7. Convex Analysis

So far, we have examined a single dual problem, minimizing expenditure subject to a utility constraint. This dual problem was associated with the primal problem of maximizing utility subject to a budget constraint. Duality can be used in other settings, and this chapter provides basic tools for studying duality in concave or convex problems.

Chapter Outline

1. Properties of Convex and Concave Functions
2. Separation Theorems
3. Supergradients
4. The Conjugate Function
5. Conjugates and Duality

The first section gives an introduction to some important results from convex analysis. By restricting our attention to differentiable concave functions on the real line we are able to quickly see some key results without technical complications.

The second section covers separation theorems. The separation theorems are powerful tools for converting functions and sets between goods space and price space. These theorems will play a key role in duality theory, and have other important applications in economics that do not obviously involve duality. That these other applications do involve moving from goods to prices suggests there is some sort of duality involved.

Supergradients, generalized derivatives, are next. Section three examines the relation between supergradients, supporting hyperplanes, and the conjugate function. We also find that concave functions are differentiable if and only if they have a unique supergradient.

The fourth section introduces the conjugate functions. Three big economic problems involve conjugate functions. The expenditure function is a conjugate function, as are the cost and profit functions. One of the insights of Fenchel (1953) was the importance of subgraphs and epigraphs in convex analysis. This shows up immediately when we try to understand the meaning of the conjugate function. It describes hyperplanes that support the subgraph or epigraph.

With the basic tools in place, section five focuses on a class of problems that includes our big three economic problems. These all involve indicator functions and support functions. We establish a theorem that gives us a considerable amount of information about these functions and their associated maximization or minimization problems.

We will apply these tools to some economic problems in the next chapter.

7.1 Introduction to Convex Analysis

We start by examining some basic concepts of convex analysis in the simple setting of differentiable concave functions defined on the real line. Thus $f\colon \mathbb{R} \to \mathbb{R}$. The Support Property Theorem tells us that if f is differentiable at x_0, then[1]

$$f(x) \le f(x_0) + f'(x_0)(x - x_0) \tag{7.1.1}$$

By the Support Property Theorem, a differentiable function is concave if and if equation 7.1.1 holds. This important inequality is a special case of the *supergradient inequality*. It characterizes concavity for differentiable functions.

The right-hand side of equation 7.1.1 can be used to define a line,

$$y = f(x_0) + f'(x_0)(x - x_0).$$

This line has slope $f'(x_0)$ and is tangent to the graph of f at the point $\big(x_0, f(x_0)\big)$.[2]

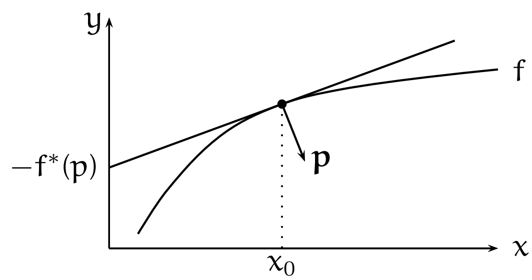

Figure 7.1.1: The tangent line at x_0 has the equation $y = f(x_0) + f'(x_0)(x - x_0)$. Because f is concave, the tangent line supports the graph of f. The graph is never above the tangent line and touches it at $\big(x_0, f(x_0)\big)$. The vector $\mathbf{p} = (p, -1)$ is perpendicular to the tangent and the vertical intercept is the concave conjugate function $f^*(p)$.

Let $p = f'(x_0)$ be the slope of the tangent line and the equation of the tangent is $p(x - x_0) = y - f(x_0)$. We now rewrite this equation in a way that that expresses it as a hyperplane in $\mathbb{R}^2$. We have

$$\mathbf{p}\cdot(x, y) = (p, -1)\cdot(x, y) = px_0 - f(x_0) = \mathbf{p}\cdot\big(x_0, f(x_0)\big). \tag{7.1.2}$$

Here the vector $\mathbf{p} = (p, -1)$ is perpendicular to the tangent line.

[1] The full theorem, with proof, is in section 31.1.1. It can also be found in Simon and Blume (1994) as Theorem 21.3.

[2] If the function f has a flat in its graph there may be multiple tangency points.

7.1.1 The Conjugate Function

The right-hand side of equation 7.1.2 is not zero unless tangent goes through the origin. It tells us how much the tangent line is offset from the origin. That value is called the *concave conjugate function* and is denoted $f^*(p) = px_0 - f(x_0)$ when $p = f'(x_0)$.

In fact, $f^*(p)$ is the negative of the vertical intercept of the tangent line. The equation of the tangent line then becomes

$$(p, -1)\cdot(x, y) = f^*(p), \quad \text{or} \quad y = px - f^*(p)$$

and the supergradient inequality is

$$(p, -1)\cdot\big(x, f(x)\big) \geq f^*(p), \quad \text{or} \quad f(x) \leq px - f^*(p) \tag{7.1.3}$$

The left-hand side is $px - f(x)$, the vertical distance between the line and the graph of f. Since the inequality is an equality at x_0, we have

$$f^*(p) = \inf\big[px - f(x)\big].$$

When $p = f'(x_0)$ the infimum is actually a minimum, attained at x_0.

We can immediately turn this around to ask whether a given p is the slope of a supporting tangent line. If the minimum does not exist for some value of p, it means that there is no supporting tangent line with slope p. In that case, $f^*(p) = -\infty$. If the minimum does exist, $-f^*(p)$ is the vertical intercept of the tangent line.

7.1.2 Recovering a Function from its Conjugate

Finally, if we start with a concave function f, and we know the conjugate function f^*, we know all of the supporting tangent lines via equation 7.1.3. It turns out that that is the information we need in order to find f itself. We merely take the envelope of all the tangent lines, as illustrated in Figure 7.1.2.

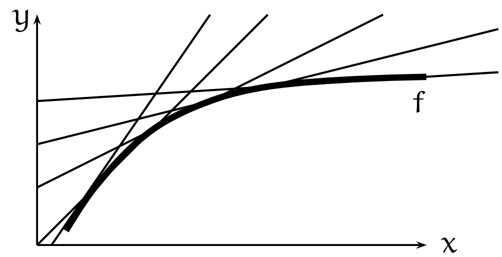

Figure 7.1.2: Since f is concave, its graph is the infimum of the tangent lines.

Convex analysis generalizes these results to the case where f is defined on $\mathbb{R}^m$ rather than $\mathbb{R}$ (this makes $\mathbf{p}$ a vector) and where f is not necessarily differentiable.

It will turn out that the envelope function is the conjugate of the conjugate—that $f^{**} = f$. We can apply these results to write expenditure, cost, and profit functions as conjugate functions, which allows us to recover the original utility and production functions as conjugates of expenditure, cost, and profit. The remainder of the chapter is devoted to that task.

7.2 Separation Theorems

Separation theorems are one of the most important mathematical tools available in economics. They play an important role in many of the key theorems of microeconomics. They pop up in duality arguments, in the welfare theorems, core equivalence, and the theory of asset pricing.

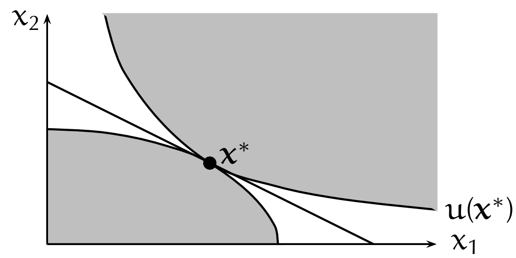

Figure 7.2.1: The optimum is at x^*. The same line is tangent to both the shaded production set and the indifference curve at x^*. We can decentralize the economy by interpreting the line as simultaneously indicating maximum profit for the producer, and as the consumer's budget constraint.

The basic idea is one that we see in intermediate micro. Suppose we have a Robinson Crusoe economy where one individual has quasiconcave preferences and a convex production possibilities set. The utility maximum is characterized by a mutual tangency of the optimal indifference curve and the production possibility set as in Figure 7.2.1. As we all know, we can use a price system to decentralize this economy. We re-interpret the tangent line as both the maximum isoprofit line for the producer and budget line for the consumer. This allows us to convert the optimal problem into an equilibrium, where the consumer maximizes utility, the firm maximizes profit, and markets clear (i.e., they choose the same point).

A separation theorem not only establishes that such a line exists in this simple case, but allows us to handle whole economies in a similar fashion.

7.2.1 Separation Theorem A

In fact, we will use four separation theorems, differing in their assumptions and the strength of the separation. The first theorem strictly separates a closed convex set from a point outside that set.

Separation Theorem A. *Suppose* $C \subset \mathbb{R}^m$ *is non-empty, closed, and convex and that* $\mathbf{x} \notin C$*. Then there is a vector* $\mathbf{p} \in \mathbb{R}^m$*,* $\mathbf{p} \neq \mathbf{0}$ *and a scalar* $\alpha \in \mathbb{R}$ *with* $\alpha < \mathbf{p}\cdot\mathbf{x}$ *and* $\mathbf{p}\cdot\mathbf{y} < \alpha$ *for all* $\mathbf{y} \in C$*.*

Proof. I claim there is a closest point in C to $\mathbf{x}$. Let $\bar{\mathbf{x}}$ be a point in C and $r = \|\bar{\mathbf{x}} - \mathbf{x}\| > 0$. Then define $B = C \cap \overline{B}_r(\mathbf{x})$. The set B is illustrated in Figure 7.2.2. As B is both closed and bounded, it is compact. By the Weierstrass Theorem, we can minimize the distance from B to $\mathbf{x}$. That also minimizes the distance from C to $\mathbf{x}$, proving the claim.

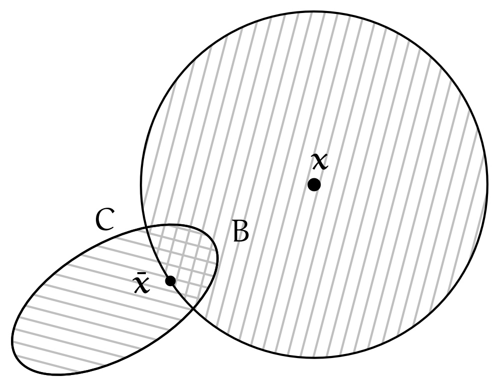

Figure 7.2.2: The set $B = C \cap \overline{B}_r(\mathbf{x})$ is cross-hatched.

7.2.2 Proof of Separation Theorem A

Proof of Theorem A. Now let $\boldsymbol{w}$ be a closest point in C to $\boldsymbol{x}$. Because C is closed, we know $\boldsymbol{w} \neq \boldsymbol{x}$. Define $\mathbf{p} = \boldsymbol{x} - \boldsymbol{w} \neq \mathbf{0}$.

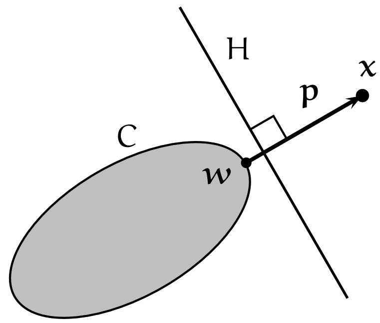

Figure 7.2.3: We now separate the convex set C from a point $\boldsymbol{x}$ by finding the nearest point in C to $\boldsymbol{x}$ ($\boldsymbol{w}$) and using the vector $\mathbf{p}$ from $\boldsymbol{w}$ to $\boldsymbol{x}$ to perform the separation. The separating hyperplane is perpendicular to $\mathbf{p}$ and given by $H = \{\boldsymbol{z} : \mathbf{p}\cdot\boldsymbol{z} = \alpha\}$. Since C lies left and below of H, $\mathbf{p}\cdot\boldsymbol{y} < \alpha$ for $\boldsymbol{y} \in C$.

For $\boldsymbol{y} \in C$, consider $\varepsilon\boldsymbol{y} + (1-\varepsilon)\boldsymbol{w} \in C$ for any ε with $1 > \varepsilon > 0$. It is at least as far from $\boldsymbol{x}$ as $\boldsymbol{w}$, so

$$\|\varepsilon\boldsymbol{y} + (1-\varepsilon)\boldsymbol{w} - \boldsymbol{x}\| \geq \|\boldsymbol{w} - \boldsymbol{x}\|$$

Squaring and expanding the left-hand side we obtain

$$\varepsilon^2\|\boldsymbol{y} - \boldsymbol{w}\|^2 + 2\varepsilon(\boldsymbol{w} - \boldsymbol{x})\cdot(\boldsymbol{y} - \boldsymbol{w}) + \|\boldsymbol{w} - \boldsymbol{x}\|^2 \geq \|\boldsymbol{w} - \boldsymbol{x}\|^2.$$

Cancelling $\|\boldsymbol{w} - \boldsymbol{x}\|^2$ from both sides and dividing by ε yields

$$\varepsilon\|\boldsymbol{y} - \boldsymbol{w}\|^2 + 2(\boldsymbol{w} - \boldsymbol{x})\cdot(\boldsymbol{y} - \boldsymbol{w}) \geq 0.$$

Let $\varepsilon \to 0$ to find $-2\mathbf{p}\cdot(\boldsymbol{y} - \boldsymbol{w}) \geq 0$. Then

$$\mathbf{p}\cdot\boldsymbol{w} \geq \mathbf{p}\cdot\boldsymbol{y}$$

for every $\boldsymbol{y} \in C$.

Now $\mathbf{p}\cdot(\boldsymbol{x} - \boldsymbol{w}) = \|\boldsymbol{x} - \boldsymbol{w}\|^2 > 0$, so $\mathbf{p}\cdot\boldsymbol{x} > \mathbf{p}\cdot\boldsymbol{w}$. Choose any α with $\mathbf{p}\cdot\boldsymbol{x} > \alpha > \mathbf{p}\cdot\boldsymbol{w}$. Then

$$\mathbf{p}\cdot\boldsymbol{x} > \alpha > \mathbf{p}\cdot\boldsymbol{w} \geq \mathbf{p}\cdot\boldsymbol{y}$$

for all $\boldsymbol{y} \in C$, completing the proof. $\square$

7.2.3 Lemma on Interiors

Before proving the second separation theorem, we need some more information about convex sets. We start with a lemma showing that a strict convex combination of a point in the interior of a convex set with a point in the closure is always in the interior of the set.

Lemma 7.2.4. *Let* $C \subset \mathbb{R}^m$ *be convex. If* $\mathbf{x} \in \operatorname{int} C$ *and* $\mathbf{y} \in \overline{C}$, *then* $(1-\alpha)\mathbf{x}+\alpha\mathbf{y} \in \operatorname{int} C$ *for* $0 \le \alpha < 1$.

Proof. Let $B = \{\mathbf{z} : \|\mathbf{z}\| < 1\}$, so $\mathbf{x} + \varepsilon B$ is the ball of radius ε around $\mathbf{x}$. Let $0 \le \alpha < 1$. Then

$$\begin{aligned}(1-\alpha)\mathbf{x} + \alpha\mathbf{y} + \varepsilon B &\subset (1-\alpha)\mathbf{x} + \alpha(C + \varepsilon B) + \varepsilon B \\ &= (1-\alpha)\left[\mathbf{x} + \varepsilon\left(\frac{1+\alpha}{1-\alpha}\right)B\right] + \alpha C\end{aligned}$$

for all $\varepsilon > 0$. The first line uses the fact that $\mathbf{y} \in \overline{C} \subset C + \varepsilon B$ and the second uses $\alpha < 1$. Now for ε small,

$$\mathbf{x} + \varepsilon\left(\frac{1+\alpha}{1-\alpha}\right)B \subset C.$$

This shows $(1-\alpha)\mathbf{x} + \alpha\mathbf{y} + \varepsilon B \subset C$ for ε small. In other words, the ε-ball about $(1-\alpha)\mathbf{x} + \alpha\mathbf{y}$ is contained in C, so $(1-\alpha)\mathbf{x} + \alpha\mathbf{y} \in \operatorname{int} C$. $\square$

7.2.4 Interiors and Closures of Convex Sets

A corollary of this is that if a convex set has an interior, the closure of the interior is the closure of the original convex set, and the interior of the closure is the interior of the original set.

Corollary 7.2.5. *Suppose* C *is convex with* $\operatorname{int} C \neq \emptyset$. *Then* $\overline{C} = \overline{(\operatorname{int} C)}$ *and* $\operatorname{int}(\overline{C}) = \operatorname{int} C$.

Proof. Now $\operatorname{int} C \subset C$ so $\overline{(\operatorname{int} C)} \subset \overline{C}$. Let $\mathbf{y} \in \overline{C}$ and take $\mathbf{x} \in \operatorname{int} C$. For all $\alpha \in [0, 1)$, $(1-\alpha)\mathbf{x} + \alpha\mathbf{y} \in \operatorname{int} C$ by Lemma 7.2.3. Letting $\alpha \to 1$ we see $\mathbf{y} \in \operatorname{cl}(\operatorname{int} C)$. Thus $\overline{C} \subset \operatorname{cl}(\operatorname{int} C)$, showing $\overline{C} = \overline{(\operatorname{int} C)}$.

For the second part, $\operatorname{int} C \subset \operatorname{int}(\overline{C})$ since $C \subset \overline{C}$. Now let $\mathbf{y} \in \operatorname{int}(\overline{C})$ and take $\mathbf{x} \in \operatorname{int} C$. Consider $(1-\alpha)\mathbf{x} + \alpha\mathbf{y}$. For $\alpha > 1$ with $\alpha - 1$ small, $(1-\alpha)\mathbf{x} + \alpha\mathbf{y} \in \operatorname{int}(\overline{C}) \subset \overline{C}$. Setting $\beta = 1/\alpha < 1$, the lemma implies $\mathbf{y} = (1-\beta)\mathbf{x} + \beta((1-\alpha)\mathbf{x} + \alpha\mathbf{y}) \in \operatorname{int} C$. This shows $\operatorname{int}(\overline{C}) \subset \operatorname{int} C$ and so $\operatorname{int}(\overline{C}) = \operatorname{int} C$. □

This may fail if the interior is empty or if the set is not convex.

► Example 7.2.6: Cases where Corollary 7.2.4 does not apply. In $\mathbb{R}^2$, let $C = B_1(\mathbf{0}) \cup \{\mathbf{x} \in \mathbb{R}^2 : x_1 \in [-2, 2], x_2 = 0\}$. Then $\operatorname{int} C = \operatorname{int} B_1(\mathbf{0})$ and $\overline{(\operatorname{int} C)} = B_1(\mathbf{0}) \neq C = \overline{C}$.

In $\mathbb{R}^2$, let $C = \{\mathbf{x} \in \mathbb{R}^2 : x_1 \in [-1, +1], x_2 = 0\}$. Then $\operatorname{int} C = \emptyset$, so $\overline{C} = C \neq \overline{(\operatorname{int} C)} = \emptyset$.

The foregoing examples show that first conclusion of Corollary 7.2.4 can fail if either of the assumptions concerning C fails.

In $\mathbb{R}$, if $C = (-1, 0) \cup (0, +1)$, $C = \operatorname{int} C$, but $\operatorname{int}(\overline{C}) = (-1, +1) \neq C$. This shows that the second conclusion of Corollary 7.2.4 can fail if C is not convex. ◄

7.2.5 Separation Theorem B

This extra information about convex sets allows us to use Separation Theorem A to prove a second separation theorem asserting that we can weakly separate two disjoint closed convex sets.

Separation Theorem B. *Suppose $A, B \subset \mathbb{R}^m$ are disjoint and convex. Then there is a vector $\mathbf{p} \neq \mathbf{0}$ and a scalar $\alpha \in \mathbb{R}$ with $\mathbf{p}\cdot\mathbf{a} \leq \alpha$ for all $\mathbf{a} \in A$ and $\alpha \leq \mathbf{p}\cdot\mathbf{b}$ for all $\mathbf{b} \in B$.*

Proof. Let $C = A - B$. Then C is convex with $\mathbf{0} \notin C$. By Corollary 7.2.4, $\mathbf{0} \notin \operatorname{int}(\overline{C}) = \operatorname{int} C \subset C$. Since $\mathbf{0} \notin \operatorname{int} C$, we can find $\mathbf{x}_n \notin \overline{C}$ with $\mathbf{x}_n \to \mathbf{0}$. Separation Theorem A yields $\mathbf{p}_n \neq \mathbf{0}$ with $0 > \mathbf{p}_n\cdot\mathbf{c}$ for all $\mathbf{c} \in C$.

Let $\mathbf{p}'_n = \mathbf{p}_n/\|\mathbf{p}_n\|$. Since the $\mathbf{p}'_n$ are bounded, then have a convergent subsequence. Let $\mathbf{p}$ be the limit. Note $\|\mathbf{p}\| = 1$. Moreover, $0 \geq \mathbf{p}\cdot\mathbf{c}$ for all $\mathbf{c} \in C$. Now $\mathbf{p}\cdot\mathbf{b} \geq \mathbf{p}\cdot\mathbf{a}$ for all $\mathbf{b} \in B$ and $\mathbf{a} \in A$. Since the $\mathbf{p}\cdot\mathbf{a}$ are bounded above, $\sup_{\mathbf{a}\in A} \mathbf{p}\cdot\mathbf{a}$ is finite. Let $\alpha = \sup \mathbf{p}\cdot\mathbf{a}$ to complete the proof. □

7.2.6 Separation Theorem C

As it stands, $\mathbf{p}\cdot\mathbf{x}$ can take the value α for points in both sets. The separating hyperplane can contain points in both sets. This is not possible if one of the sets is open. We can then sharpen Separation Theorem B to get strict separation.

Separation Theorem C. *Suppose $A, B \subset \mathbb{R}^m$ are disjoint and convex and that B is open. Then there is a vector $\mathbf{p} \neq \mathbf{0}$ and a scalar $\alpha \in \mathbb{R}$ with $\mathbf{p}\cdot\mathbf{a} \leq \alpha$ for all $\mathbf{a} \in A$ and $\alpha < \mathbf{p}\cdot\mathbf{b}$ for all $\mathbf{b} \in B$.*

Proof. Separation Theorem B gives us a vector $\mathbf{p} \neq \mathbf{0}$ and $\alpha \in \mathbb{R}$ with $\mathbf{p}\cdot\mathbf{a} \leq \alpha$ for all $\mathbf{a} \in A$ and $\mathbf{p}\cdot\mathbf{b} \geq \alpha$ for all $\mathbf{b} \in B$.

Suppose there is a $\mathbf{b}_0 \in B$ with $\mathbf{p}\cdot\mathbf{b}_0 = \alpha$. Since B is open, we can find $\varepsilon > 0$ with $B_\varepsilon(\mathbf{b}_0) \subset B$. Set

$$\mathbf{b}_0 - \frac{\varepsilon}{2}\frac{\mathbf{p}}{\|\mathbf{p}\|^2} \in B_\varepsilon(\mathbf{b}_0).$$

Since this is in B,

$$\alpha \leq \mathbf{p}\cdot\left(\mathbf{b}_0 - \frac{\varepsilon}{2}\frac{\mathbf{p}}{\|\mathbf{p}\|^2}\right) = \alpha - \frac{\varepsilon}{2},$$

which is impossible. This contradiction shows that $\mathbf{p}\cdot\mathbf{b}_0 = \alpha$ is impossible, so $\mathbf{p}\cdot\mathbf{b} > \alpha$ for all $\mathbf{b} \in B$. $\square$

7.2.7 Separation of Comprehensive Convex Sets

In some cases of economic interest the separating vector has a natural sign. A set $A \subset \mathbb{R}^m$ is *comprehensive* if $\mathbf{x} \in A$ and $\mathbf{x}' \leq \mathbf{x}$ implies $\mathbf{x}' \in A$. A set $A \subset \mathbb{R}^m$ is *anti-comprehensive* if $\mathbf{x} \in A$ and $\mathbf{x}' \geq \mathbf{x}$ implies $\mathbf{x}' \in A$.

When preferences are monotonic, the upper contour set is anti-comprehensive. Free disposal will often ensure that production sets are comprehensive. If besides being convex, either A is anti-comprehensive or B is comprehensive, then the separating price vector $\mathbf{p}$ must be positive.

Corollary 7.2.7. *Suppose $A, B \subset \mathbb{R}^m$ are disjoint and convex and that either A is comprehensive or B is anti-comprehensive. If there is a vector $\mathbf{p} \neq \mathbf{0}$ and a scalar $\alpha \in \mathbb{R}$ with $\mathbf{p}\cdot\mathbf{a} \leq \alpha$ for all $\mathbf{a} \in A$ and $\alpha \leq \mathbf{p}\cdot\mathbf{b}$ for all $\mathbf{b} \in B$, then $\mathbf{p} > \mathbf{0}$.*

Proof. Suppose A is comprehensive. Take $\mathbf{a} \in A$. Then $\mathbf{a} - t\boldsymbol{e}_i \in A$ for all $t > 0$. It follows that $\mathbf{p}\cdot(\mathbf{a} - t\boldsymbol{e}_i) \geq \alpha$. Dividing by t we obtain $(\mathbf{p}\cdot\mathbf{a})/t - p_i \leq \alpha/t$. Let $t \to \infty$ to see that $p_i \geq 0$ for each $i = 1, \dots, m$.

A similar argument using $\mathbf{b} + t\boldsymbol{e}_i$ works when B is anti-comprehensive. □

7.2.8 Separation Theorem D

Separation Theorem D is also an easy consequence of Separation Theorem B. It strongly separates a point not in the interior of a convex set from the interior of the set, and weakly separates the point from the entire convex set. This is often applied when the point being separated is on the boundary of the convex set.

Separation Theorem D. *Suppose* $C \subset \mathbb{R}^m$ *is non-empty and convex and that* $\mathbf{x} \notin \operatorname{int} C$. *Then there is a vector* $\mathbf{p} \neq \mathbf{0}$ *with* $\mathbf{p}\cdot\mathbf{x} < \mathbf{p}\cdot\mathbf{y}$ *for all* $\mathbf{y} \in \operatorname{int} C$ *and* $\mathbf{p}\cdot\mathbf{x} \leq \mathbf{p}\cdot\mathbf{y}$ *for all* $\mathbf{y} \in C$.

Proof. Apply Separation Theorem B to $A = C$ and $B = \{\mathbf{x}\}$ to obtain $\mathbf{p} \neq \mathbf{0}$ with $\mathbf{p}\cdot\mathbf{x} \geq \mathbf{p}\cdot\mathbf{c}$ for all $\mathbf{c} \in C$.

Now suppose $\mathbf{c} \in \operatorname{int} C$. For $\varepsilon > 0$ small enough, $\mathbf{c} + \varepsilon\mathbf{p} \in \operatorname{int} C$. Then $\mathbf{p}\cdot\mathbf{x} \geq \mathbf{p}\cdot\mathbf{c} + \varepsilon\|\mathbf{p}\|^2 > \mathbf{p}\cdot\mathbf{c}$. $\square$

7.2.9 Inverse Hicksian Demand MENTION

One application of Separation Theorem D is to the problem of inverting Hicksian demand. That is, given a consumption vector, can we find a price vector that yields that consumption as a Hicksian demand?

Separation Theorem D allows us to answer this in the affirmative.

Theorem 7.2.8. *Suppose a preference order $\succsim$ is continuous, convex, and locally non-satiated on a closed, convex consumption set $\mathfrak{X}$. If $\mathbf{x} \in \mathfrak{X}$, there is a $\mathbf{p} \neq \mathbf{0}$ with $\mathbf{p}\cdot\mathbf{x} \le \mathbf{p}\cdot\mathbf{x}'$ for all $\mathbf{x}' \in R(\mathbf{x}) = \{\mathbf{x}' : \mathbf{x}' \succsim \mathbf{x}\}$.*

Proof. By Corollary 2.5.5, the weakly preferred set $R(\mathbf{x})$ is the closure of the strictly preferred set $P(\mathbf{x})$. This implies $P(\mathbf{x}) = \operatorname{int} R(\mathbf{x})$, so $\mathbf{x} \notin \operatorname{int} R(\mathbf{x})$. The set $R(\mathbf{x})$ is convex by convexity of $\succsim$.

We now apply Separation Theorem D to find a $\mathbf{p} \neq \mathbf{0}$ with $\mathbf{p}\cdot\mathbf{x}' \ge \mathbf{p}\cdot\mathbf{x}$ for all $\mathbf{x}' \in R(\mathbf{x})$. □

When $\succsim$ is also represented by a utility function u, we can set $\bar{u} = u(\mathbf{x})$. Then Theorem 7.2.7 gives us a $\mathbf{p}$ with $\mathbf{x} \in \mathbf{h}(\mathbf{p}, \bar{u})$. Reasoning the other way, if $\bar{u}$ is given and $\bar{u}$ is in the range of u, there is a $\mathbf{p}$ with $\mathbf{x} \in \mathbf{h}(\mathbf{p}, \bar{u})$.

7.2.10 More about Inverse Hicksian Demand MENTION

There are three things to note here. First is that the price vector $\mathbf{p}$ in Theorem 7.2.7 is not unique. Any positive scalar multiple of $\mathbf{p}$ will also do the job. One way of dealing with this is to normalize prices so that $\mathbf{p}\cdot\mathbf{x} = 1$. This method is used later when examining the distance function. This method does not guarantee uniqueness, but if the utility function is differentiable and has non-zero derivative, the normalized $\mathbf{p}$ will be unique by Theorem 31.4.4. If there is a kink in the indifference curve, we cannot expect a unique price, even with normalization.

The second thing is that $\mathbf{p}$ might not be positive. Locally non-satiated preferences need not be monotonic, and if utility is decreasing in some direction at $\mathbf{x}$, the supporting price vector will not be positive. This happens when $u(x_1, x_2) = x_2 - (x_1 - 1)^2$ at the point $(3, 4)$. The supporting price vectors are positive multiples of $Du(3, 4) = (-4, 1)$.

Finally, in Examples 5.3.2 and 5.3.3, we use Hicksian demands to recover the utility function from the expenditure function. Implicit in this was the idea that the Hicksian demands trace out the indifference surface, which Theorem 7.2.7 shows.

7.3* Supergradients

Jan. 32, 2023

In ordinary calculus, we use the derivative to define a tangent line. When functions are concave, that tangent line is a supporting hyperplane. We are currently in the opposite situation. We have defined supporting hyperplanes for any functions with affine upper bounds (including concave functions), but do not have a corresponding generalized derivative. The supergradient is that generalized derivative.

We can close in on the supergradient by considering the inequalities that define the supporting hyperplanes. Suppose we can solve the minimization problem that defines the conjugate function. We find an $\mathbf{x}^*$ that minimizes $\mathbf{p}\cdot\mathbf{x} - f(\mathbf{x})$. Then

$$f^*(\mathbf{p}) = \mathbf{p}\cdot\mathbf{x}^* - f(\mathbf{x}^*) \le \mathbf{p}\cdot\mathbf{x} - f(\mathbf{x})$$

for all $\mathbf{x}$. Rearranging, we obtain $f(\mathbf{x}) \le f(\mathbf{x}^*) + \mathbf{p}\cdot(\mathbf{x} - \mathbf{x}^*)$. Here we were given $\mathbf{p}$ and needed $\mathbf{x}^*$ to obtain this inequality. If we turn the problem around (duality!) and ask what $\mathbf{p}$ satisfy the inequality for a given $\mathbf{x}^*$ we have the supergradient problem.

In order to develop the theory of supergradients, we need some basic concepts concerning concave functions and the real numbers.

Extended Real Numbers. Let $\mathbb{R}^* = \mathbb{R} \cup \{-\infty, +\infty\} = [-\infty, +\infty]$ denote the *extended real numbers*.

We have already used the extended real numbers when defining the expenditure function. After all, $e(\mathbf{p}, \bar{u}) = +\infty$ when the utility level $\bar{u}$ is not attainable.

Effective Domain. Let $A \subset \mathbb{R}^m$ be convex. When $f\colon A \to \mathbb{R}^*$ is concave, we define the *effective domain* of f by $\operatorname{dom} f = \{\mathbf{x} : f(\mathbf{x}) > -\infty\}$ and when f is convex, the *effective domain* of f is $\operatorname{dom} f = \{\mathbf{x} : f(\mathbf{x}) < +\infty\}$

The effective domain of a concave function includes all points where it is not negatively infinite. However, points where f is positively infinite are allowed. The definition is reversed for convex functions because f is convex if and only if $-f$ is concave.

Proper Function. A concave function f is *proper* if $f(\mathbf{x}) < +\infty$ for all $\mathbf{x}$ and $f(\mathbf{x}) > -\infty$ for some $\mathbf{x}$. Similarly, a convex function is *proper* if $f(\mathbf{x}) > -\infty$ for all $\mathbf{x}$ and $f(\mathbf{x}) < +\infty$ for some $\mathbf{x}$.

When f is a proper and concave function, its effective domain is non-empty and coincides with the set of points where f is finite.

It is easy to see that the expenditure function is proper if there is an $\bar{\mathbf{x}}$ with $u(\bar{\mathbf{x}}) \ge \bar{u}$. If no such $\bar{\mathbf{x}}$ exists, $e(\mathbf{p}, \bar{u}) = +\infty$ and the expenditure function is not proper.

7.3.1 Supergradients and Subgradients

Supergradient and Subgradient. For any $f\colon A \to \mathbb{R}^*$ where $A \subset \mathbb{R}^m$ is convex, a vector $\mathbf{p} \in \mathbb{R}^m$ is a *supergradient* of f at $\mathbf{x} \in \operatorname{dom} f$ if

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \mathbf{p}\cdot(\mathbf{y} - \mathbf{x})$$

for all $\mathbf{y} \in \operatorname{dom} f$.

The set of all supergradients of f at $\mathbf{x}$ is called the *superdifferential* and denoted $\partial^* f(\mathbf{x})$ and the inequality $f(\mathbf{y}) \leq f(\mathbf{x}) + \mathbf{p}\cdot(\mathbf{y} - \mathbf{x})$ is called the *supergradient inequality*. We say that f is *superdifferentiable* at $\mathbf{x}$ if $\partial^* f(\mathbf{x})$ is not empty.

A vector $\mathbf{p}$ is a *subgradient* of f at $\mathbf{x}$ if

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \mathbf{p}\cdot(\mathbf{y} - \mathbf{x})$$

for all $\mathbf{y} \in \operatorname{dom} f$. The *subdifferential*, the set of all subgradients of f at $\mathbf{x}$ is denoted $\partial_* f(\mathbf{x})$ and the inequality $f(\mathbf{y}) \geq f(\mathbf{x}) + \mathbf{p}\cdot(\mathbf{y} - \mathbf{x})$ is called the *subgradient inequality*. We say that f is *subdifferentiable* at $\mathbf{x}$ if $\partial_* f(\mathbf{x})$ is not empty.

Since $\partial^* f = -\partial_*(-f)$, any statement about supergradients of concave functions is easily translated to a statement about subgradients of convex functions. Supergradients are appropriate for concave functions, and we will mostly use supergradients. Anything is a supergradient when $f(\mathbf{x}) = +\infty$, and there are no supergradients at $f(\mathbf{x}) = -\infty$ unless $f = -\infty$ everywhere.

Functions that are not concave often do not have supergradients.

► **Example 7.3.1: A Function Without a Supergradient.** On $\mathbb{R}_+$, let $f(x) = x^2$. Now $f(x+y) - f(x) - py = (2x - p)y + y^2$. The right hand side is positive for y large enough, contradicting the supergradient inequality. ◄

7.3.2 Concavity and Supergradients

When a function is concave and differentiable, the supergradient and the derivative are the same, as shown by the following proposition.

Proposition 7.3.2. *If* f *is concave and differentiable at* $\mathbf{x}$*, the derivative* $Df(\mathbf{x})$ *is the only supergradient of* f *at* $\mathbf{x}$*.*

Proof. First suppose $\mathbf{p}$ is a supergradient. Then $f(\mathbf{x}+\varepsilon\mathbf{y}) \le f(\mathbf{x}) + \varepsilon\mathbf{p}\cdot\mathbf{y}$ by the supergradient inequality. It follows that $[f(\mathbf{x}+\varepsilon\boldsymbol{e}_i) - f(\mathbf{x})]/\varepsilon \le p_i$ for all $\varepsilon > 0$. Letting $\varepsilon \to 0$, we find $\partial f/\partial x_i \le p_i$. Similarly, $[f(\mathbf{x}+\varepsilon\boldsymbol{e}_i) - f(\mathbf{x})]/\varepsilon \ge p_i$ for $\varepsilon < 0$, which yields $\partial f/\partial x_i \ge p_i$. Combining these two results shows $p_i = \partial f/\partial x_i$, so $\mathbf{p} = Df(\mathbf{x})$.

Next suppose f is concave and differentiable. We follow the proof from the Support Property Theorem.

$$f\big(\mathbf{x}+\varepsilon(\mathbf{y}-\mathbf{x})\big) = f\big((1-\varepsilon)\mathbf{x}+\varepsilon\mathbf{y}\big) \ge \varepsilon f(\mathbf{x}) + (1-\varepsilon)f(\mathbf{x}).$$

We can rearrange to obtain

$$f\big(\mathbf{x}+\varepsilon(\mathbf{y}-\mathbf{x})\big) - f(\mathbf{x}) \ge \varepsilon[f(\mathbf{y}) - f(\mathbf{x})].$$

Dividing by $\varepsilon > 0$ and taking the limit yields $Df(\mathbf{x})(\mathbf{y}-\mathbf{x}) \ge f(\mathbf{y}) - f(\mathbf{x})$, which is the supergradient inequality. Similar calculations with inequalities reversed apply when f is convex. $\square$

7.3.3 Non-Differentiable Function with Supergradients

If f is not differentiable at a point, it may still have supergradients.

► Example 7.3.3: Supergradients, but No Derivative. The concave function $f(x) = -|x|$ has many supergradients at $x = 0$, but no derivative there. For $y > 0$, $f(y) = -y \leq f(0) + p(y-0) = py$ whenever $-1 \leq p$ For $y < 0$, $f(y) = y \leq py$ whenever $1 \geq p$. It follows that $p \in \partial^* f(0)$ if and only if $p \in [-1, 1]$. When $x > 0$, the function is differentiable with $\partial^* f(x) = \{-1\}$ while if $x < 0$, $\partial^* f(x) = +1$. Thus

$$\partial^* f(x) = \begin{cases} +1 & \text{for } x < 0 \\ [-1, +1] & \text{for } x = 0 \\ -1 & \text{for } x > 0. \end{cases}$$

◄

7.3.4 Right and Left Derivatives SKIPPED

You will notice that in the example the supergradient at zero lies within the range of the derivative to the left ($f'(x) = +1$ for $x < 0$) and to the right ($f'(x) = -1$ for $x > 0$). For concave functions on $\mathbb{R}$, there will always be such a range. It is most easily expressed in terms of the left-hand and right-hand derivatives.

Let $f'(x-)$ and $f'(x+)$ denote the *left-hand* and *right-hand derivatives* of f,

$$f'(x-) = \lim_{\varepsilon \uparrow 0} \frac{f(x+\varepsilon) - f(x)}{\varepsilon} \quad \text{and} \quad f'(x+) = \lim_{\varepsilon \downarrow 0} \frac{f(x+\varepsilon) - f(x)}{\varepsilon}.$$

7.3.5 Difference Quotients of Concave Functions SKIPPED

Before examining the derivatives further, we prove a lemma concerning difference quotients of concave functions.

Lemma 7.3.4. *Suppose* $f\colon \mathbb{R} \to \mathbb{R}$ *is concave and* $x < y < z$*. Then*

$$\frac{f(y) - f(x)}{y - x} \geq \frac{f(z) - f(x)}{z - x} \geq \frac{f(z) - f(y)}{z - y}. \tag{7.3.4}$$

Lemma 7.3.4 is illustrated in Figure 7.3.5, where the three difference quotients of equation 7.3.4 are, from left to right, the slopes of segments AB, AC, and BC.

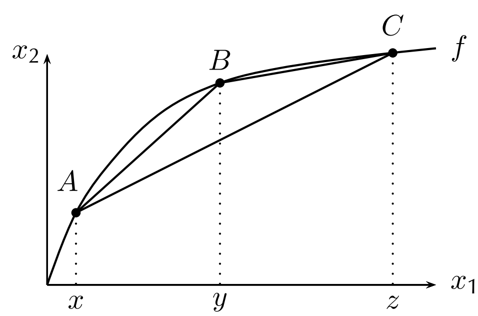

Figure 7.3.5: The argument in Lemma 7.3.4 is clear on a diagram. Because f is concave, the chord AB is steeper than AC and AC is steeper than BC. Since the difference quotients are the slopes of the chords, the lemma follows.

7.3.6 Proof of Lemma 7.3.4 SKIPPED

Figure 7.3.5 gave us the intuition for Lemma 7.3.4, which we now prove.

Lemma 7.3.4. *Suppose* $f\colon \mathbb{R} \to \mathbb{R}$ *is concave and* $x < y < z$*. Then*

$$\frac{f(y)-f(x)}{y-x} \geq \frac{f(z)-f(x)}{z-x} \geq \frac{f(z)-f(y)}{z-y}. \tag{7.3.4}$$

Proof. Let $x < y < z$ and define

$$\alpha = \frac{z-y}{z-x} \quad \text{so} \quad 1-\alpha = \frac{y-x}{z-x}.$$

By hypothesis, $0 < \alpha < 1$. This allows us to write y as a convex combination of x and z:

$$\alpha x + (1-\alpha)z = z + \alpha(x-z) = z + (y-z) = y.$$

Since f is concave,

$$f(y) \geq \alpha f(x) + (1-\alpha)f(z). \tag{7.3.5}$$

Then $f(y) - f(x) \geq (1-\alpha)[f(z)-f(x)]$. Substituting the expression for $(1-\alpha)$ and rearranging, we find

$$\frac{f(y)-f(x)}{y-x} \geq \frac{f(z)-f(x)}{z-x}. \tag{7.3.6}$$

Now rewrite equation (7.3.5) as

$$\alpha[f(z)-f(x)] \geq [f(z)-f(y)].$$

Using $\alpha = (z-y)/(z-x)$, we find

$$\frac{f(z)-f(x)}{z-x} \geq \frac{f(z)-f(y)}{z-y} \tag{7.3.7}$$

Combine equations 7.3.6 and 7.3.7 to complete the proof. □

7.3.7 Right and Left Derivatives of Concave Functions SKIPPED

Lemma 7.3.6. *Suppose* $x \in \operatorname{int} I$ *where* $I \subset \mathbb{R}$ *is an interval and* $f\colon I \to \mathbb{R}$ *is concave. Then* $f'(x-)$ *and* $f'(x+)$ *exist in* $\mathbb{R}^*$ *and* $f'(x-) \geq f'(x+)$.

Proof. Choose $\varepsilon > 0$ with $(x-\varepsilon, x+\varepsilon) \subset I$. Take ε_1, ε_2 with $\varepsilon > \varepsilon_2 > \varepsilon_1 > 0$. Set $y = x + \varepsilon_1$, $z = x + \varepsilon_2$ and apply equation (7.3.6) from Lemma 7.3.4 to obtain

$$\frac{f(x+\varepsilon_1) - f(x)}{\varepsilon_1} \geq \frac{f(x+\varepsilon_2) - f(x)}{\varepsilon_2}.$$

In other words, the difference quotient $[f(x+\varepsilon) - f(x)]/\varepsilon$ increases as ε decreases. Since the difference quotient is increasing, it has a limit, possibly $+\infty$. It follows that $f'(x+)$ exists. Similarly, the difference quotients for $f'(x-)$ decrease as $\varepsilon \uparrow 0$, so $f'(x-)$ also exists and is possibly $-\infty$.

Now rewrite equation (7.3.6) for the case $y < x < z$.

$$\frac{f(y) - f(x)}{y - x} \geq \frac{f(z) - f(y)}{z - y}$$

Letting $y \uparrow x$ and $z \downarrow x$, we obtain $f'(x-) \geq f'(x+)$. $\square$

7.3.8 Directional Derivatives SKIPPED

When studying convex or concave functions, we will find both one-sided and (two-sided) directional derivatives useful. We start with the ordinary directional derivative, which takes the limit of difference quotients as $\varepsilon \to 0$ from both directions.

Directional Derivative. Let f be a real-valued function $f\colon \mathbb{R}^m \to \mathbb{R}$. The *directional derivative* of f at $\boldsymbol{x}$ in direction $\boldsymbol{v}$ is defined as

$$\mathcal{D}f(\boldsymbol{x};\boldsymbol{v}) = \lim_{\varepsilon\to 0} \frac{f(\boldsymbol{x}+\varepsilon\boldsymbol{v}) - f(\boldsymbol{x})}{\varepsilon}.$$

The one-sided directional derivative only allows $\varepsilon > 0$ when taking the limit of difference quotients, rather than considering all ε.

One-sided Directional Derivative. Let f be a real-valued function $f\colon \mathbb{R}^m \to \mathbb{R}$. The *one-sided directional derivative* of f at $\boldsymbol{x}$ in direction $\boldsymbol{v}$ is defined as

$$\vec{\mathcal{D}}f(\boldsymbol{x};\boldsymbol{v}) = \lim_{\varepsilon\to 0^+} \frac{f(\boldsymbol{x}+\varepsilon\boldsymbol{v}) - f(\boldsymbol{x})}{\varepsilon}.$$

The arrow on top of the $\mathcal{D}$ indicates we have a one-sided directional derivative.

Another way to think about the directional derivatives of f, is that we form a function $F(t) = f(\boldsymbol{x}+t\boldsymbol{v})$, $F\colon \mathbb{R} \to \mathbb{R}$ and take its derivative to obtain the directional derivative. This is a very general notion of a derivative, and makes sense whenever the operations of vector addition and scalar multiplication are continuous. It works for functions defined on any topogical vector space. This generalization is known as the *Câteaux derivative*.

7.3.9 One-sided Directional Derivatives of Concave Functions SKIPPED

When f is defined on $\mathbb{R}^m$, we can use Lemma 7.3.6 to establish a similar inequality for one-sided directional derivatives. It's worth noting that for $f\colon \mathbb{R} \to \mathbb{R}$, although $\vec{\mathcal{D}}f(x) = f'(x+; +1)$, the other one-sided directional derivative is not $f'(x-)$, but rather is its negative. That is $\vec{\mathcal{D}}f(x) = -f'(x-; -1)$.

Proposition 7.3.7. *Suppose* $f\colon \mathbb{R}^m \to \mathbb{R}$ *is concave. Then* $\vec{\mathcal{D}}f(\boldsymbol{x}; \boldsymbol{v})$ *and* $\vec{\mathcal{D}}f(\boldsymbol{x}; -\boldsymbol{v})$ *exist and obey* $-\vec{\mathcal{D}}f(\boldsymbol{x}; -\boldsymbol{v}) \geq \vec{\mathcal{D}}f(\boldsymbol{x}; \boldsymbol{v})$.

Proof. Define $\varphi(t) = f(\boldsymbol{x} + t\boldsymbol{v})$. Then $\varphi'(0+) = \vec{\mathcal{D}}f_{\boldsymbol{v}}(\boldsymbol{x})$ and $\varphi'(0-) = -\vec{\mathcal{D}}f(\boldsymbol{x}; -\boldsymbol{v})$, so $-\vec{\mathcal{D}}f(\boldsymbol{x}; -\boldsymbol{v}) \geq \vec{\mathcal{D}}f(\boldsymbol{x}; \boldsymbol{v})$ by Lemma 7.3.6. $\square$

7.4** The Conjugate Function

The conjugate function plays an important role in convex analysis and in economics. It sets up a duality between functions on goods space ($\mathbb{R}^m$) and functions on its dual price space (also $\mathbb{R}^m$). Thus functions of goods, such as utility or production, may be converted to functions of prices, such as expenditure or cost. Further, functions of prices can be converted back to functions of goods. In many cases, the conversion process is bidirectional. All the information contained in the function of goods (e.g., utility) is not only present in the function of prices (e.g., expenditure), but can be recovered.

One consequence of this duality is that when functions are continuous and concave (or convex) we have the option of specifying certain economic objects in terms of goods or prices. Consumer preferences can be defined by utility or by expenditure. A firm's production can be defined in terms of a production function or via a cost function. There is no difference between them. They are two ways of looking at the same information.

There are two flavors of conjugate function, concave and convex. The concave conjugate is usually used for concave functions (or functions with affine upper bounds) while the convex conjugate is usually used for convex functions (or functions with affine lower bounds).

The conjugate function plays a key role in economic duality. As we will see, expenditure, cost, and profit functions are not only conjugate functions, but conjugate functions of a particular type. Many of their properties follow from that fact.

7.4.1 Definition of Conjugate Function

Conjugate Functions. If f is an extended real-valued function on a convex set A, the *concave conjugate* of f is[3]

$$f^*(\mathbf{p}) = \inf_{x \in A} \{\mathbf{p} \cdot \mathbf{x} - f(\mathbf{x})\}$$

and the *convex conjugate* of f is

$$f_*(\mathbf{p}) = \sup_{x \in A} \{\mathbf{p} \cdot \mathbf{x} - f(\mathbf{x})\}.$$

The convex conjugate may be obtained from the concave conjugate as follows:

$$\begin{aligned} -f_*(\mathbf{p}) &= -\sup_A \{\mathbf{p} \cdot \mathbf{x} - f(\mathbf{x})\} \\ &= \inf_A \{-\mathbf{p} \cdot \mathbf{x} + f(\mathbf{x})\} \\ &= (-f)^*(-\mathbf{p}). \end{aligned}$$

Thus $f_*(\mathbf{p}) = -(-f)^*(-\mathbf{p})$. We will usually state theorems in terms of concave conjugates, but each of them can be converted into a similar theorem for convex conjugates. The function f need not be concave for a concave conjugate to be defined.

The concave conjugate is the infimum of functions that are upper semicontinuous and concave in $\mathbf{p}$. By Propositions 31.1.1 and 30.4.4, the conjugate is also a concave upper semicontinuous function. For similar reasons, the convex conjugate is always convex and lower semicontinuous.

[3] See Rockafellar (1970), esp. sections 12 and 30.

7.4.2 Indicator Functions

Conjugate functions frequently arise in economics, usually as the conjugate of an convex or concave indicator function.

Indicator Function. If A is a convex set, the *concave indicator function* of A is

$$\mathbb{I}_A(\boldsymbol{x}) = \begin{cases} 0 & \text{if } \boldsymbol{x} \in A \\ -\infty & \text{if } \boldsymbol{x} \notin A. \end{cases}$$

The effective domain of the indicator $\mathbb{I}_A$ is the set A. An indicator function $\mathbb{I}_A$ is proper if and only if A is non-empty.

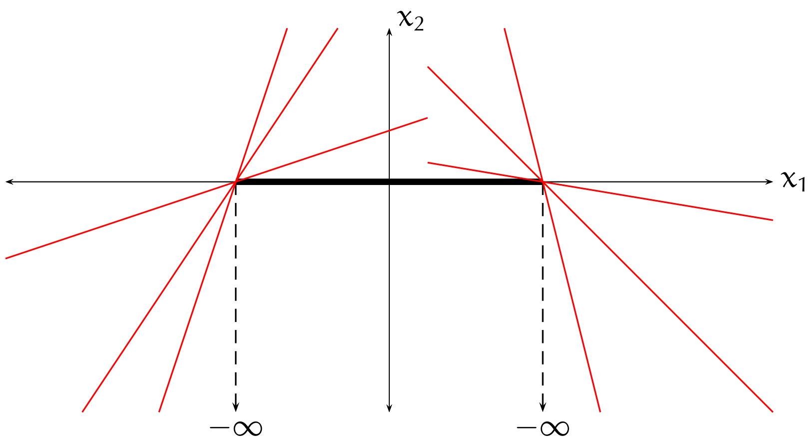

Figure 7.4.1: The indicator function for $[-2, +2]$ is shown, with supporting hyperplanes in red. For $x \in (-2, +2)$, the indicator is constant and the supergradient is 0. At the left corner, anything with a positive slope supports the graph, and $\partial^*\mathbb{I}(-2) = [0, +\infty)$. At the right corner, we need negative slopes for support, so $\partial^*\mathbb{I}(2) = (-\infty, 0]$. This is illustrated by segments of sample supporting tangents.

There is also a *convex indicator function*, $-\mathbb{I}_A$, which takes the value $+\infty$ on points outside of A. It's the concave indicator function turned upside-down.[4]

[4] These concave and convex indicator functions should not be confused with the indicator functions of measure theory, which take the values 0 and 1.

7.4.3 Conjugate Functions in Economics

► Example 7.4.2: Economic Conjugates. Let u be a utility function on $\mathfrak{X}$. Let $U(\bar{u}) = \{x \in \mathfrak{X} : u(x) \geq \bar{u}\}$ be the upper contour set and consider

$$\mathbf{p}\cdot\mathbf{x} - \mathbb{I}_{U(\bar{u})} = \begin{cases} \mathbf{p}\cdot\mathbf{x} & \text{when } u(\mathbf{x}) \geq \bar{u} \\ +\infty & \text{when } u(\mathbf{x}) < \bar{u} \text{ or } \mathbf{x} \notin \operatorname{dom} u \end{cases}$$

It follows that $\mathbb{I}^*_{U(\bar{u})}(\mathbf{p}) = e(\mathbf{p}, \bar{u})$, that the conjugate of the indicator of any upper contour set of utility is the the expenditure function.

One example from producer theory is the cost function (see Chapter 6). Suppose $f\colon \mathbb{R}^m_+ \to \mathbb{R}$ is a production function. The cost function $c_q(\mathbf{w}) = c(\mathbf{w}, q)$ is the concave conjugate of the concave indicator function of the upper contour set $F(q) = \{\mathbf{x} \in \mathbb{R}^m_+ : f(\mathbf{x}) \geq q\}$. Here $c_q(\mathbf{w}) = \mathbb{I}^*_{F(q)}(\mathbf{w})$.

Another example from producer theory is the profit function (see Chapter 14). Suppose Y is a production set and $\mathbf{p}$ a vector of prices. The profit function is the *convex* conjugate of the convex indicator function of the production set Y, $\pi(\mathbf{p}) = (-\mathbb{I}_Y)_*(\mathbf{p})$. ◄

As concave conjugates, the cost and expenditure functions are automatically concave and upper semicontinuous in prices while the profit function is convex and lower semicontinuous in prices as a convex conjugate.

7.4.4 Support Functions

The conjugate of an convex or concave indicator function is called a *support function*. The cost, expenditure, and profit functions are not just conjugates, they are support functions.

One easily established property of support functions is that they are always homogeneous of degree one.

Proposition 7.4.3. *Let $\mathbb{I}_A$ be a concave indicator function. Then the corresponding support function $\mathbb{I}_A^*$ is homogeneous of degree one.*

Proof. Let $t > 0$. Then $\mathbb{I}_A^*(t\mathbf{p}) = \inf\{t\mathbf{p}\cdot\mathbf{x} - \mathbb{I}_A(\mathbf{x})\} = \inf_{\mathbf{x}\in A} t\mathbf{p}\cdot\mathbf{x} = t\inf_{\mathbf{x}\in A} \mathbf{p}\cdot\mathbf{x} = t\mathbb{I}_A^*(\mathbf{p})$, establishing homogeneity. □

It follows that the cost, expenditure, and profit functions are all homogeneous of degree one in prices.

7.4.5 Conjugates and Supporting Hyperplanes

To understand the conjugate function a little better, note that $f^*(\mathbf{p}) \le \mathbf{p}\cdot\mathbf{x} - f(\mathbf{x})$ for all $\mathbf{x}$. Then $f(\mathbf{x}) \le -f^*(\mathbf{p}) + \mathbf{p}\cdot\mathbf{x}$. It is easy to see that this holds even if f and f^* are infinite.[5] For fixed $\mathbf{p}$, graph both f and $-f^*(\mathbf{p}) + \mathbf{p}\cdot\mathbf{x}$ as functions of $\mathbf{x}$. Then the graph of f lies just below the graph of the affine function $-f^*(\mathbf{p}) + \mathbf{p}\cdot\mathbf{x}$. Moreover, the two graphs either touch, or touch asymptotically.

Another way of thinking about this uses the region below the graph of f, the *subgraph* of f, $\operatorname{sub} f = \{(\mathbf{x}, y) \in \mathbb{R}^{m+1} : y \le f(\mathbf{x})\}$. It says that sub f is contained in the half-space $H_+ = \{(\mathbf{x}, y) : (\mathbf{p}, -1)\cdot(\mathbf{x}, y) \ge f^*(\mathbf{p})\}$ and that no smaller half-space of the form $\{(\mathbf{x}, y) : (\mathbf{p}, -1)\cdot(\mathbf{x}, y) \ge \alpha\}$ contains sub f. We refer to the set $H = \{(\mathbf{y}, z) : (\mathbf{p}, -1)(\mathbf{y}, z) = f^*(\mathbf{p})\}$ as a *supporting hyperplane* of sub f. More formally:

Supporting Hyperplane. Let $A \subset \mathbb{R}^m$. A hyperplane $H(\mathbf{p}, \alpha) = \{\mathbf{x} : \mathbf{p}\cdot\mathbf{x} = \alpha\}$ is a *supporting hyperplane* of A if $\mathbf{p} \ne \mathbf{0}$ and $A \subset H_+(\mathbf{p}, \alpha) = \{\mathbf{x} : \mathbf{p}\cdot\mathbf{x} \ge \alpha\}$ and if there is no larger α' with $A \subset H_+(\mathbf{p}, \alpha')$.

The convex conjugate f_* defines a minimal half-space that contains the region above the graph of f, the *epigraph* defined by $\operatorname{epi} f = \{(\mathbf{x}, y) \in \mathbb{R}^{m+1} : y \ge f(\mathbf{x})\}$.

[5] Note that $f^*(\mathbf{p}) = +\infty$ implies $f(\mathbf{x}) = -\infty$.

7.4.6 Asymptotic Support

If there is a $\mathbf{x}$ with $(\mathbf{x}, f(\mathbf{x})) \in H$, the supporting hyperplane of the subgraph of f is tangent to the subgraph at $(\mathbf{x}, f(\mathbf{x}))$. If there is no such $\mathbf{x}$, the hyperplane supports occurs asymptotically.

► Example 7.4.4: Asymptotic Supporting Hyperplane. Let $f\colon \mathbb{R} \to \mathbb{R}$ be given by $f(x) = -\infty$ for $x \leq 0$ and $f(x) = -1/x$ for $x > 0$. At $x > 0$, the only supergradient is $p = 1/x^2$, the derivative. However, a short computation shows $f^*(0) = 0$ even though 0 is not a supergradient. This indicates that the hyperplane $H = \{(0, -1), \mathbf{0}\}$ asymptotically supports the graph of f, as it obviously does.

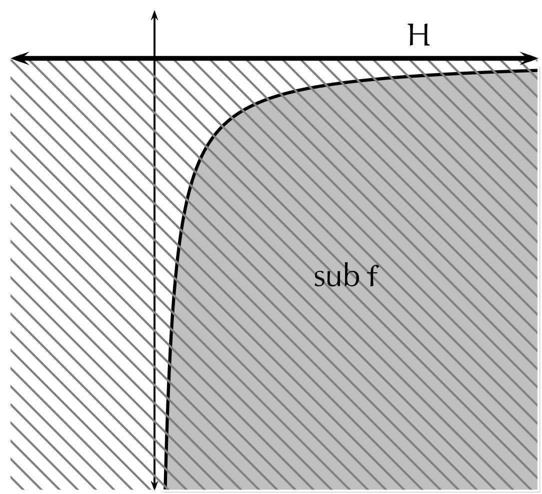

Figure 7.4.4: The hyperplane H asymptotically supports the subgraph of this function. The set H_+ is hatched, showing that it contains the subgraph of f. Although the subgraph and H do not touch, they are asymptotic as $x_1 \to +\infty$. No lower parallel hyperplane will lie above the subgraph.

◄

7.5** Conjugates and Duality

In the previous two sections we have looked at supporting hyperplanes, the conjugate function, and supergradients. Here, we tie them all together. The key results are the Young-Fenchel, Conjugate Duality, and Support Function Theorems. Once these are in hand, we will be ready to revisit our basic economic problems: cost minimization, expenditure minimization, and profit maximization.

7.5.1 Young-Fenchel Equality and Inequality

Two fundamental (and simple!) results relating supergradients and conjugates are the Young-Fenchel inequality and Young-Fenchel equality. Both are collected together in the Young-Fenchel Theorem.

Young-Fenchel Theorem. *If $f^*(\mathbf{p})$ and $f(\mathbf{x})$ are not oppositely infinite, then $f^*(\mathbf{p}) + f(\mathbf{x}) \leq \mathbf{p}\cdot\mathbf{x}$ (the Young-Fenchel inequality). Further, $f^*(\mathbf{p}) + f(\mathbf{x}) = \mathbf{p}\cdot\mathbf{x}$ if and only if $\mathbf{p}$ is a supergradient for f at $\mathbf{x}$ (the Young-Fenchel Equality).*[6]

Proof. By definition, $f^*(\mathbf{p}) \leq \mathbf{p}\cdot\mathbf{x} - f(\mathbf{x})$. Unless $f(\mathbf{x})$ and $f^*(\mathbf{p})$ are oppositely infinite, we can add $f(\mathbf{x})$ to both sides to obtain the Young-Fenchel inequality: $f^*(\mathbf{p}) + f(\mathbf{x}) \leq \mathbf{p}\cdot\mathbf{x}$.

If $f^*(\mathbf{p}) + f(\mathbf{x}) = \mathbf{p}\cdot\mathbf{x}$, $\mathbf{p}\cdot\mathbf{x} - f(\mathbf{x}) = f^*(\mathbf{p}) \leq \mathbf{p}\cdot\mathbf{y} - f(\mathbf{y})$ for all $\mathbf{y}$. This is the supergradient inequality. It shows $\mathbf{p}$ is a supergradient at $\mathbf{x}$.

If p is a supergradient, $f(\mathbf{y}) \leq f(\mathbf{x}) + \mathbf{p}\cdot(\mathbf{y} - \mathbf{x})$, so $\mathbf{p}\cdot\mathbf{x} - f(\mathbf{x}) \leq \mathbf{p}\cdot\mathbf{y} - f(\mathbf{y})$ for all $\mathbf{y}$. Taking the infimum over all $\mathbf{y}$ yields $\mathbf{p}\cdot\mathbf{x} - f(\mathbf{x}) \leq f^*(\mathbf{p})$. Combining this with the Young-Fenchel inequality shows $f^*(\mathbf{p}) + f(\mathbf{x}) = \mathbf{p}\cdot\mathbf{x}$. □

One way to view the conjugate function is that it tells us whether a given $\mathbf{p}$ is a supergradient somewhere. The Young-Fenchel Equality tells us that when $\mathbf{p}$ is a supergradient of f at $\mathbf{x}$, $f^*(\mathbf{p}) = f(\mathbf{x}) - \mathbf{p}\cdot\mathbf{x}$. Supergradients correspond to supporting hyperplanes that actually touch the graph of f. Supporting hyperplanes that are asymptotic are also identified by the conjugate function, but do not correspond to any supergradient at any point.

[6] Another version of the Young-Fenchel Theorem relates the convex conjugate f_* and subgradients of f.

7.5.2 Calculating the Conjugate: The Legendre Transformation

The close relation between the conjugate and supergradients means that the Young-Fenchel Theorem can often be used to calculate the conjugate function. The supergradient of $\mathbb{I}_{[0,1]}$ is

$$\partial^*\mathbb{I}_{[0,1]}(x) = \begin{cases} [0,\infty), & \text{if } x = 0 \\ 0, & \text{if } x \in (0,1) \\ (-\infty, 0], & \text{if } x = 1. \end{cases}$$

By the Young-Fenchel Equality, $\mathbb{I}^*_{[0,1}(p) + \mathbb{I}_{[0,1]}(0) = p \cdot 0 = 0$ for $p \geq 0$. Thus $\mathbb{I}^*_{[0,1]}(p) = 0$ when $p \geq 0$. For $p \leq 0$, $\mathbb{I}^*_{[0,1]}(p) + \mathbb{I}_{[0,1]}(1) = p \cdot 1 = p$. Thus $\mathbb{I}^*_{[0,1]}(p) = p$ when $p \leq 0$.

► Example 7.5.1: Calculating the Conjugate. Consider $f(\mathbf{x}) = -|x|$. Here $\partial^* f(\mathbf{x}) = 1$ for $x < 0$, $[-1, 1]$ at 0 and -1 for $x > 0$. For $p \in A = [-1, 1]$, $f^*(\mathbf{p}) + f(0) = 0$. Thus $f^*(\mathbf{p}) = 0$. If $p > 1$, let $x \to -\infty$ to see $f^*(\mathbf{p}) = -\infty$. If $p < -1$, let $x \to \infty$ shows $f^*(\mathbf{p}) = -\infty$. Thus $f^* = \mathbb{I}_A$. ◄

The Young-Fenchel Equality is even more powerful when f is differentiable and concave. If there is an $\mathbf{x}$ with $Df|_{\mathbf{x}} = \mathbf{p}$, we write it as $\mathbf{x}(\mathbf{p}) = (Df)^{-1}(\mathbf{p})$.[7] Then $\mathbf{p}$ is a supergradient at $\mathbf{x}(\mathbf{p})$. By the Young-Fenchel Equality, $f^*(\mathbf{p}) = \mathbf{p} \cdot \mathbf{x}(\mathbf{p}) - f(\mathbf{x})$. This method of calculating the conjugate is the *Legendre transform*. It can be written

$$f^*(\mathbf{p}) = \mathbf{p} \cdot \mathbf{x}(\mathbf{p}) - f\big(\mathbf{x}(\mathbf{p})\big)$$

or

$$f^*(\mathbf{p}) = \mathbf{p} \cdot (Df)^{-1}(\mathbf{p}) - f\big((Df)^{-1}(\mathbf{p})\big).$$

The following examples show how the Legendre transform can be used.

► Example 7.5.2: Legendre Transform. Take $f(x) = -x^2$. Then $Df = -2x = p$, so $x = -p/2$. Thus $f^*(p) = -p^2/2 + p^2/4 = -p^2/4$. ◄

► Example 7.5.3: Another Legendre Transform. Let $f(x, y) = -x^2 - y^2$. Then $Df = (2x, 2y) = -(p, q)$, so $(x, y) = -(p/2, q/2)$ and $f^*(p, q) = -(p, q) \cdot (p/2, q/2) - f(-p/2, -q/2) = -(p^2 + q^2)/4$. ◄

[7] This is the inverse of the function $\mathbf{x} \mapsto Df(\mathbf{x})$, not the inverse of the matrix Df.

7.5.3 The Conjugate Duality Theorem

The Young-Fenchel Theorem is one of the keys to the Conjugate Duality theorem, which gives further information about the relationship between supergradients and conjugates.

Conjugate Duality Theorem. *Suppose* $f\colon E \to \mathbb{R}^*$ *is proper, concave and upper semicontinuous. The following are equivalent:*

(1) $\mathbf{x}$ *minimizes* $\mathbf{p}\cdot\mathbf{y} - f(\mathbf{y})$.
(2) $\mathbf{p} \in \partial^* f(\mathbf{x})$.
(3) $f(\mathbf{x}) + f^*(\mathbf{p}) = \mathbf{p}\cdot\mathbf{x}$.
(4) $f^{**}(\mathbf{x}) + f^*(\mathbf{p}) = \mathbf{p}\cdot\mathbf{x}$.
(5) $\mathbf{x} \in \partial^* f^*(\mathbf{p})$.
(6) $\mathbf{p}$ *minimizes* $\mathbf{q}\cdot\mathbf{x} - f^*(\mathbf{q})$.

Proof. Condition (1) can be restated as $\mathbf{p}\cdot\mathbf{x} - f(\mathbf{x}) = \inf[\mathbf{p}\cdot\mathbf{y} - f(\mathbf{y})] = f^*(\mathbf{p})$. Thus (1) and (3) are equivalent. By the same reasoning, applied to f^*, (4) and (6) are equivalent. By the Young-Fenchel Equality, (2) and (3) are equivalent, as are (4) and (5).

Now, since f is concave and upper semicontinuous, $f^{**} = f$ by corollary 7.4.11. Then condition (3) is equivalent to condition (4). It follows that all six conditions are equivalent. □

For all functions, the first three conditions are equivalent, and imply the last three, which are equivalent. The implication follows since $p \in \partial^* f(\mathbf{x})$ implies $(-p, 1)$ supports sub f and thus sub f^{**}. Therefore p is a supergradient for f^{**} also and $f^{**}(\mathbf{x}) + f^*(\mathbf{p}) = px$. Concavity and upper semicontinuity are required for forging the link between (3) and (4).

7.5.4 Properties of the Support Function

Before proceeding we note some additional properties of support functions, and supergradients.

Proposition 7.5.4. *If $\mathbb{I} = \mathbb{I}_A$ is an concave indicator function, then the support function $\mathbb{I}^*$ is homogeneous of degree one and the correspondence $\mathbf{p} \twoheadrightarrow \{\mathbf{x} : \mathbf{p}\cdot\mathbf{x} - \mathbb{I}(\mathbf{x}) = \mathbb{I}^*(\mathbf{p})\}$ is homogeneous of degree zero.*

Proof. Note $\mathbb{I}^*(\mathbf{p}) = \inf_{\mathbf{x}\in A}\{px - \mathbb{I}(\mathbf{x})\} = \inf_{\mathbf{x}\in A} px$. Thus $\mathbb{I}^*(t\mathbf{p}) = \inf_{\mathbf{x}\in A} t\mathbf{p}\cdot\mathbf{x} = t\mathbb{I}^*(\mathbf{p})$ for $t > 0$.

Further $\mathbf{x} \in A$ obeys $\mathbf{p}\cdot\mathbf{x} = \mathbb{I}^*(\mathbf{p})$ if and only if $(tp)x = \mathbb{I}^*(tp)$ for $t > 0$. The correspondence $\mathbf{p} \twoheadrightarrow \{\mathbf{x} : \mathbf{p}\cdot\mathbf{x} - \mathbb{I}(\mathbf{x}) = \mathbb{I}^*(\mathbf{p})\}$ is thus homogeneous of degree 0. □

Another useful property of supergradients follows directly from the definition.

Lemma 7.5.5. *Let f be proper and concave. If $\mathbf{p} \in \partial^* f(\mathbf{x})$ and $\mathbf{p}' \in \partial^* f(\mathbf{x}')$, then $(\mathbf{p}' - \mathbf{p})\cdot(\mathbf{x}' - \mathbf{x}) \leq 0$.*

Proof. The supergradient inequality tells us that, $f(\mathbf{x}') \leq f(\mathbf{x}) + \mathbf{p}\cdot(\mathbf{x}' - \mathbf{x})$ and $f(\mathbf{x}) \leq f(\mathbf{x}') + \mathbf{p}'\cdot(\mathbf{x} - \mathbf{x}')$. Add these together, simplify and rearrange to get the result. □

Support functions enjoy a number of properties. Many of them follow from the Conjugate Duality Theorem, while others derive from the definition of a support function. We collect them together in the Support Function Theorem.

7.5.5 The Support Function Theorem

Support Function Theorem. *Let* A *be a convex set in* $\mathbb{R}^m$ *and* $\mathbb{I}_A(\mathbf{x})$ *its concave indicator function. Let* $f(\mathbf{p}) = \mathbb{I}_A^*(\mathbf{p})$ *be the associated support function and* $\mathbf{x}(\mathbf{p}) = \{\mathbf{x} \in A : \mathbf{p}\cdot\mathbf{x} = f(\mathbf{p})\}$ *be the set of solutions to* $\min_{\mathbf{x}\in A} \mathbf{p}\cdot\mathbf{x}$ *(if any). Then:*

(1) *The support function* $f(\mathbf{p})$ *is concave, upper semicontinuous and homogeneous of degree one. In addition,* f *is continuous on* int(dom f) *and weakly increasing in* $\mathbf{p}$.
(2) *If* A *is closed and convex,* $f^*(\mathbf{x}) = \mathbb{I}_A(\mathbf{x})$. *Thus* $A = \{\mathbf{x} : f^*(\mathbf{x}) = 0\}$. *Moreover,* $A = \{\mathbf{x} : \mathbf{p}\cdot\mathbf{x} \geq f(\mathbf{p})$ *for all* $\mathbf{p}\}$.
(3) *If* A *is convex, the minimizers obey* $\mathbf{x}(\mathbf{p}) = \partial^* f(\mathbf{p})$. *Moreover, if the minimizer is also unique, the directional derivative of* f *at* $\mathbf{p}$ *exists and is given by* $\mathcal{D}f(\mathbf{p};\mathbf{v}) = \mathbf{x}(\mathbf{p})\cdot\mathbf{v}$.
(4) *If* $\mathbf{x}_i \in \mathbf{x}(\mathbf{p}_i)$ *for* $i = 0, 1$, *then* $(\mathbf{x}_1 - \mathbf{x}_0)\cdot(\mathbf{p}_1 - \mathbf{p}_0) \leq 0$.
(5) $\mathbf{x}(\mathbf{p})$ *is homogeneous of degree zero.*
(6) *If* A *is convex and* $f(\mathbf{p})$ *is* $\mathcal{C}^2$, *then* $\mathbf{x}(\mathbf{p})$ *is differentiable and the matrix* $D\mathbf{x}(\mathbf{p}) = D^2 f(\mathbf{p})$ *is symmetric and negative semi-definite. Moreover,* $\big[D\mathbf{x}(\mathbf{p})\big]\mathbf{p} = \big[D^2 f(\mathbf{p})\big]\mathbf{p} = 0$.

Proof. Conditions (1) and (5) follow from Proposition 7.5.4 since f is a support function.

When A is closed and convex, $f^* = \mathbb{I}_A^{**} = \mathbb{I}_A$ by Corollary 7.4.11, establishing the first part of (2). The fact that $f^* = \mathbb{I}_A$ immediately yields $A = \{\mathbf{x} : f^*(\mathbf{x}) = 0\}$. By the Young-Fenchel inequality, $\mathbb{I}_A(\mathbf{x}) \leq \mathbf{p}\cdot\mathbf{x} - f(\mathbf{x})$. So if $\mathbf{x} \in A$, $\mathbf{p}\cdot\mathbf{x} - f(\mathbf{x})$ for all $\mathbf{p}$. If $\mathbf{x} \notin A$, $-\infty = \mathbb{I}_A(\mathbf{x}) = f^*(\mathbf{x}) = \inf\{\mathbf{p}\cdot\mathbf{x} - f(\mathbf{p})\}$. Thus there is a $\mathbf{p}^*$ with $\mathbf{p}^*\cdot\mathbf{x} - f(\mathbf{p}^*) < 0$. It follows that $A = \{\mathbf{x} : \mathbf{p}\cdot\mathbf{x} \geq f(\mathbf{p})$ for all $\mathbf{p}\}$, finishing part (2).

The first part of (3) follows from the Conjugate Duality Theorem. For the second part of (3), the uniqueness of the supergradient allows us to invoke Proposition 7.3.14.

The fact that $\mathbf{x}_i \in \partial^* f(\mathbf{x}_i)$ combines with Lemma 7.5.5 to yield the inequality in (4). Finally, if f is twice differentiable, $d\mathbf{x} = d^2 f$ which must be symmetric. Further, since f is concave, $d^2 f$ is also negative semi-definite. The last statement in (6) follows from Euler's Theorem. □

7.5.6 The Cost Theorem

One application is to cost functions. Given a production function f, define the *input set* $\mathcal{Z}(q)$ by $\mathcal{Z}(q) = \{\boldsymbol{z} \in \mathbb{R}^m_+ : f(\boldsymbol{z}) \ge q\}$. The input set is the upper contour set for f. It is the collection of input vectors that produce output at least as big as q. We can recast the cost minimization problem of a firm facing factor prices $\boldsymbol{w} \in \mathbb{R}^m$ with minimum output q, as finding $\inf\{\boldsymbol{w}\boldsymbol{\cdot}\boldsymbol{z} - \mathbb{I}_{\mathcal{Z}(q)}\} = \mathbb{I}^*_{\mathcal{Z}(q)}$. When the infimum is a minimum, the solutions, the conditional factor demands, are denoted $\boldsymbol{z}_q(\boldsymbol{w})$ or $\boldsymbol{z}(\boldsymbol{w}, q)$. The minimum cost to produce output q is $\mathbb{I}^*_{\mathcal{Z}(q)}(\mathbf{p})$ and is denoted $c(\boldsymbol{w}, q)$ or $c_q(\boldsymbol{w})$. We then have the following result.

Cost Theorem. *Let f be a continuous production function. Then:*

(1) *If $\boldsymbol{w} \gg 0$ and there is a $\hat{\boldsymbol{z}}$ with $f(\hat{\boldsymbol{z}}) \ge q$, there is a $\boldsymbol{z}(\boldsymbol{w}, q)$ solving the cost minimization problem for factor prices $\boldsymbol{w}$ and output q.*

(2) *The cost function $c_q(\boldsymbol{w})$ is concave, upper semicontinuous and homogeneous of degree one. In addition, c_q is continuous on $\operatorname{int}(\operatorname{dom} c_q)$ and weakly increasing in $\boldsymbol{w}$ and q.*

(3) *If f is concave, the cost function obeys $(c_q)^*(\boldsymbol{z}) = \mathbb{I}_{\mathcal{Z}(q)}(\boldsymbol{z})$. Thus $\mathcal{Z}(q) = \{\boldsymbol{z} : (c_q)^*(\boldsymbol{z}) = 0\}$. Moreover, $\mathcal{Z}(q) = \{\boldsymbol{z} : \boldsymbol{w}\boldsymbol{\cdot}\boldsymbol{z} \ge c(\boldsymbol{w}, q)$ for all $\boldsymbol{w}\}$.*

(4) Shephard's Lemma. *The factor demand correspondence obeys $\boldsymbol{z}_q(\boldsymbol{w}) = \partial^* c_q(\boldsymbol{w})$. Moreover, if there is a unique minimizer, the directional derivative of c_q exists at $\boldsymbol{w}$ and is given by $\mathcal{D}c_q(\boldsymbol{w}; \boldsymbol{v}) = \boldsymbol{z}_q(\mathbf{p})\boldsymbol{\cdot}\boldsymbol{v} = \boldsymbol{z}(\boldsymbol{w}, q)\cdot \boldsymbol{v}$.*

(5) *Law of Factor Demand. If $\boldsymbol{z}_i \in \boldsymbol{z}_q(\boldsymbol{w}_i)$ for $i = 0, 1$, then $(\boldsymbol{z}_1 - \boldsymbol{z}_0)\cdot(\boldsymbol{w}_1 - \boldsymbol{w}_0) \le 0$.*

(6) *The factor demand $\boldsymbol{z}_q(\boldsymbol{w})$ is homogeneous of degree zero.*

(7) *If $c_q(\boldsymbol{w})$ is $\mathcal{C}^2$, then $\boldsymbol{z}_q(\boldsymbol{w})$ is differentiable and the matrix $D\boldsymbol{z}_q(\boldsymbol{w}) = D^2 c_q(\boldsymbol{w})$ is symmetric and negative semi-definite. Moreover, $\big[D\boldsymbol{z}_q(\boldsymbol{w})\big]\boldsymbol{w} = \big[D^2 c_q(\boldsymbol{w})\big]\boldsymbol{w} = 0$.*

(8) *If f is quasiconcave then $\boldsymbol{z}(\boldsymbol{w}, q)$ is a convex set. If*

(9) *If f is homogeneous of degree $\gamma > 0$, then $c(\boldsymbol{w}, q)$ and $\boldsymbol{z}(\boldsymbol{w}, q)$ are homogeneous of degree $1/\gamma$ in q.*

(10) *If f is concave, then $c(\boldsymbol{w}, q)$ is convex in q. f is strictly quasiconcave, then $\boldsymbol{z}(\boldsymbol{w}, q)$ is a single point. In this case c_q has directional derivatives as in (4).*

Proof. Property (1) is of the Cost Theorem is Theorem 6.2.4.

Properties (2)-(7) follow immediately from the Support Function Theorem. Notice that $\mathcal{Z}(q)$ is closed and convex because f is concave and continuous.

The remaining properties, (8)–(10) were already shown in the Basic Cost Theorem of Chapter 6. □

There is a similar result for the expenditure function, which we will save for Chapter 8.

January 28, 2023