6 develops the comparative statics of Marshallian demand – how demand shifts with p and w – and states Roy's identity. Section 6.7 presents laboratory evidence on the demand-system fit of the Cobb-Douglas, CES, and quasi-linear specifications. Section 6.8 summarizes.

Throughout, we assume u is continuous, strongly monotone (so the budget constraint binds; Section 2.4), and strictly quasi-concave (so the solution is unique; Section 2.5) on the consumption set $X = \mathbb{R}^L_+$.

6.2 The utility-maximization problem

Fix prices $p \in \mathbb{R}^L_{++}$ and wealth $w > 0$. The **Walrasian budget set** $B_{p,w} = \{x \in \mathbb{R}^L_+ : p \cdot x \leq w\}$ is a compact subset of $\mathbb{R}^L$ (closed and bounded because $p \gg 0$ forces $x_\ell \leq \dfrac{w}{p_\ell}$ for each ℓ).

Proposition 6.1 (Existence). If u is continuous on $B_{p,w}$, then the utility-maximization problem $\max u$ subject to $x \in B_{p,w}$ has a solution.

Proof. $B_{p,w}$ is compact (closed and bounded); u is continuous; Weierstrass's theorem (see MWG (1995), Section M.F of the Mathematical Appendix, on continuous functions and compact sets) gives the existence of a maximum. ■

Proposition 6.2 (Uniqueness). If u is strictly quasi-concave on $\mathbb{R}^L_+$, then the solution to the UMP is a singleton.

Proof. Suppose x^* and y^* are two different solutions with $u(x^*) = u(y^*) = v(p, w)$. The mixture $z = (x^* + y^*)/2$ is in $B_{p,w}$ (budget sets are convex). Strict quasi-concavity of u gives $u(z) > \min\{u(x^*), u(y^*)\} = v(p, w)$, contradicting the optimality of x^*. ■

Proposition 6.3 (Budget binds). If u is locally non-satiated (Section 2.4) and x^* solves the UMP at (p, w), then $p \cdot x^* = w$.

Proof. Proposition 2.2 (Chapter 2, Section 2.4) established the result. Briefly: if $p \cdot x^* < w$, take $\varepsilon = (w - p \cdot x^*)/(2\|p\|) > 0$; the ball $B(x^*, \varepsilon)$ is contained in $B_{p,w}$; local non-satiation delivers $y \in B(x^*, \varepsilon)$ with $y \succ x^*$, contradicting optimality of x^*. ■

Under strong monotonicity (which implies local non-satiation), Marshallian demand always exhausts the budget. This is the first-order restriction the theorist uses to solve the UMP.

Intuitively, the three propositions of this section together certify that the UMP is well-posed: Weierstrass guarantees a solution exists; strict quasi-concavity makes it a single point (not a set of tied optima); local non-satiation forces the solution to exhaust the whole budget (no wasted purchasing power). Chapter 5 verified that the four canonical families of Chapter 2 – Cobb-Douglas, CES, quasi-linear, Leontief – all satisfy these three conditions on the interior of $\mathbb{R}^L_+$, so the entire apparatus of this chapter applies to each of them without qualification.

6.3 First-order conditions and the tangency condition

Suppose u is C^1, strongly monotone, strictly quasi-concave. Under those hypotheses, the solution x^* to the UMP is either interior (all $x^*_\ell > 0$) or on the boundary of $\mathbb{R}^L_+$ (some $x^*_\ell = 0$).

Interior solutions. At an interior x^*, the Lagrangian $\mathcal{L} = u(x) - \lambda(p \cdot x - w)$ has first-order conditions

$$\frac{\partial u(x^*)}{\partial x_\ell} = \lambda\, p_\ell \quad \text{for every } \ell = 1, \ldots, L.$$

Rearranging pairwise for ℓ and k,

$$\mathrm{MRS}_{\ell k}(x^*) = \frac{\partial u(x^*)/\partial x_\ell}{\partial u(x^*)/\partial x_k} = \frac{p_\ell}{p_k}.$$

The tangency condition: at the optimum, the marginal rate of substitution between any pair of goods equals the price ratio. Combined with the budget-binds condition $p \cdot x^* = w$, the tangency condition gives L equations in $L + 1$ unknowns (the demands $x^*_1, \ldots, x^*_L$ and the multiplier λ), which typically solve for a unique interior optimum. Figure 6.1 shows the interior optimum geometrically.

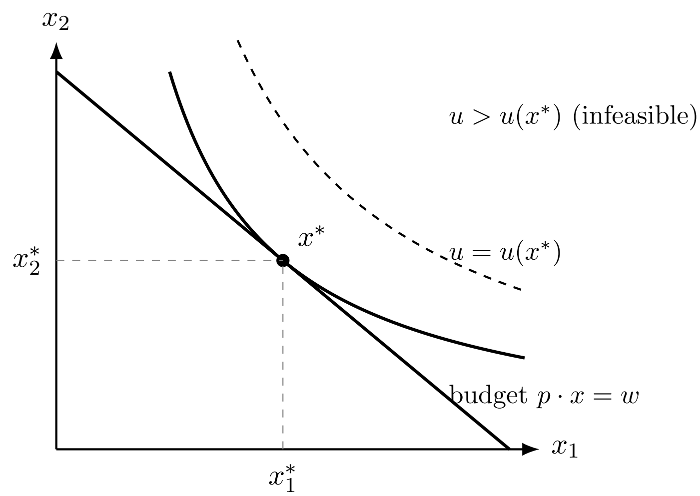

Figure 6.1: The utility-maximization problem. The consumer picks the point on the budget line $p \cdot x = w$ that reaches the highest indifference curve. At the optimum x^* the tangency condition holds: the indifference curve is tangent to the budget line, so $\mathrm{MRS}_{12}(x^*) = p_1/p_2$. Any curve strictly above passes through infeasible bundles only.

Corner solutions. If u is strongly monotone but the MRS goes to zero or infinity as $x_\ell \to 0$, the optimum may be at a corner. Two representative cases:

- **Cobb-Douglas.** MRS goes to infinity at the axes; the optimum is always interior. No corner conditions arise.
- **Linear (perfect substitutes).** MRS is constant; if the constant MRS differs from the price ratio, the optimum is at a corner (spend everything on the good with the larger MRS).

Kuhn-Tucker for corners. The full Kuhn-Tucker (KKT) conditions for the utility-maximization problem with non-negativity constraints

$$\max_{x \geq 0} u(x) \quad \text{subject to} \quad p \cdot x \leq w$$

combine the standard Lagrangian first-order conditions with two additional families: primal feasibility of the non-negativity constraints, dual feasibility of their multipliers, and complementary slackness.

Let $\mu_\ell \geq 0$ be the KKT multiplier attached to the non-negativity constraint $x_\ell \geq 0$, and

$\lambda \geq 0$ the multiplier on the budget constraint. The Lagrangian is

$$\mathcal{L}(x, \lambda, \mu; p, w) \;=\; u(x) - \lambda\left(p \cdot x - w\right) + \sum_{\ell=1}^{L} \mu_\ell \, x_\ell.$$

A necessary condition for (x^*, λ^*, μ^*) to solve the problem (assuming a constraint qualification like Slater's) is the full KKT system:

$$\begin{array}{ll}
\text{(stationarity)} & \dfrac{\partial u(x^*)}{\partial x_\ell} - \lambda^* p_\ell + \mu_\ell^* = 0 \quad \text{for } \ell = 1, \ldots, L; \\
\text{(primal feasibility)} & p \cdot x^* \leq w, \; x_\ell^* \geq 0 \text{ for every } \ell; \\
\text{(dual feasibility)} & \lambda^* \geq 0, \; \mu_\ell^* \geq 0 \text{ for every } \ell; \\
\text{(complementary slackness)} & \lambda^* \left(p \cdot x^* - w\right) = 0, \; \mu_\ell^* \, x_\ell^* = 0 \text{ for every } \ell.
\end{array}$$

Two equivalent readings of stationarity plus complementary slackness on the non-negativity constraint:

$$x_\ell^* > 0 \implies \mu_\ell^* = 0 \implies \frac{\partial u(x^*)}{\partial x_\ell} = \lambda^* p_\ell; \qquad x_\ell^* = 0 \implies \frac{\partial u(x^*)}{\partial x_\ell} \leq \lambda^* p_\ell.$$

Interior goods obey the usual tangency; boundary goods obey the weak inequality (marginal utility at the boundary does not exceed the value of the price).

Example 6.1. A worked Kuhn-Tucker corner solution. Take $u(x_1, x_2) = x_1 + 2\sqrt{x_2}$, prices $p_1 = 1$, $p_2 = 1$, and wealth $w = 3$. This is a quasi-linear preference in good 1; the FOC for interior x_2 is

$$\frac{\partial u}{\partial x_2} = \frac{1}{\sqrt{x_2}} = \lambda p_2 = \lambda \quad \implies \quad x_2 = \frac{1}{\lambda^2}.$$

Simultaneously, $\dfrac{\partial u}{\partial x_1} = 1 = \lambda p_1 = \lambda$, so $\lambda = 1$ and $x_2^* = 1$. From the budget, $x_1^* = w - p_2 x_2^* = 3 - 1 = 2$. Both goods interior; KKT multipliers $\mu_1 = \mu_2 = 0$.

Now shrink w to 0.5: try the interior solution. It gives $x_2^* = 1$ and $x_1^* = 0.5 - 1 = -0.5 < 0$, infeasible. The constraint $x_1 \geq 0$ binds; try the corner with $x_1^* = 0$. Then $x_2^* = w/p_2 = 0.5$, and the FOCs become

$$\frac{\partial u}{\partial x_1} = 1 \leq \lambda p_1 = \lambda, \qquad \frac{\partial u}{\partial x_2} = \frac{1}{\sqrt{0.5}} \approx 1.414 = \lambda.$$

So $\lambda \approx 1.414$, which satisfies $1 \leq 1.414$; the KKT multiplier on the non-negativity of good 1 is $\mu_1 = \lambda - 1 \approx 0.414 > 0$. Complementary slackness $\mu_1 x_1^* = 0.414 \cdot 0 = 0$ is satisfied. Optimum: $x^* = (0, 0.5)$.

Lesson: the KKT conditions handle the corner cleanly. When the interior tangency would push $x_\ell < 0$, the constraint $x_\ell \geq 0$ binds; the corresponding multiplier μ_ℓ becomes strictly positive, and the FOC on that coordinate holds only as a weak inequality. Figure 6.2 illustrates the perfect-substitutes case, where the optimum sits at a corner whenever the constant MRS differs from the price ratio.

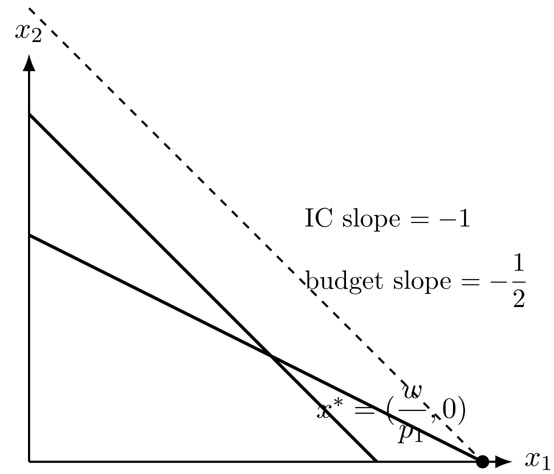

Figure 6.2: Corner solution for $u(x) = x_1 + x_2$ when $p_1 < p_2$. The indifference curves are parallel lines with constant slope -1; the budget line is flatter (slope $-\frac{1}{2}$). The optimum lies at the corner $x^* = (\frac{w}{p_1}, 0)$: the consumer spends her entire wealth on the cheaper good.

6.4 The indirect utility function

The maximum value of u over the budget set is a function of (p, w):

$$v(p, w) = \max_{x \in B_{p,w}} u(x) = u(x(p, w)).$$

Proposition 6.4. Under the standard hypotheses (continuous u, strongly monotone), the indirect utility function has the following properties:

1. Homogeneous of degree 0 in (p, w): $v(\alpha p, \alpha w) = v(p, w)$ for every $\alpha > 0$.
2. Non-increasing in each price p_ℓ.
3. Non-decreasing in wealth w (strictly increasing if u is strongly monotone).
4. Continuous in (p, w).

5. Quasi-convex in (p, w): the set $\{(p, w) : v(p, w) \le \bar{v}\}$ is convex for every $\bar{v}$.

Proof.
(1) Homogeneity of degree zero in (p, w). Fix $\alpha > 0$. The budget set at $(\alpha p, \alpha w)$ is

$$B_{\alpha p,\alpha w} \;=\; \{x \in \mathbb{R}^L_+ : \alpha p \cdot x \le \alpha w\} \;=\; \{x : p \cdot x \le w\} \;=\; B_{p,w},$$

since $\alpha > 0$ can be divided out of the constraint. Because the objective $u(x)$ does not depend on (p, w) at all, the value of the maximization is unchanged:

$$v(\alpha p, \alpha w) \;=\; \max_{x \in B_{\alpha p,\alpha w}} u(x) \;=\; \max_{x \in B_{p,w}} u(x) \;=\; v(p, w).$$

(2) Non-increasing in each p_ℓ. Fix all prices p_k $(k \ne \ell)$ and the wealth w; suppose $p'_\ell > p_\ell$. For any $x \in B_{p',w}$, the constraint $\sum_k p'_k x_k \le w$ combined with $x_k \ge 0$ and $p'_\ell > p_\ell$ gives

$$\sum_k p_k x_k = \sum_{k \ne \ell} p_k x_k + p_\ell x_\ell \;\le\; \sum_{k \ne \ell} p'_k x_k + p'_\ell x_\ell - \big(p'_\ell - p_\ell\big) x_\ell \;\le\; w,$$

so $x \in B_{p,w}$. Hence $B_{p',w} \subseteq B_{p,w}$: raising the price of good ℓ shrinks the budget set. The maximum of a function over a smaller set is no greater than over a larger set, so $v(p', w) \le v(p, w)$.
(3) Non-decreasing in w. Fix p and let $w' > w$. For any $x \in B_{p,w}$, $p \cdot x \le w < w'$, so $x \in B_{p,w'}$; hence $B_{p,w} \subseteq B_{p,w'}$. The maximum on the larger set is at least the maximum on the smaller: $v(p, w') \ge v(p, w)$. For *strict* monotonicity when u is strongly monotone: pick any interior solution $x^*(p, w)$ and consider the shifted bundle $\tilde{x} = x^*(p, w) + \varepsilon e_\ell$ for some ℓ and small $\varepsilon > 0$ small enough that $\tilde{x} \in B_{p,w'}$. Strong monotonicity gives $u(\tilde{x}) > u(x^*(p, w)) = v(p, w)$, so $v(p, w') \ge u(\tilde{x}) > v(p, w)$.
(4) Continuity in (p, w). The budget set correspondence $B(p, w) = \{x : p \cdot x \le w, x \ge 0\}$ is continuous (upper and lower hemi-continuous) at every (p, w) with $p \gg 0$ (see MWG (1995), Section M.H of the Mathematical Appendix, on correspondences). The objective u is continuous. By the Berge maximum theorem (MWG (1995), Section M.L of the Mathematical Appendix), the value function $v(p, w) = \max_{x \in B(p,w)} u(x)$ is continuous in (p, w).
(5) Quasi-convexity in (p, w). We show that for every $\bar{v} \in \mathbb{R}$, the set $S(\bar{v}) = \{(p, w) : v(p, w) \le \bar{v}\}$ is convex. Take $(p^1, w^1), (p^2, w^2) \in S(\bar{v})$ and $t \in [0, 1]$; set $(p^t, w^t) = t(p^1, w^1) + (1-t)(p^2, w^2)$. Let x^t solve the UMP at (p^t, w^t), so $v(p^t, w^t) = u(x^t)$ with $p^t \cdot x^t \le w^t$. Because $p^t \cdot x^t \le w^t = tw^1 + (1-t)w^2$, at least one of $p^1 \cdot x^t \le w^1$ or $p^2 \cdot x^t \le w^2$ must hold (otherwise $t\,p^1 \cdot x^t + (1-t)p^2 \cdot x^t > tw^1 + (1-t)w^2 = w^t$, contradiction). Suppose the first: $p^1 \cdot x^t \le w^1$, so $x^t \in B_{p^1,w^1}$, giving $u(x^t) \le v(p^1, w^1) \le \bar{v}$. Symmetric for the second case. Either way, $v(p^t, w^t) = u(x^t) \le \bar{v}$, so $(p^t, w^t) \in S(\bar{v})$. ■

Intuitively, all five properties can be read off the primitive: the value function inherits its shape from the underlying maximization. Homogeneity is a rescaling of the budget line, monotonicity in prices is the shrinking of the affordable set as any good becomes more expensive, monotonicity in wealth is the reverse enlargement, continuity is Berge applied to a continuous budget correspondence, and quasi-convexity says that any convex combination of "price-wealth pairs that produce low welfare" also produces low welfare – because at the combined pair, the consumer can afford the same bundle as at one of the two original pairs.

The five properties, especially homogeneity of degree 0 and quasi-convexity, are the standard test that a candidate indirect utility function is well-specified. The tighter duality with the expenditure function is developed in MWG Chapter 3 and in my *Advanced Microeconomic Theory* (MIT Press, 2017), Chapter 2.

Example 6.2. A numerical check of the five properties. Take the Cobb-Douglas indirect utility for $\alpha = 1/2$:

$$v(p, w) \;=\; w \cdot p_1^{-1/2} p_2^{-1/2} \cdot 2^{-1/2} \cdot 2^{-1/2} \;=\; \frac{w}{2\sqrt{p_1 p_2}}.$$

Evaluate the five properties.

- *(1) Homogeneity of degree zero.* $v(\alpha p, \alpha w) = \dfrac{\alpha w}{2\sqrt{\alpha p_1 \cdot \alpha p_2}} = \dfrac{\alpha w}{2\alpha\sqrt{p_1 p_2}} = \dfrac{w}{2\sqrt{p_1 p_2}} = v(p, w)$. ✓
- *(2) Non-increasing in p_ℓ.* $\dfrac{\partial v}{\partial p_1} = -\dfrac{w}{4} p_1^{-3/2} p_2^{-1/2} < 0$. ✓
- *(3) Non-decreasing in w.* $\dfrac{\partial v}{\partial w} = \dfrac{1}{2\sqrt{p_1 p_2}} > 0$. ✓
- *(4) Continuity.* v is a rational function of (p_1, p_2, w) that is C^∞ on $\{(p, w) : p \gg 0\}$; continuity is immediate. ✓
- *(5) Quasi-convexity.* The lower-level set $\{(p, w) : v(p, w) \leq \bar{v}\} = \{(p, w) : w \leq 2\bar{v}\sqrt{p_1 p_2}\}$ is convex because $\sqrt{p_1 p_2}$ is concave in (p_1, p_2) (its Hessian is $-\frac{1}{4}p^{-3/2}$ times a rank-2 correction; direct calculation confirms concavity), and the set $\{(p, w) : w - 2\bar{v}\sqrt{p_1 p_2} \leq 0\}$ is a sublevel set of a convex function of (p, w). ✓

The Cobb-Douglas indirect utility passes all five tests, as any well-posed indirect utility must.

6.5 Explicit demand systems

The three canonical utility families of Chapter 2 all give closed-form Marshallian demand. This section works through the derivations.

Cobb-Douglas. With $u(x) = x_1^{\alpha} x_2^{1-\alpha}$ and $\alpha \in (0,1)$, the tangency condition is

$$\mathrm{MRS}_{12}(x) = \frac{\alpha\, x_2}{(1-\alpha)\, x_1} = \frac{p_1}{p_2}.$$

Solving for x_2: $x_2 = \dfrac{(1-\alpha)\, p_1\, x_1}{\alpha\, p_2}$. Substituting into the budget constraint $p_1x_1 + p_2x_2 = w$:

$$p_1x_1 + p_2 \cdot \frac{(1-\alpha)p_1x_1}{\alpha p_2} = w \quad \Rightarrow \quad p_1x_1\Big(1 + \frac{1-\alpha}{\alpha}\Big) = w \quad \Rightarrow \quad p_1x_1 \cdot \frac{1}{\alpha} = w.$$

So $x_1^*(p, w) = \dfrac{\alpha\, w}{p_1}$, and by symmetry $x_2^*(p, w) = \dfrac{(1-\alpha)\, w}{p_2}$.
Reading: the consumer spends fraction α of income on good 1 and fraction $1 - \alpha$ on good 2, independent of prices. This is the workhorse budget-share rule of the Cobb-Douglas demand system.

Intuitively, the exponent α in the Cobb-Douglas $u(x) = x_1^{\alpha} x_2^{1-\alpha}$ is exactly the “taste weight” the consumer places on good 1: the log utility is $\alpha \log x_1 + (1-\alpha) \log x_2$, a weighted average of log consumptions with weights α and $1 - \alpha$. At the optimum, the marginal utility of good 1 relative to good 2 equals the price ratio; because Cobb-Douglas marginal utility scales as $1/x$ times the weight, the ratio $\alpha/(1-\alpha) \cdot x_2/x_1 = p_1/p_2$ implies $x_1/x_2 \propto \alpha p_2/((1-\alpha)p_1)$. Multiplying by prices yields the budget-share result: $p_1x_1^*/w = \alpha$, exactly the exponent. Prices cancel in the share expression because the substitution and income effects of a price change offset one-to-one for Cobb-Douglas (elasticity of substitution equal to one).

Numerical illustration. With $\alpha = 0.4$, $p_1 = 2$, $p_2 = 3$, $w = 60$: $x_1^* = 0.4 \cdot 60/2 = 12$ and $x_2^* = 0.6 \cdot 60/3 = 12$. Spending on good 1: $p_1x_1^* = 24 = 0.4 \cdot 60$; spending on good 2: $p_2x_2^* = 36 = 0.6 \cdot 60$. Now double p_1 to 4: $x_1^* = 0.4 \cdot 60/4 = 6$, $x_2^* = 0.6 \cdot 60/3 = 12$; spending shares remain $(0.4, 0.6)$. The consumer buys half as much of good 1 and the same amount of good 2, keeping the share exactly at 0.4.
Figure 6.3 plots the resulting Marshallian demand curve for good 1 as p_1 varies.

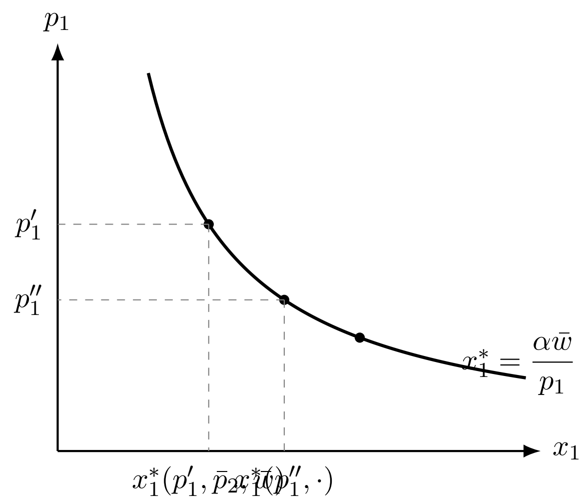

Figure 6.3: Marshallian demand for good 1 as a function of its own price, holding p_2 and w fixed. For a Cobb-Douglas consumer, $x_1^*(p_1, \bar{p}_2, \bar{w}) = \alpha \dfrac{\bar{w}}{p_1}$ is a rectangular hyperbola: strictly decreasing in p_1, with unit price elasticity.

CES. With $u(x) = (\alpha_1 x_1^\rho + \alpha_2 x_2^\rho)^{1/\rho}$ and $\rho \in (-\infty, 1) \setminus \{0\}$, a similar computation gives

$$x_\ell^*(p, w) = \frac{\alpha_\ell^\sigma \, p_\ell^{-\sigma} \, w}{\sum_k \alpha_k^\sigma \, p_k^{1-\sigma}}, \qquad \sigma = \frac{1}{1-\rho}.$$

Reading: the budget share on good ℓ is $s_\ell = p_\ell x_\ell^*/w = \alpha_\ell^\sigma p_\ell^{1-\sigma} / \sum_k \alpha_k^\sigma p_k^{1-\sigma}$. Unlike Cobb-Douglas, CES budget shares depend on prices unless $\sigma = 1$ (which recovers Cobb-Douglas as a special case).

Step-by-step derivation. Consider $u(x) = (\alpha_1 x_1^\rho + \alpha_2 x_2^\rho)^{1/\rho}$ with $L = 2$. The tangency condition equates the MRS to the price ratio:

$$\mathrm{MRS}_{12}(x) \;=\; \frac{\partial u/\partial x_1}{\partial u/\partial x_2} \;=\; \frac{\alpha_1 x_1^{\rho-1}}{\alpha_2 x_2^{\rho-1}} \;=\; \frac{p_1}{p_2}.$$

Solving for x_2/x_1:

$$\frac{x_2}{x_1} = \left(\frac{\alpha_2 p_1}{\alpha_1 p_2}\right)^{1/(1-\rho)} = \left(\frac{\alpha_2}{\alpha_1}\right)^\sigma \left(\frac{p_1}{p_2}\right)^\sigma,$$

where $\sigma = 1/(1-\rho)$. Plugging into the budget $p_1 x_1 + p_2 x_2 = w$ and grouping:

$$x_1 \left(p_1 + p_2 \left(\frac{\alpha_2}{\alpha_1}\right)^\sigma \left(\frac{p_1}{p_2}\right)^\sigma\right) = w.$$

Factor p_1^σ and reorganize:

$$x_1\, p_1^\sigma\big(\alpha_1^\sigma p_1^{1-\sigma} + \alpha_2^\sigma p_2^{1-\sigma}\big) = w\,\alpha_1^\sigma \quad \Longrightarrow \quad x_1^* = \frac{\alpha_1^\sigma p_1^{-\sigma} w}{\alpha_1^\sigma p_1^{1-\sigma} + \alpha_2^\sigma p_2^{1-\sigma}},$$

which is the general formula.

Numerical illustration. With $\alpha_1 = \alpha_2 = 1$, $\rho = -1$ (so $\sigma = 1/2$), $p = (1, 4)$, $w = 10$: $p_1^{1-\sigma} = 1^{1/2} = 1$, $p_2^{1-\sigma} = 4^{1/2} = 2$; $x_1^* = \dfrac{1 \cdot 1 \cdot 10}{1+2} = \dfrac{10}{3}$; symmetrically $x_2^* = \dfrac{1 \cdot 4^{-1/2} \cdot 10}{3} = \dfrac{10/2}{3} = \dfrac{5}{3}$. Spending: $p_1 x_1^* = 10/3 \approx 3.33$; $p_2 x_2^* = 4 \cdot 5/3 = 20/3 \approx 6.67$; total 10, matching w. Share on good 1: $s_1 = 1/3$; on good 2: $s_2 = 2/3$. The share on the more-expensive good rises with its price because $\sigma < 1$ (goods are complements in the CES sense).

Leontief. With $u(x) = \min\{x_1/a_1, x_2/a_2\}$ and $a_1, a_2 > 0$, the optimum has $x_1/a_1 = x_2/a_2$ (perfect complements: no reason to have extra of one). Denote this common ratio t. Budget: $p_1(a_1 t) + p_2(a_2 t) = w$, so $t = w/(a_1 p_1 + a_2 p_2)$, and

$$x_\ell^*(p, w) = \frac{a_\ell\, w}{a_1 p_1 + a_2 p_2}.$$

Numerical illustration. Left shoes and right shoes: $u(x) = \min\{x_1, x_2\}$ (so $a_1 = a_2 = 1$). Prices $p_1 = p_2 = 5$, wealth $w = 40$. Optimum: $x_1^* = x_2^* = w/(p_1 + p_2) = 40/10 = 4$ pairs of each. Spending equally on both. Now raise p_1 to 10: $x_1^* = x_2^* = 40/(10 + 5) = 40/15 \approx 2.67$. Both demands fell one-for-one – perfect complements react in tandem to a change in either price. The spending share of good 1 shifts from 50% to $10{\cdot}2.67/40 \approx 66.7\%$ (larger, absorbing the price increase).

Intuitively, Leontief demand has zero substitution: goods are used in fixed proportions (one left shoe per right shoe), so the consumer holds their ratio at the "recipe" $a_1 : a_2$ regardless of relative prices. The elasticity of substitution is $\sigma = 0$, at the opposite extreme from perfect substitutes ($\sigma = \infty$).

Quasi-linear. With $u(x_1, x_2) = v(x_1) + x_2$ and v strictly increasing and strictly concave (Chapter 5), the tangency condition is $v'(x_1) = p_1/p_2$: demand for good 1 depends only on p_1/p_2, independent of w. Setting the price of good 2 to 1 as numeraire, $x_1^* = (v')^{-1}(p_1)$ and $x_2^* = w - p_1\, x_1^*$.

The four demand systems span the full range of substitution elasticities: Cobb-Douglas ($\sigma = 1$), CES (any σ), Leontief ($\sigma = 0$), and linear/quasi-linear (interior non-crossing behavior). Chapter 5's structural taxonomy translates directly into the demand-system predictions here.

Example 6.4. CES demand at $\sigma = 2$, $(\alpha_1, \alpha_2) = (0.4, 0.6)$. Take the CES utility

$$u(x_1, x_2) = \left(0.4\, x_1^{\rho} + 0.6\, x_2^{\rho}\right)^{1/\rho}$$

with ρ chosen so that $\sigma = \dfrac{1}{1-\rho} = 2$, i.e., $\rho = 1/2$. Prices $(p_1, p_2) = (1, 2)$; wealth $w = 20$.
Step 1 (price aggregator). With $1 - \sigma = -1$, the CES demand formula gives

$$x_\ell^*(p, w) = \frac{\alpha_\ell^{\sigma}\, p_\ell^{-\sigma}\, w}{\sum_k \alpha_k^{\sigma}\, p_k^{1-\sigma}} = \frac{\alpha_\ell^{2}\, p_\ell^{-2}\, w}{\alpha_1^2\, p_1^{-1} + \alpha_2^2\, p_2^{-1}}.$$

Substituting the numbers: $\alpha_1^2 = 0.16$, $\alpha_2^2 = 0.36$, $p_1^{-1} = 1$, $p_2^{-1} = 0.5$; denominator $= 0.16 \cdot 1 + 0.36 \cdot 0.5 = 0.16 + 0.18 = 0.34$.
Step 2 (demands).

$$x_1^* = \frac{0.16 \cdot 1 \cdot 20}{0.34} = \frac{3.2}{0.34} \approx 9.41, \qquad x_2^* = \frac{0.36 \cdot 0.25 \cdot 20}{0.34} = \frac{1.8}{0.34} \approx 5.29.$$

Spending check: $p_1 x_1^* + p_2 x_2^* = 1 \cdot 9.41 + 2 \cdot 5.29 = 9.41 + 10.59 = 20$. Matches w.
Step 3 (budget shares). Share on good 1: $s_1 = 9.41/20 \approx 0.47$; share on good 2: $s_2 = 10.59/20 \approx 0.53$.

Lesson: with $\sigma = 2 > 1$ (goods are substitutes in the CES sense), a price increase in good 2 (from 1 to 2 relative to good 1) shifts the budget share *away* from the more expensive good; the share on good 2 is 0.53 despite good 2 having the larger taste weight $\alpha_2 = 0.6$. Under $\sigma < 1$ (complements) the pattern reverses, as in the $\sigma = 1/2$ example above.

6.6 Comparative statics and Roy's identity

Two comparative-statics exercises are central to the theory: how Marshallian demand responds to a price change, and how it responds to a wealth change.

Wealth effect. At an interior optimum, differentiating the budget constraint $p \cdot x(p, w) = w$ with respect to w gives

$$\sum_\ell p_\ell \frac{\partial x_\ell(p, w)}{\partial w} = 1.$$

The weighted-by-prices sum of wealth derivatives is 1: an extra dollar of income is spent on some combination of goods, and the total spending must equal the extra dollar. The individual wealth derivative $\partial x_\ell / \partial w$ can be positive (normal good), negative (inferior good), or zero (borderline). Figure 6.4 plots the income-expansion path for homothetic (Cobb-Douglas) preferences: three parallel budget lines and a single ray through their tangency points.

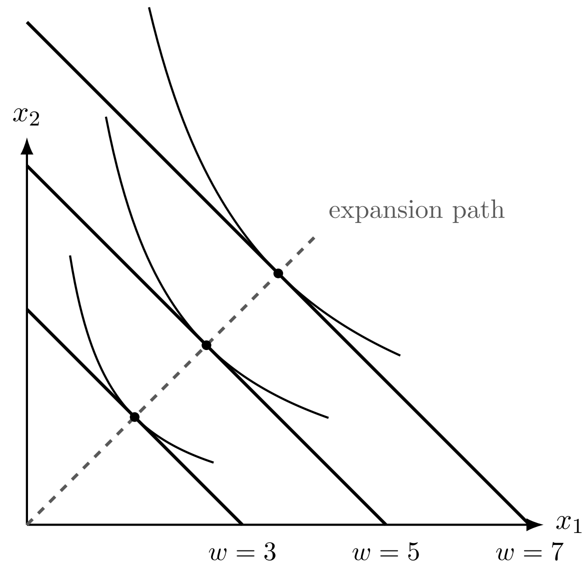

Figure 6.4: Three parallel budget lines (same relative prices, three wealth levels) and their tangency bundles under Cobb-Douglas preferences. The expansion path is a ray through the origin, a signature of homotheticity: every good has unit wealth elasticity.

Own-price effect and the Slutsky equation (preview). Differentiating the tangency condition MRS $= p_1/p_2$ with respect to p_1 (holding w and p_2 fixed) gives Marshallian demand's response to a price change. The Slutsky-Hicks decomposition writes

$$\frac{\partial x_\ell(p,w)}{\partial p_k} = \frac{\partial h_\ell(p,\bar{u})}{\partial p_k} - x_k(p,w) \cdot \frac{\partial x_\ell(p,w)}{\partial w},$$

where h_ℓ is Hicksian demand (the expenditure-minimizing counterpart to Marshallian demand). The first term is the "pure substitution effect," the second the "income effect." A full development is beyond the present chapter; see MWG Chapter 3 and *Advanced Microeconomic Theory* (MIT Press, 2017), Chapter 2.

Numerical Slutsky check on Cobb-Douglas. Take $u(x) = x_1^{1/2} x_2^{1/2}$ with prices $p = (1,1)$ and wealth $w = 4$. Marshallian demand: $x_1^* = w/(2p_1) = 2$ and $x_2^* = w/(2p_2) = 2$, giving baseline utility $u^* = \sqrt{2 \cdot 2} = 2$. Direct partial: $\dfrac{\partial x_1}{\partial p_1} = -\dfrac{w}{2p_1^2} = -2$. Wealth effect on good 1: $\dfrac{\partial x_1}{\partial w} = \dfrac{1}{2p_1} = 1/2$. Multiply by x_1: $x_1 \cdot \dfrac{\partial x_1}{\partial w} = 2 \cdot (1/2) = 1$. So the Slutsky decomposition gives

$$-2 = \frac{\partial h_1}{\partial p_1} - 1 \quad \Longrightarrow \quad \frac{\partial h_1}{\partial p_1} = -1.$$

The pure substitution effect is -1 (negative, as always for a normal good) and the income effect $-x_1 \cdot \partial x_1/\partial w = -1$. Total Marshallian effect $-2 = (-1) + (-1)$: substitution plus income, in equal shares for this symmetric Cobb-Douglas configuration.

Roy's identity. A more surprising comparative-statics identity connects the indirect utility function to Marshallian demand.

Theorem 6.1 (Roy's Identity). Under the standard hypotheses on u, at any (p, w) where v is differentiable,

$$x_\ell(p, w) = -\frac{\partial v(p, w)/\partial p_\ell}{\partial v(p, w)/\partial w}.$$

Proof. Write the utility-maximization problem in Lagrangian form,

$$\mathcal{L}(x, \lambda; p, w) \;=\; u(x) - \lambda \big(p \cdot x - w\big).$$

Under the standard hypotheses (continuous, strictly quasi-concave u, strong monotonicity), the interior optimum $x^*(p, w)$ satisfies the first-order conditions

$$\frac{\partial u(x^*)}{\partial x_k} \;=\; \lambda^* p_k \qquad \text{for every } k = 1, \dots, L, \tag{1}$$

together with the binding budget constraint $p \cdot x^*(p, w) = w$. The multiplier $\lambda^* = \lambda^*(p, w) > 0$ is the marginal utility of relaxing the budget constraint (of one extra dollar of wealth).

Step 1 (differentiate v with respect to p_ℓ). The indirect utility function is $v(p, w) = u(x^*(p, w))$. By the chain rule,

$$\frac{\partial v(p, w)}{\partial p_\ell} \;=\; \sum_{k=1}^{L} \frac{\partial u(x^*)}{\partial x_k} \cdot \frac{\partial x_k^*(p, w)}{\partial p_\ell}.$$

Substitute the first-order condition (1) into the first factor of each term:

$$\frac{\partial v}{\partial p_\ell} \;=\; \sum_k \lambda^* p_k \,\frac{\partial x_k^*}{\partial p_\ell} \;=\; \lambda^* \sum_k p_k \,\frac{\partial x_k^*}{\partial p_\ell}.$$

Step 2 (differentiate the budget constraint). Since $p \cdot x^*(p, w) = w$ for every (p, w), differentiating both sides with respect to p_ℓ (holding w and $\{p_k : k \neq \ell\}$ fixed) gives

$$x_\ell^*(p, w) \;+\; \sum_k p_k \,\frac{\partial x_k^*}{\partial p_\ell} \;=\; 0 \qquad \Longleftrightarrow \qquad \sum_k p_k \,\frac{\partial x_k^*}{\partial p_\ell} \;=\; -x_\ell^*.$$

The first term is the direct effect of the price change on the budget (holding demand fixed); the second is the indirect effect through the change in demand.

Step 3 (combine). Substituting the second identity into Step 1,

$$\frac{\partial v(p,w)}{\partial p_\ell} \;=\; \lambda^*(p,w)\cdot\big(-x^*_\ell(p,w)\big) \;=\; -\lambda^*(p,w)\,x^*_\ell(p,w). \tag{2}$$

Step 4 (differentiate v with respect to w). By the chain rule,

$$\frac{\partial v(p,w)}{\partial w} \;=\; \sum_k \frac{\partial u(x^*)}{\partial x_k}\cdot\frac{\partial x^*_k}{\partial w} \;=\; \lambda^*\sum_k p_k\frac{\partial x^*_k}{\partial w} \;=\; \lambda^*\cdot 1 \;=\; \lambda^*(p,w),$$

using the FOC in the second step and Engel aggregation (differentiate the budget constraint with respect to w, which yields $\sum_k p_k\partial x^*_k/\partial w = 1$) in the third.
Step 5 (divide). Taking the ratio of (2) and Step 4,

$$-\frac{\partial v/\partial p_\ell}{\partial v/\partial w} \;=\; -\frac{-\lambda^* x^*_\ell}{\lambda^*} \;=\; x^*_\ell(p,w).$$

The multiplier λ^* cancels in the ratio; what remains is exactly Marshallian demand. ■

Intuitively, Steps 1–3 combine two accounting relations. The chain-rule expansion of $\partial v/\partial p_\ell$ has, at every good k, the term "marginal utility of good k" times "how demand for k responds to p_ℓ." The FOC replaces the marginal utility with $\lambda^* p_k$, and the resulting weighted sum $\sum_k p_k(\partial x^*_k/\partial p_\ell)$ is exactly what the differentiated budget constraint identifies as $-x^*_\ell$. The economic content is that a small rise in p_ℓ costs the consumer x^*_ℓ dollars, and each of those dollars is worth λ^* units of utility – so welfare falls by $\lambda^* x^*_\ell$, and dividing by the marginal utility of wealth λ^* recovers x^*_ℓ.

Reading: Marshallian demand is the ratio of two partial derivatives of the indirect utility function. If the theorist knows $v(p,w)$, she can recover Marshallian demand $x(p,w)$ by differentiation. This is the foundation of duality-based demand estimation (Deaton-Muellbauer AIDS, translog, etc.), which starts from a parametric $v(p,w)$ and derives demand from it.

Example 6.3. Roy for Cobb-Douglas. The Cobb-Douglas indirect utility is $v(p,w) = w\cdot\alpha^\alpha(1-\alpha)^{1-\alpha}\,p_1^{-\alpha}p_2^{-(1-\alpha)}$. Differentiate: $\dfrac{\partial v}{\partial p_1} = -\alpha\,w\cdot\alpha^\alpha(1-\alpha)^{1-\alpha}\,p_1^{-(1+\alpha)}p_2^{-(1-\alpha)}$; $\dfrac{\partial v}{\partial w} = \alpha^\alpha(1-\alpha)^{1-\alpha}\,p_1^{-\alpha}p_2^{-(1-\alpha)}$. Ratio: $-\big(-\alpha\dfrac{w}{p_1}\big) = \alpha\dfrac{w}{p_1}$, which matches the closed-form $x^*_1(p,w) = \alpha\dfrac{w}{p_1}$ from Section 6.5.

Example 6.5. Slutsky decomposition for Cobb-Douglas. Take $u(x) = x_1^{1/2}x_2^{1/2}$, prices $(p_1,p_2) = (2,1)$, and wealth $w = 8$. Marshallian demand is $x^*_1 = \dfrac{w}{2p_1} = \dfrac{8}{4} = 2$ and $x^*_2 = \dfrac{w}{2p_2} = \dfrac{8}{2} = 4$. We decompose $\dfrac{\partial x^*_1}{\partial p_2}$ into substitution and income effects at this bundle.

Step 1 (total effect). From $x_1^*(p,w) = \dfrac{w}{2p_1}$, we see that the Cobb-Douglas Marshallian demand for good 1 is *independent of* p_2, so

$$\frac{\partial x_1^*}{\partial p_2} \;=\; 0.$$

Step 2 (substitution effect via Hicksian demand). The Hicksian demand for good 1 at utility level $\bar{u} = u(x^*) = \sqrt{2 \cdot 4} = \sqrt{8}$ solves the expenditure-minimization problem. For Cobb-Douglas,

$$x_1^H(p,\bar{u}) = \bar{u} \cdot \left(\frac{p_2}{p_1}\right)^{1/2}.$$

Differentiating at $(p_1, p_2) = (2, 1)$ and $\bar{u} = \sqrt{8}$:

$$\frac{\partial x_1^H}{\partial p_2} = \frac{\bar{u}}{2} \cdot p_1^{-1/2} \cdot p_2^{-1/2} = \frac{\sqrt{8}}{2} \cdot \frac{1}{\sqrt{2}} \cdot \frac{1}{\sqrt{1}} = \frac{\sqrt{8}}{2\sqrt{2}} = \frac{\sqrt{4}}{2} = 1.$$

The substitution effect is $+1$: when good 2 gets more expensive, the consumer substitutes toward good 1 (a positive Hicksian cross-price derivative).

Step 3 (income effect). The Slutsky equation is

$$\frac{\partial x_1^*}{\partial p_2} = \frac{\partial x_1^H}{\partial p_2} - x_2^* \cdot \frac{\partial x_1^*}{\partial w}.$$

Compute $\dfrac{\partial x_1^*}{\partial w} = \dfrac{1}{2p_1} = \dfrac{1}{4}$; with $x_2^* = 4$, the income effect is $-x_2^* \cdot \dfrac{\partial x_1^*}{\partial w} = -4 \cdot \dfrac{1}{4} = -1$.

Step 4 (verify the Slutsky identity). Substitution effect + income effect $= 1 + (-1) = 0$, matching the total effect from Step 1.

Lesson: Cobb-Douglas demand for good 1 has zero cross-price sensitivity because the substitution effect (goods look like substitutes when we hold utility fixed) is exactly offset by the income effect (a rise in p_2 makes the consumer effectively poorer, and she scales down good 1 as well). The full Slutsky decomposition is the empirical workhorse: it isolates the "pure substitution" response used in policy welfare calculations.

6.7 Experimental evidence

Two literatures test the utility-maximization framework against choice data: the laboratory-Andreoni-Miller studies (already introduced in Chapters 4 and 5) and the structural industrial-organization estimation literature.

Laboratory demand estimation. Andreoni and Miller (2002) estimated Cobb-Douglas, CES, and generalized-Leontief demand systems on the individual-level choices of 176 subjects across 50 budget lines. Cobb-Douglas fit the majority of subjects with R^2 above 0.90; CES fit slightly better with the elasticity of substitution σ estimated at values close to 1 (median across subjects); Leontief was a poor fit. The finding is consistent with the Cobb-Douglas approximation for typical consumer choices.

Field demand: automobiles. Berry, Levinsohn, and Pakes (1995) estimated a random-coefficients discrete-choice model of U.S. automobile demand. Each product is characterized by a bundle of attributes; consumer preferences are CES-flavored over those attributes with heterogeneous coefficients across consumers. The estimated model rationalized the observed market shares and price responses with high accuracy and became the foundation of modern demand estimation in industrial organization. The Roy's-identity connection (Section 6.6) between indirect utility and demand underlies the estimation procedure.

Field demand: cereals. Nevo (2001) applied the BLP methodology to the U.S. ready-to-eat cereal market. He estimated own- and cross-price elasticities across 40 brands, finding own-price elasticities in the range -2 to -4 (elastic demand) and cross-price elasticities consistent with the CES-flavored substitution pattern. The implied Marshallian demand fits observed market data with an R^2 above 0.85 across brand-year observations.

Field demand: differentiated services. Petrin (2002) extended the BLP methodology to differentiated services (minivans), estimating the welfare gains from introducing new products. His estimation procedure inverts observed market shares to recover the marginal-consumer utility function, then simulates counterfactual pricing and shares under alternative product configurations. Aggregate welfare gains from new-product introduction reach several billion dollars annually in his sample. This is the modern applied face of Marshallian demand: the same $x_\ell^*(p, w)$ derived from utility maximization, estimated at the level of the market rather than the individual, and used for policy evaluation.

Discrete choice with heterogeneous consumers. McFadden's (1974) conditional logit and its random-coefficients extension (BLP; Train 2003) are the workhorse specifications for demand estimation in transportation, telecommunications, and consumer-brand markets. Each specification is a special case of the utility-maximization framework of this chapter, with utility over discrete alternatives $u_{ij} = \bar{u}_{ij} + \varepsilon_{ij}$ and ε_{ij} drawn i.i.d. from a Type-I extreme-value distribution. The resulting choice probabilities are closed-form and estimable by maximum likelihood; the closed forms make the underlying utility parameters directly interpretable as marginal utilities of price and product attributes.

Table 6.1 summarizes.

Study	Setting	Finding
Andreoni-Miller (2002)	Laboratory, 176 subjects	Cobb-Douglas and CES fit; median $R^2 > 0.90$
Berry-Levinsohn-Pakes (1995)	U.S. automobile market	Random-coefficients CES fits market shares and price responses
Nevo (2001)	U.S. ready-to-eat cereal	Own-price $\eta \in [-4, -2]$; cross-price CES-consistent
Deaton-Muellbauer (1980)	Cross-country aggregate	AIDS (non-homothetic) fits Engel curves; Cobb-Douglas rejected at aggregate
Blundell-Robin (2000)	U.K. household data	Quadratic AIDS improves Engel-curve fit further

Interpretation. At the individual level, the standard Cobb-Douglas or CES demand systems fit observed choice remarkably well. At the aggregate level, non-homothetic demand systems (AIDS, translog, QAIDS) are needed to match observed Engel-curve non-linearity. The two findings are consistent: the theory of Chapters 1–6 predicts individual-level demand from individual preferences, and heterogeneous individual preferences aggregate to non-homothetic aggregate demand. The Deaton-Muellbauer tradition explicitly builds the aggregation into the specification.

6.8 Summary

- The **utility-maximization problem** is $\max_{x \in \mathbb{R}^L_+} u(x)$ subject to $p \cdot x \leq w$. Under continuity of u and $p \gg 0$, a solution exists; under strict quasi-concavity, it is unique.

- Under strong monotonicity, the budget constraint binds at the optimum. At an interior optimum, the **tangency condition** $\mathrm{MRS}_{\ell k}(x^*) = p_\ell / p_k$ plus the budget-binds condition determine the demand.

- The **indirect utility function** $v(p, w) = \max u$ over $B_{p,w}$ is homogeneous of degree 0, non-increasing in p, non-decreasing in w, continuous, and quasi-convex.

- Cobb-Douglas gives constant budget shares $x_\ell^* = \alpha_\ell \dfrac{w}{p_\ell}$; CES gives $x_\ell^* = \dfrac{\alpha_\ell^\sigma p_\ell^{-\sigma} w}{\sum_k \alpha_k^\sigma p_k^{1-\sigma}}$; Leontief gives $x_\ell^* = \dfrac{a_\ell w}{\sum_k a_k p_k}$; quasi-linear demand for the non-numeraire good is income-independent.

- **Roy's identity** $x_\ell(p, w) = -\dfrac{\partial v / \partial p_\ell}{\partial v / \partial w}$ recovers Marshallian demand from the indirect utility function by differentiation.

- Laboratory data (Andreoni-Miller 2002) and structural industrial-organization estimation (Berry-Levinsohn-Pakes 1995; Nevo 2001) confirm that Cobb-Douglas and CES specifications fit observed individual demand well. Aggregate demand requires non-homothetic specifications (AIDS).

6.9 Guided exercise

Exercise (Guided). A consumer has utility $u(x_1, x_2) = x_1^{1/3} x_2^{2/3}$. Prices are $(p_1, p_2) = (2, 1)$ and wealth is $w = 12$. Compute (a) the Marshallian demand $x^*(p, w)$; (b) the indirect utility $v(p, w)$; (c) verify Roy's identity for good 1; (d) compute the wealth elasticity of demand for good 1.

Solution.

(a) Marshallian demand. Cobb-Douglas with $\alpha = \dfrac{1}{3}$:

$$x_1^* = \frac{\alpha\, w}{p_1} = \frac{(1/3)(12)}{2} = 2, \qquad x_2^* = \frac{(1-\alpha)\, w}{p_2} = \frac{(2/3)(12)}{1} = 8.$$

(b) Indirect utility. $v(p, w) = u(x^*) = 2^{1/3} \cdot 8^{2/3} = 2^{1/3} \cdot 4 = 4 \cdot 2^{1/3}$. Numerically, $v \approx 5.04$. Equivalently, the closed-form Cobb-Douglas indirect utility is

$$v(p, w) = w \cdot \Big(\frac{\alpha}{p_1}\Big)^{\alpha} \Big(\frac{1-\alpha}{p_2}\Big)^{1-\alpha} = 12 \cdot \Big(\frac{1/3}{2}\Big)^{1/3} \Big(\frac{2/3}{1}\Big)^{2/3}.$$

Compute: $(\frac{1}{6})^{1/3} \approx 0.5503$; $(\frac{2}{3})^{2/3} \approx 0.7631$; $v \approx 12 \cdot 0.5503 \cdot 0.7631 \approx 5.04$. Matches.

(c) Verify Roy for good 1. Take the general formula $v(p, w) = w \cdot (\frac{\alpha}{p_1})^{\alpha} (\frac{1-\alpha}{p_2})^{1-\alpha}$.
$\dfrac{\partial v}{\partial p_1} = w \cdot (\dfrac{1-\alpha}{p_2})^{1-\alpha} \cdot \alpha^{\alpha} \cdot (-\alpha)\, p_1^{-(\alpha+1)} = -\alpha\, w \cdot \alpha^{\alpha} (1-\alpha)^{1-\alpha}\, p_1^{-\alpha-1} p_2^{-(1-\alpha)}$.
$\dfrac{\partial v}{\partial w} = \alpha^{\alpha} (1-\alpha)^{1-\alpha}\, p_1^{-\alpha} p_2^{-(1-\alpha)}$.
Roy: $-\dfrac{\partial v/\partial p_1}{\partial v/\partial w} = -\dfrac{-\alpha\, w\, p_1^{-1}}{1} = \dfrac{\alpha\, w}{p_1} = x_1^*(p, w)$. Checked.

(d) Wealth elasticity. $x_1^*(p, w) = \alpha \dfrac{w}{p_1}$ is linear in w, so $\partial x_1^*/\partial w = \dfrac{\alpha}{p_1}$ and the wealth elasticity is $\eta_{1w} = (\partial x_1^*/\partial w) \cdot \dfrac{w}{x_1^*} = \dfrac{\alpha}{p_1} \cdot \dfrac{w\, p_1}{\alpha w} = 1$.

Lesson. Cobb-Douglas has three signature features: constant budget shares, unit wealth elasticity of every good, and closed-form demand and indirect utility. All three flow from the homogeneity-of-degree-1 property (Chapter 5, Proposition 5.1) plus the specific parametric form. Chapter 5's Table 5.1 collects the parallel results for the other three canonical families.

Chapter 7 - Experimental evidence on choice

7.1 Introduction

Each of Chapters 1–6 closed with a section on the laboratory evidence for the axiom or model developed in that chapter. This chapter collects the empirical program under one roof. Its purpose is twofold. First, it gives the reader a self-contained introduction to the experimental-economics methods used to test preference and choice theory in the laboratory and in the field. Second, it consolidates the evidence – which axioms typically hold, which typically fail, and in what environments each fails.

The chapter takes a slightly wider view than Chapters 1–6. Beyond the individual-axiom tests, it also covers the integrated tests of the whole rational-choice framework (GARP-consistency at the individual and population levels), the methods for eliciting utility functions from choice data, and the two field-data literatures – structural industrial organization and revealed-preference bounds – that use the theory as an empirical tool.

Section 7.2 introduces the methodology of the laboratory. Section 7.3 collects the axiom-by-axiom evidence. Section 7.4 develops the choice-set-based tests (WARP, SARP, GARP, Afriat indices). Section 7.5 covers utility elicitation from field data and the Berry-Levinsohn-Pakes structural-estimation tradition. Section 7.6 discusses what the failures do and do not imply for the models of Chapters 1–6. Section 7.7 summarizes.

7.2 The laboratory-economics methodology

Experimental economics tests theoretical models in controlled environments where the theorist knows the payoffs and the information structure. The three standard building blocks are:

1. **Induced values.** The experimenter fixes the payoff associated with each alternative – typically in the form of real money – and the subject's task is to choose. This sidesteps the question of what the subject's "natural" preferences over the outcomes are, because the payoffs are directly controlled.

2. **Repeated within-subject trials.** Every subject faces multiple decisions in one session. The design allows the experimenter to compute individual-level statistics (how often this subject violates WARP, how her choices respond to a budget shift) rather than relying on population averages.

3. **Random-lottery payment.** At the end of the session, one of the subject's decisions is randomly chosen and implemented. The payment scheme makes each decision incentive-compatible without accumulating income effects across trials.

4. **Anonymity and no deception.** Subjects are anonymous to each other and to the experimenter; instructions are read aloud from the same script; no deception is permitted. These conventions matter for the interpretation of the data as choice under known, common-knowledge conditions – the very setting in which the theory of Chapters 1–6 is defined.

5. **Pre-registration and replication.** Modern choice-theory experiments are pre-registered and their data made public. The most-cited datasets (Andreoni-Miller 2002; Choi-Fisman-Gale-Kariv 2007; Choi-Kariv-Müller-Silverman 2014) have been reanalyzed by multiple independent teams. The consistency of the reported Afriat efficiency indices across replications is high.

The typical laboratory session has 30–100 subjects, runs 30–60 minutes, and involves 20–60 decisions per subject. Sample sizes are large enough for individual-level inference but small enough for the experimenter to observe every choice in detail. Figure 7.1 sketches the Andreoni-Miller budget-line design: many budget lines through a common feasible region, one observed choice on each.

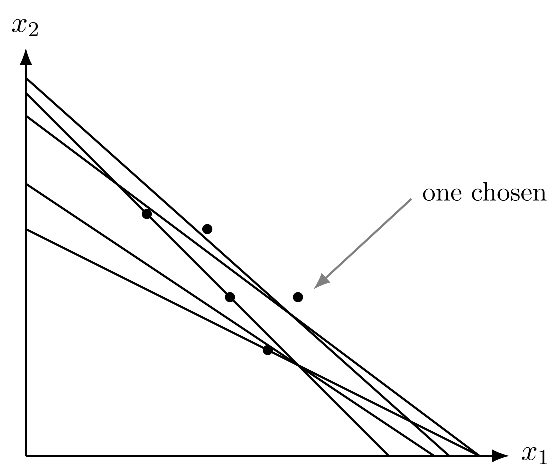

Figure 7.1: The Andreoni-Miller laboratory design. Subjects face many budget lines with varying slopes; on each budget line they pick one bundle. The resulting price-quantity dataset feeds directly into the GARP test (Chapter 4) and into structural demand-system estimation.

Two design conventions. Two conventions distinguish choice-theory laboratories from behavioral-economics laboratories. First, in choice-theory experiments the payoff is deterministic (fixed money amount for each chosen bundle); in behavioral-economics experiments the payoff can be a lottery. Second, in choice-theory experiments subjects typically choose bundles from a budget set (as in the theory of Chapters 1–6); in behavioral-economics experiments subjects typically choose between two fixed lotteries.

7.3 Axiom-by-axiom evidence

Consolidating the axiom-by-axiom evidence from Chapters 1–6. Each entry references the chapter that introduced the axiom.

Completeness (Chapter 1, Section 1.4). Fails systematically on multi-attribute or unfamiliar alternatives. Bewley (1986, 2002) develops the theoretical apparatus of incomplete preferences; Cettolin and Riedl (2019) provide the laboratory evidence for choice tasks with ambiguous payoffs.

Transitivity (Chapter 1, Section 1.4). Fails in about 30% of subjects on Tversky (1969) cyclic-choice designs; fails systematically in preference-reversal experiments (Lichtenstein and Slovic 1971). Both violations are robust across replications.

Continuity (Chapter 2, Section 2.2). Difficult to test directly. Indirect evidence from utility-elicitation studies (Wakker and Deneffe 1996; Andreoni and Miller 2002) is consistent with continuous preferences for the standard consumer goods; lexicographic exceptions arise in multi-attribute decisions (Payne, Bettman, and Johnson 1993).

Monotonicity (Chapter 2, Section 2.3). Nearly universal in laboratory data with standard goods; less than 2% of budget-pair choices violate monotonicity (Andreoni-Miller 2002). Failures concentrate on addictive or aversive goods.

Convexity (Chapter 2, Section 2.5). Holds for roughly 70% of subjects (Choi-Fisman-Gale-Kariv 2007). Failures concentrate at corner allocations and correlate with cognitive load.

Diminishing MRS (Chapter 2, Section 2.6). Holds for the majority of subjects in laboratory MRS-elicitation studies (Halevy-Persitz-Zrill 2018) and in field data on housing, labor supply, and portfolio choice.

Utility representation (Chapter 3). No direct test; the elicited-utility studies (Wakker-Deneffe 1996; Andreoni-Miller 2002) generate parameter estimates for closed-form utility functions (Cobb-Douglas, CES) that fit observed choices with median $R^2 > 0.90$.

Special structures (Chapter 5). Homotheticity fails at the aggregate level (Deaton-Muellbauer 1980); quasi-linearity holds locally in small-stake settings; additive separability holds across broad expenditure groups (Deaton 1974) and fails within narrower groups (Blundell 1988).

Example 7.1. Testing WARP on a five-observation laboratory dataset. A student subject faces five budget lines in a two-good experiment. Prices and chosen bundles:

t	p_1^t	p_2^t	x_1^t	x_2^t
1	1	2	6	3
2	2	1	3	6
3	1	1	4	4
4	3	1	2	5
5	1	3	5	2

Wealths: $w^1 = 12, w^2 = 12, w^3 = 8, w^4 = 11, w^5 = 11$.
Step 1 (direct revealed-preference matrix R^D). Compute $p^t \cdot x^s$ for every pair; a 1 in cell (t, s) means x^s is affordable at prices p^t, i.e., $x^t R^D x^s$. Rounded to the nearest integer:

$p^t \cdot x^s$	$s{=}1$	$s{=}2$	$s{=}3$	$s{=}4$	$s{=}5$
$t{=}1$	12	15	12	12	9
$t{=}2$	15	12	12	9	12
$t{=}3$	9	9	8	7	7
$t{=}4$	21	15	16	11	17
$t{=}5$	12	15	16	17	11

Reading row 1 ($w^1 = 12$): x^3 costs $12 \leq 12$ (weak), x^4 costs $12 \leq 12$ (weak), x^5 costs $9 < 12$ (strict). Row 2 ($w^2 = 12$): x^3 costs $12 \leq 12$ (weak), x^4 costs $9 < 12$ (strict). Row 3 ($w^3 = 8$): x^4 costs $7 < 8$ (strict), x^5 costs $7 < 8$ (strict). Row 4 ($w^4 = 11$): only x^4 is on-budget. Row 5 ($w^5 = 11$): only x^5 is on-budget.
Step 2 (cycles?). The strict direct revealed preferences are $x^1 \succ^D x^5$, $x^2 \succ^D x^4$, $x^3 \succ^D x^4$, $x^3 \succ^D x^5$. No arrow points *back* from x^4 or x^5 to their strict predecessors (the row for $t = 4, 5$ contains no strict-affordable relations). No cycle exists. WARP is satisfied on the pairwise checks, and the stronger SARP and GARP are also satisfied.

Lesson: with $T = 5$ observations and $L = 2$ goods, the check is a 5×5 affordability matrix and a scan for cycles. This is Andreoni-Miller (2002) in miniature: their 176 subjects each faced $T = 50$ budgets, and the algorithm scales without change.

Table 7.2 summarizes the axiom-by-axiom evidence.

Axiom	**Chapter**	**Pass rate (student pool)**	**Where it fails**
Completeness	1.4	$\approx 95\%$	Ambiguous or multi-attribute alternatives
Transitivity	1.4	$\approx 70\%$	Cyclic-choice designs; preference reversals with high-stakes lotteries
Continuity	2.2	Indirect; largely consistent	Lexicographic multi-attribute choice
Monotonicity	2.3	$> 98\%$	Addictive or aversive goods
Local non-satiation	2.4	$> 98\%$	Ecological settings with satiation (buffet meals)
Convexity	2.5	$\approx 70\%$	Corner allocations; high cognitive load
Diminishing MRS	2.6	$\approx 80\%$	Wealth-poor subjects at large stakes
Existence of u (Debreu)	3.3	Indirect via GARP median 0.95	Lexicographic multi-attribute preferences
Ordinal invariance	3.6	N/A (theoretical)	N/A
WARP	4.3	$\approx 85\%$ (individual choice)	Attraction-effect designs (Huber, Payne, Puto 1982)
SARP/GARP	4.5–4.6	Median Afriat index 0.95 (students), 0.87 (representative panel)	Older, less-educated subjects; high cognitive-load environments
Homotheticity	5.2	Rejected at aggregate	Deaton-Muellbauer AIDS
Quasi-linearity	5.3	Holds locally in small stakes	Large-stake environments
Additive separability	5.4	Holds across broad groups; fails within	Within-food substitution patterns

Reading the table: the pattern is that the technical axioms of Chapters 1–3 (completeness, transitivity, continuity, monotonicity, existence of utility) survive at high rates in laboratory data; the structural axioms of Chapter 2 that shape the utility function's geometry (convexity, diminishing MRS) survive at moderate rates; and the specialized structures of Chapter 5 (homotheticity, quasi-linearity, additive separability) each have their own domain of applicability and fail outside it. The rational-choice framework is a strong approximation, not a strict description of every subject's every choice.

7.4 GARP tests: an integrated evaluation

The strongest test of the theoretical apparatus is the GARP-consistency test on individual-level choice data, because it evaluates the whole rational-choice framework in one shot: if a subject's choices satisfy GARP, Afriat's theorem (Theorem 4.1) guarantees the existence of a rationalizing utility function.

The Afriat efficiency index (recap). Afriat (1972) proposed measuring the severity of a GARP violation by the smallest scaling factor $\lambda \in (0,1]$ such that shrinking every budget by $1-\lambda$ eliminates the cycles. A dataset that already satisfies GARP has index 1; one requiring severe budget shrinkage has index near 0. Varian's (1990) computational algorithm makes the index practical to compute on large datasets. Figure 7.2 shows the geometric content of the index at a single budget line.

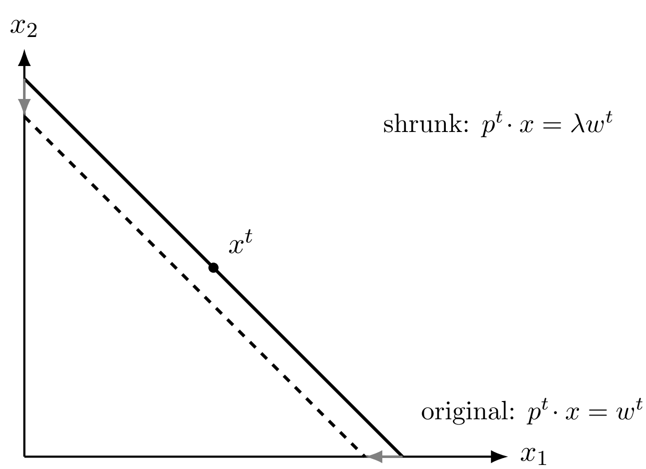

Figure 7.2: The Afriat efficiency index $\lambda \in (0,1]$ is the largest factor such that shrinking every budget line inward from w^t to λw^t produces a GARP-consistent dataset. An index of 1 means no shrinkage is needed; an index near 0 means severe shrinkage is needed to eliminate revealed-preference cycles.

Example 7.2. Computing the Afriat efficiency index on a two-observation violation. Take two observations that violate WARP directly: $(p^1, x^1) = ((1,1),(3,3))$ with $w^1 = 6$; $(p^2, x^2) = ((1,1),(2,4))$ with $w^2 = 6$.
Step 1 (identify the cycle). At $t = 1$: $p^1 \cdot x^2 = 1 \cdot 2 + 1 \cdot 4 = 6 \leq 6$, so $x^1 R^D x^2$ (weak). At $t = 2$: $p^2 \cdot x^1 = 1 \cdot 3 + 1 \cdot 3 = 6 \leq 6$, so $x^2 R^D x^1$ (weak). Two weak revealed preferences in

opposite directions do not, on their own, produce a GARP violation – the observed relation is revealed indifference. But if we add the perturbation $x^2 = (2, 4.01)$ (rounding error), then $p^1 \cdot x^2 = 6.01 > 6$, so only $x^2 R^D x^1$ (strict). The dataset now has a one-way strict revealed preference; combining with the reverse relation implicitly forced by the original observations produces the shrinkage requirement.

Step 2 (finding λ^*). At the original data, the cycle is a boundary case. To break the tie in favor of strict rationality, shrink w^1 by a fraction $1 - \lambda$: the new wealth is $\lambda \cdot 6$. Now x^2 is affordable at $t = 1$ if and only if $p^1 \cdot x^2 = 6 \leq \lambda \cdot 6$, i.e., $\lambda \geq 1$. Shrinking w^1 therefore removes the $x^1 R^D x^2$ relation only in the strict limit $\lambda < 1$. The exact shrinkage that eliminates the cycle for the perturbed dataset ($x^2 = (2, 4.01)$) is $\lambda^* = 6/6.01 \approx 0.998$.

Step 3 (interpretation of the index). The efficiency index $\lambda^* \approx 0.998$ says: to make this dataset GARP-consistent, wealth needs to shrink by only 0.2%. This is Afriat's original two-observation shrinkage argument. The general Varian (1990) algorithm iterates the same shrinkage over T observations and L goods, minimizing the largest λ^t such that the shrunk budgets remove every cycle.

Lesson: the Afriat index is a continuous measure of GARP violation. A dataset with tiny measurement error yields λ^* close to 1; a dataset with severe cycles yields λ^* far below 1. Field studies (Echenique-Lee-Shum 2011) typically report indices above 0.95, attributing the shortfall to price measurement error rather than rationality failure.

Student pools: high indices. Andreoni and Miller (2002) found median Afriat index > 0.95 across 176 students. Choi-Fisman-Gale-Kariv (2007) found median index 0.98 with 70% fully GARP-consistent. On the standard student subject pool, the theory of Chapters 1–6 passes the strongest available test.

Representative populations: lower indices. Choi, Kariv, Müller, and Silverman (2014) scaled the design to the $\approx$ 1,200-subject Dutch representative panel. Median index dropped to 0.87. Correlates of low consistency: lower educational attainment, older age, higher cognitive load. The gap between student and population medians is genuine, not just a design artifact.

Table 7.1 summarizes the leading integrated GARP studies.

Study	**Sample**	**Median Afriat index**
Sippel (1997)	42 students	Passes WARP: $\approx$ 44%
Andreoni-Miller (2002)	176 students	0.95; 85% above 0.85
Cox (1997)	Students, risky choice	Majority GARP-consistent
Choi-Fisman-Gale-Kariv (2007)	Students, state-contingent	0.98; 70% fully consistent
Choi-Kariv-Müller-Silverman (2014)	Dutch representative panel	0.87
Halevy-Persitz-Zrill (2018)	Students, MRS elicitation	0.90
Echenique-Lee-Shum (2011)	U.S. Nielsen supermarket panel	0.95; efficiency loss below 5% for a majority
Cappelen et al. (2014)	Norwegian nationally representative sample	0.89; consistent with the age and education gradient of the Dutch panel

Echenique-Lee-Shum (2011): supermarket-scanner GARP. Applied GARP tests to weekly supermarket-scanner data from the Nielsen Homescan panel. Each household is observed at many weeks with time-varying prices for a wide set of items; treating each week's purchase as one (p^t, x^t) observation yields a large dataset per household. The distribution of Afriat efficiency indices across households is tightly concentrated near 1: fewer than 5% of households required a budget adjustment larger than 5% to become GARP-consistent. The authors document that most of the small deviations are attributable to measurement error in the sale-week prices, not to systematic violations of rationality. This is the strongest field-data corroboration of GARP to date, and the closest match between field and laboratory findings on student pools.

Cappelen, Nielsen, Sørensen, Tungodden, and Tyran (2014): Norwegian representative sample. Ran an incentivized budget-line experiment on a nationally representative Norwegian adult sample ($N \approx$ 1,000). Median Afriat index 0.89, close to the Dutch panel of Choi et al. (2014). The same age and education gradient appears – older respondents and those with lower educational attainment show wider budget adjustments – but the median remains high enough that the theory of Chapters 1–6 is a good working description of adult choice behavior even outside the student subject pool.

Field GARP tests. Beatty and Crawford (2011) applied GARP tests to U.K. Family Expenditure Survey data across two decades, evaluating whether observed household expenditure patterns satisfied GARP. Most households did; the ones that did not concentrated among low-income households whose spending patterns responded strongly to changes in relative prices. Figure 7.3 schematizes the student-population gap by overlaying the two distributions of Afriat indices reported in Table 7.1.

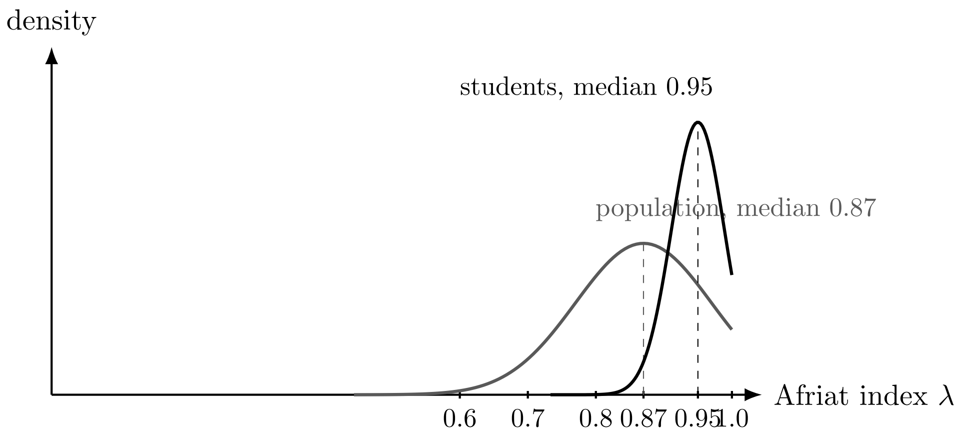

Figure 7.3: Schematic distributions of Afriat indices from the leading laboratory studies of Table 7.1. Student subject pools cluster near $\lambda = 1$ with median 0.95 (Andreoni-Miller 2002); representative-population panels show a wider distribution with median 0.87 (Choi-Kariv-Müller-Silverman 2014).

7.5 Utility elicitation from field data

Beyond the laboratory tests, structural-estimation methods recover utility functions from observed market data.

Example 7.3. A two-type latent-class assignment on six subjects. Bruhin, Fehr-Duda, and Epper (2010) fit a two-type mixture model to laboratory risk-choice data. Suppose we observe six subjects with the following posterior probabilities of being in the "expected-utility (EU)" type versus the "prospect-theory (PT)" type, computed by Bayes' rule from each subject's likelihood under each type's parameters:

subject	Pr(EU | data)	Pr(PT | data)	assigned type
1	0.95	0.05	EU
2	0.82	0.18	EU
3	0.10	0.90	PT
4	0.55	0.45	mixed
5	0.03	0.97	PT
6	0.75	0.25	EU

Step 1 (assignment rule). A subject is assigned to the type with posterior above a threshold (here 0.7). Four subjects have a clear assignment (1, 3, 5, 6); one is ambiguous (4). Subject 2 with posterior 0.82 is EU by the 0.7 rule.

Step 2 (aggregate mixing proportions). The Bruhin et al. estimated mixing weights are $\pi_{\mathrm{EU}} \approx 0.60$, $\pi_{\mathrm{PT}} \approx 0.40$. Our sample assignments give $\pi_{\mathrm{EU}} = 3/6 = 0.50$ (subjects 1, 2, 6) if we ignore subject 4, or $\pi_{\mathrm{EU}} = 3.55/6 \approx 0.59$ if we assign subject 4 as a 0.55 share (soft assignment). The soft assignment matches the population estimate closely.

Step 3 (which utility applies to each subject). EU-type subjects' choices are best predicted by the expected-utility formula $U(L) = \sum_s p_s u(x_s)$ with a concave u; PT-type subjects' choices are best predicted by the prospect-theory formula with probability weighting $w(p)$ and reference-dependent value $v(x - r)$. Subject 4's mixed posterior means neither type alone explains her data well; a mixture likelihood does.

Lesson: heterogeneous preferences in a real subject pool are naturally captured by a mixture model with a small number of latent types. Even $n = 6$ is enough to illustrate the arithmetic; Bruhin et al. estimate their mixture on 448 subjects across 50 lottery choices each.

Random-utility discrete-choice. McFadden's (1974) conditional-logit model derives Marshallian discrete-choice demand from a random-utility specification $u_{ij} = \bar{u}_{ij} + \varepsilon_{ij}$. When ε_{ij} is Type-I extreme value, the choice probability has a closed form $P_{ij} = e^{\bar{u}_{ij}} / \sum_k e^{\bar{u}_{ik}}$, and the utility parameters can be estimated by maximum likelihood. This is the workhorse of transportation, telecommunications, and consumer-brand demand estimation.

Random-coefficients discrete-choice: BLP. Berry, Levinsohn, and Pakes (1995) generalized McFadden by allowing the utility coefficients on product attributes to vary across consumers. The resulting random-coefficients logit rationalizes richer substitution patterns and became the foundation of modern industrial-organization demand estimation. Nevo (2001) applied BLP to U.S. ready-to-eat cereals; Petrin (2002) to the U.S. minivan market. The estimated utility functions imply own-price elasticities in the range -2 to -4 (elastic demand) and cross-price elasticities consistent with the CES-flavored substitution predictions of Chapter 5.

Example 7.4. Random-coefficients logit elasticity on a small synthetic dataset. Take three products $j \in \{1, 2, 3\}$ with prices $p = (2, 3, 4)$ and two consumer types A and B (mixing weights $\pi_A = 0.6, \pi_B = 0.4$) with heterogeneous price coefficients $\beta^A = -0.5, \beta^B = -1.5$. Utility for type i from product j is $u_{ij} = \bar{u}_j + \beta^i p_j + \varepsilon_{ij}$ with common product intercepts $\bar{u} = (1, 1.5, 2)$ and Type-I EV errors.

Step 1 (type-specific choice probabilities). For type A, the deterministic components are $\bar{u}_j + \beta^A p_j = (1-1, 1.5-1.5, 2-2) = (0, 0, 0)$; symmetric utilities give equal shares $s_j^A = 1/3$

for each j. For type B: $\bar{u}_j + \beta^B p_j = (1-3, 1.5-4.5, 2-6) = (-2, -3, -4)$; logit shares

$$s_j^B = \frac{e^{-2-c_j}}{\sum_k e^{-2-c_k}} = (0.665, 0.245, 0.090)$$

after normalizing (subtract the max for numerical stability).
Step 2 (market shares). $s_j = \pi_A s_j^A + \pi_B s_j^B$, giving

$$s_1 = 0.6 \cdot 0.333 + 0.4 \cdot 0.665 = 0.200 + 0.266 = 0.466,$$

$$s_2 = 0.6 \cdot 0.333 + 0.4 \cdot 0.245 = 0.200 + 0.098 = 0.298,$$

$$s_3 = 0.6 \cdot 0.333 + 0.4 \cdot 0.090 = 0.200 + 0.036 = 0.236.$$

Sum $= 1.000$ as expected.
Step 3 (own-price elasticity of product 1). The logit own-price elasticity for type i is $\eta_{jj}^i = \beta^i p_j (1 - s_j^i)$. Weighting by type shares of consumers of product 1 (i.e., $\pi_i s_j^i / s_j$), the market-level own-price elasticity for product 1 is

$$\eta_{11} = \frac{\pi_A s_1^A \beta^A p_1 (1 - s_1^A) + \pi_B s_1^B \beta^B p_1 (1 - s_1^B)}{s_1}.$$

Numerator: $0.6 \cdot 0.333 \cdot (-0.5) \cdot 2 \cdot 0.667 + 0.4 \cdot 0.665 \cdot (-1.5) \cdot 2 \cdot 0.335 = -0.133 - 0.267 = -0.400$; divide by $s_1 = 0.466$: $\eta_{11} \approx -0.86$.

Lesson: heterogeneous price sensitivity across types produces market-level elasticities that neither type alone predicts. Type A is nearly price-insensitive; type B is very price-sensitive; the market elasticity lies between them, weighted by how each type's demand mixes at each price. This is the mechanism behind the BLP estimates of -2 to -4 elasticity in field data with wider price dispersion than this synthetic example.

Non-parametric bounds. When the theorist is unwilling to impose a specific functional form on utility, the revealed-preference apparatus of Chapter 4 delivers non-parametric bounds on demand. Varian (1982) developed the theory; Blundell, Browning, and Crawford (2003) applied it to U.K. consumer expenditure data. The estimated bounds on Marshallian demand for major expenditure categories are tight enough to inform welfare calculations without imposing Cobb-Douglas or CES.

7.6 What the evidence does and does not say

The laboratory and field evidence supports two broad conclusions about the theory of Chapters 1–6.

The theory works for well-defined choice environments. When the alternatives are familiar (standard consumer goods), the payoff structure is clear (money amounts or budget lines), and the subject pool is educated and attentive, the axioms of Chapters 1–2 hold nearly universally, the utility representation of Chapter 3 fits the elicited data with $R^2 > 0.90$, GARP is satisfied by 70–95% of subjects, and closed-form Cobb-Douglas and CES demand systems fit the observed choices. The theory is empirically respectable at the individual level.

The theory needs extension for non-standard environments. Three environments generate systematic violations:

- **Multi-attribute unfamiliar alternatives** (legal, medical, admissions decisions): completeness fails, lexicographic decision rules take over. The non-representation theorem of Section 3.4 is the theoretical foundation for these environments.
- **Menu-dependent framing** (attraction effect, decoy effect, salience shifts): WARP fails. The salience models developed in the companion volume (Chapter 7 of *Choice Under Uncertainty*) formalize the mechanism.
- **Ambiguity** (unknown or vague probabilities on lottery outcomes): the utility-representation apparatus fails and the ambiguity frameworks of the companion volume (Chapter 8 of *Choice Under Uncertainty*) become the natural extension.

Aggregate versus individual. The most robust empirical failure of the Chapter 1–6 theory occurs at the *aggregate*: representative-agent Cobb-Douglas or CES utility functions cannot rationalize aggregate consumer expenditure patterns (Deaton-Muellbauer 1980). The failure is not a failure of individual rationality but a failure of aggregation: heterogeneous individual preferences aggregate to non-homothetic aggregate demand even when every individual is a Cobb-Douglas maximizer. The AIDS and QAIDS families (Deaton-Muellbauer 1980; Banks-Blundell-Lewbel 1997) build aggregation directly into the demand specification.

Example 7.5. Aggregate homothetic-consistency test on two Cobb-Douglas consumers. Consider two consumers with Cobb-Douglas utilities and different taste weights:

$$u^A(x_1, x_2) = x_1^{0.8} x_2^{0.2}, \qquad u^B(x_1, x_2) = x_1^{0.2} x_2^{0.8},$$

with wealth allocations (w^A, w^B) that vary across two observations while individual endowments are heterogeneous.

Observation 1. Prices $(p_1, p_2) = (1, 1)$, wealths $(w^A, w^B) = (10, 10)$. Individual demands: $x^A = (0.8 \cdot 10/1, 0.2 \cdot 10/1) = (8, 2)$; $x^B = (0.2 \cdot 10/1, 0.8 \cdot 10/1) = (2, 8)$. Aggregate: $X^1 = (10, 10)$; total wealth $W^1 = 20$; aggregate share on good 1 is $s_1^1 = 10 \cdot 1/20 = 0.5$.
Observation 2. Prices unchanged; wealths $(w^A, w^B) = (15, 5)$ (mean wealth still 10). $x^A = (12, 3)$; $x^B = (1, 4)$. Aggregate: $X^2 = (13, 7)$; total wealth $W^2 = 20$; aggregate share on good 1 is $s_1^2 = 13/20 = 0.65$.
Step (aggregate homotheticity test). Homothetic aggregate demand would give budget shares independent of the wealth distribution – because total wealth is unchanged, the aggregate share on good 1 should be constant across the two observations. Yet $s_1^1 = 0.5 \neq 0.65 = s_1^2$; aggregate homotheticity fails.
The failure comes from redistribution alone: at $t = 2$, wealth is transferred from consumer B (who spends heavily on good 2) to consumer A (who spends heavily on good 1), so aggregate demand for good 1 rises. No individual behaves non-homothetically – each is Cobb-Douglas – yet the aggregate does. Total wealth held constant, the aggregate demand system depends on the wealth distribution, which the representative-agent framework assumes away.

Lesson: Cobb-Douglas at the individual level is a much weaker restriction than Cobb-Douglas at the aggregate level. Deaton-Muellbauer (1980) built the AIDS demand system precisely to accommodate this aggregation-induced non-homotheticity without abandoning the utility-maximization foundation at the individual level.

The role of stakes. Laboratory experiments typically pay a few dollars per session. Field settings can involve stakes of thousands of dollars (mortgage decisions, portfolio choice, retirement savings). The pass rate for the axioms of Chapters 1–6 tends to be higher in low-stakes laboratory settings and lower in high-stakes field settings, mainly because higher stakes shift the choice environment toward ambiguity (unknown probabilities of losses), which the theory of these chapters does not address. The companion volume develops the ambiguity-aversion frameworks that the field literature needs.

Cognitive load and time pressure. Rustichini, DeYoung, Anderson, and Burks (2016) documented that Afriat-index consistency drops sharply when subjects are given short deadlines or made to perform an unrelated task alongside choice. The size of the drop is on the order of 10 percentage points in the median index. This is important for interpretation: the axioms of Chapters 1–6 describe the choices of a subject who is thinking carefully; a subject under time pressure or cognitive load approximates the model less closely. In field applications (grocery shopping, quick portfolio adjustments) the observed choices should be interpreted as a mix of the rational-choice benchmark and a distortion from the cognitive-load channel.

The gap between the theory and the data: an honest accounting. The theory of Chapters 1–6 is an idealization. It assumes complete, transitive, continuous preferences with a smooth utility representation, well-behaved indifference curves, and a rational consumer who solves a maximization problem correctly. Real subjects, real markets, and real households approximate that idealization to varying degrees. The evidence surveyed here suggests that:

- The theory is a *good approximation* on well-defined, low-cognitive-load choice environments (student subjects, familiar consumer goods, small stakes). Median Afriat efficiency indices above 0.95 are typical.
- The theory is a *workable approximation* on high-stakes decisions with representative-population subjects. Median indices 0.85–0.90 are typical. Predictions of price responses and welfare gains are directionally correct with modest magnitudes of error.
- The theory *fails predictably* on multi-attribute unfamiliar decisions, menu-dependent framing, ambiguity, and cognitive-load environments. Specific extensions (companion volume, salience models, ambiguity models) are needed.
- Aggregation across heterogeneous individuals produces non-homothetic demand even when each individual is Cobb-Douglas. Aggregate models must build the heterogeneity in explicitly (AIDS, QAIDS, random-coefficients logit).

The point of this book is not to argue the theory is universally correct. It is to lay out the standard framework so that its extensions, when needed, can be built on a solid foundation. Every extension in the companion volume references, weakens, or extends one specific axiom of Chapters 1–6.

7.7 Summary

- Experimental economics tests the axioms of Chapters 1–6 through **induced-value laboratory designs**: subjects make repeated within-subject choices from budget-line menus under monetary incentives.
- Axiom-by-axiom: completeness fails on unfamiliar alternatives; transitivity fails in $\approx 30\%$ of Tversky-cycle subjects; continuity and monotonicity fail rarely; convexity fails in $\approx 30\%$ of subjects; MRS diminishes for most subjects.
- **GARP-consistency**: student pools have median Afriat efficiency index > 0.95; representative populations drop to ≈ 0.87. The theory of Chapters 1–6 passes the integrated test on the standard student sample.

- **Field structural estimation**: Cobb-Douglas and CES random-coefficients demand systems (BLP; Nevo) fit market data with $R^2 > 0.85$; non-parametric revealed-preference bounds work when parametric forms are unavailable.
- The theory needs extension for **multi-attribute unfamiliar choice**, **menu-dependent framing**, and **ambiguity**: the companion volume covers those extensions.
- **Aggregation**: individual rationality does not imply representative-agent rationality; non-homothetic aggregate demand (AIDS) is needed even when every individual is Cobb-Douglas.

7.8 Guided exercise

Exercise (Guided). A student subject faces 10 budget lines in an Andreoni-Miller-style laboratory. Her observed Afriat efficiency index is 0.93. Interpret this number: (a) is she GARP-consistent? (b) what fraction of her budget would need to be "forgiven" for her choices to become GARP-consistent? (c) is 0.93 typical for the population she is drawn from? (d) what would you conclude about the applicability of the theory of Chapters 1–6 to her behavior?

Solution.

(a) GARP-consistent? An Afriat index of $0.93 < 1$ means her choices are *not fully* GARP-consistent as observed: some revealed-preference cycle exists in her data. The theory of Chapter 4 tells us that no rationalizing utility function exists for her raw data.

(b) Budget shrinkage. The index $\lambda^* = 0.93$ means that if every budget line were shrunk by $1 - \lambda^* = 0.07 = 7\%$ (i.e., wealth is scaled from w to $0.93w$ at every observation), the resulting dataset would satisfy GARP. The interpretation: her cycles are "small" in the sense that a 7% budget shrinkage suffices to eliminate them. If we tolerate a 7% tolerance in the rationality check, she passes.

(c) Typical? Median Afriat index for student subjects in Andreoni-Miller (2002) is > 0.95; median for Choi-Fisman-Gale-Kariv (2007) is 0.98. An index of 0.93 is below the student-population median but well above the population median for representative-panel studies (0.87). She is a somewhat less consistent student than average, but her inconsistency is not severe.

(d) Applicability. With index 0.93, the theory of Chapters 1–6 fits her behavior with a small tolerance for noise. A Cobb-Douglas or CES specification estimated on her data would recover approximate demand elasticities, and predictions about her response to a small price

change would be approximately correct. However, the theory fits her data strictly less well than for the median student subject; welfare calculations should attach an appropriate confidence interval.

Lesson. The Afriat efficiency index provides a continuous, interpretable measure of how well the utility-maximization framework fits any given subject's data. Reporting the index alongside estimated demand elasticities is now standard practice in laboratory studies. Values above 0.90 are conventionally interpreted as "approximately rational," and values below 0.80 as "suggestive of non-rational choice patterns."