Topic 9: Support Functions

9.1 Support functions

The Separating Hyperplane Theorem 8.3.1 is the basis for a number of results concerning closed convex sets. Given any set A in $\boldsymbol{R}^{\mathrm{m}}$ its closed convex hull $\overline{\mathrm{co}}\, A$ is by definition the intersection of all closed convex sets that include A. But Theorem 8.3.4 sharpens this result to

$$\overline{\mathrm{co}}\, A = \bigcap\{H : A \subset H \text{ and } H \text{ is a closed half space}\}.$$

So an already closed convex set is the intersection of all the closed half spaces that include it.

The **support function** of a set A is a handy way to summarize all the closed half spaces that include A. There are two ways to define support functions, and I will give them names inspired by economics.[1]

The **cost function** μ_A of A is defined by

$$\mu_A(p) = \inf\{p \cdot x : x \in A\}.$$

The **profit function** π_A of A is defined by

$$\pi_A(p) = \sup\{p \cdot x : x \in A\}.$$

Clearly

$$\pi_A(p) = -\mu_A(-p).$$

We allow for the case that $\mu_A(p) = -\infty$ or $\pi_A(p) = \infty$. The set of points where the support function of a nonempty set is finite is a convex cone. Note that $\mu_\varnothing$ is the improper concave function $+\infty$. Also note that the infimum or supremum may not actually be attained even if it is finite. For instance, consider the closed convex set $A = \{(x, y) \in \boldsymbol{R}^2_{++} : xy \geqslant 1\}$, and let $p = (0, 1)$. Then $\mu_A(p) = 0$ even though $p \cdot (x, y) = y > 0$ for all $(x, y) \in A$. If A is compact, then of course μ_A and π_A are always finite, and the infimum and supremum are achieved as a maximum and minimum.

It is possible to restate Corollary 8.3.2 as follows.

9.1.1 Corollary *Let C be a nonempty closed convex subset of a Hilbert space. Assume that the point x does not belong to C. Then there exists a nonzero p such that*

$$p \cdot x < \mu_C(p)$$

[1] Fenchel [3] and Roko and I [1, p. 288] call the profit function the support function, while Mas-Colell, Whinston, and Green [4] use the cost function.

and letting $p' = -p$, we have

$$p' \cdot x > \pi_C(p').$$

Theorem 8.3.4 yields the following description of $\overline{\text{co}}\,A$ in terms of μ_A and π_A.

9.1.2 Theorem *For any set A in $\boldsymbol{R}^{\text{m}}$,*

$$\overline{\text{co}}\,A = \Big\{x \in \boldsymbol{R}^{\text{m}} : (\forall p \in \boldsymbol{R}^{\text{m}})\ [\,p \cdot x \geqslant \mu_A(p)\,]\Big\},$$

and

$$\overline{\text{co}}\,A = \Big\{x \in \boldsymbol{R}^{\text{m}} : (\forall p \in \boldsymbol{R}^{\text{m}})\ [\,p \cdot x \leqslant \pi_A(p)\,]\Big\},$$

Moreover, $\mu_A = \mu_{\overline{\text{co}}\,A}$ and $\pi_A = \pi_{\overline{\text{co}}\,A}$.

Proof: I shall just prove the results about the cost function μ. Observe that

$$C := \Big\{x \in \boldsymbol{R}^{\text{m}} : (\forall p \in \boldsymbol{R}^{\text{m}})\ [\,p \cdot x \geqslant \mu_A(p)\,]\Big\} = \bigcap_{p \in \boldsymbol{R}^{\text{m}}} \{p \geqslant \mu_A(p)\}$$

is an intersection of closed half spaces. By definition, if $x \in A$, then $p \cdot x \geqslant \mu_A(p)$, that is, $A \subset \{p \geqslant \mu_A(p)\}$. Thus by Theorem 8.3.4, $\overline{\text{co}}\,A \subset C$.

For the reverse inclusion, suppose $x \notin \overline{\text{co}}\,A$. By Corollary 9.1.1 there is a nonzero p such that $\mu_A(p) > p \cdot x$, so $x \notin C$.

To see that $\mu_A = \mu_{\overline{\text{co}}\,A}$ first note that $\mu_A \geqslant \mu_{\overline{\text{co}}\,A}$ since $A \subset \overline{\text{co}}\,A$. The first part of the theorem implies $\mu_{\overline{\text{co}}\,A} \geqslant \mu_A$. ∎

9.1.3 Theorem *Let A be a nonempty closed convex set in $\boldsymbol{R}^{\text{m}}$.*

Them the profit function π_A is a regular (proper and lower semicontinuous) convex and homogenous function on $\boldsymbol{R}^{\text{m}}$.

The cost function μ_A is a regular (proper and upper semicontinuous) concave and homogeneous function on $\boldsymbol{R}^{\text{m}}$.

Proof: Homogeneity of the cost and profit functions is obvious.

Each x defines a linear (and therefore both concave and convex) function ℓ_x via $\ell_x \colon p \mapsto p \cdot x$. Moreover each ℓ_x is continuous, and therefore both upper and lower semicontinuous.

Now $\mu_A = \inf_{x \in A} \ell_x$, so by Exercise 1.3.3 (4), it is concave and by Proposition A.8.4 it is upper semicontinuous. Similarly, π_A is convex and lower semicontinuous.

Let x belong to A (it is nonempty). Then for any p, $\pi_A(p) \geqslant p \cdot x$, so $\pi_A(p)$ is never $-\infty$. And $pi_A(0) = 0$, so $\operatorname{dom} \pi_A$ is nonempty. Therefore π_A is a proper convex function. Similarly, μ_A is a proper concave function. ∎

Theorem 9.1.2 asserts that we can recover the A (more precisely its closed convex hull) from the profit function π_A. Theorem 9.1.3 asserts that π_A is a homogeneous regular convex function. Suppose we take an arbitrary homogeneous regular convex function $\boldsymbol{R}^{\text{m}}$, is it the support function of some nonempty closed convex set? Yes. (And there is an analogous result for regular concave functions.)

9.1.4 Theorem *Let $f\colon \boldsymbol{R}^{\mathrm{m}} \to \boldsymbol{R}^{\sharp}$ be a homogeneous regular convex function. Then f is the profit function π_{C_f} of the nonempty closed convex set*

$$\begin{aligned} C_f &= \Big\{x \in \boldsymbol{R}^{\mathrm{m}} : (\forall p \in \boldsymbol{R}^{\mathrm{m}}) \ [\, p \cdot x \leqslant f(p)\,]\Big\} \\ &= \Big\{x \in \operatorname{dom} f : (\forall p \in \boldsymbol{R}^{\mathrm{m}}) \ [\, p \cdot x \leqslant f(p)\,]\Big\}. \end{aligned}$$

Proof: The first thing is to note that if $p \notin \operatorname{dom} f$, then $f(p) = \infty$, so $p \cdot x < f(p)$ for all x, so the two definitions of C_f agree.

Now we show that C_f is closed and convex. Since f is proper, $\operatorname{dom} f$ is a nonempty convex cone, so it contains 0 and by homogeneity $f(0) = 0$. Now $\{x : 0 \cdot x \leqslant 0\} = \boldsymbol{R}^{\mathrm{m}}$. For each nonzero $p \in \operatorname{dom} f$, the set $\{x : p \cdot x \leqslant f(p)\}$ is a closed hyperplane. Since C_f is the intersection of these sets, it is closed and convex.

The proof that C_f is nonempty is subtle, and is a byproduct of the following argument.

Fix some $\bar p \in \operatorname{dom} f$. If $x \in C_f$, then by definition $\bar p \cdot x \leqslant f(\bar p)$, so

$$\pi_{C_f}(\bar p) = \sup_{x \in C_f} \bar p \cdot x \leqslant f(\bar p).$$

We need to show that the supremum is actually equal to $f(\bar p)$. This means that for every $\varepsilon > 0$, we widh to find some $y \in C_f$ so that $\bar p \cdot y$ is within ε of $f(p)$, that is,

$$f(\bar p) - \varepsilon \leqslant \bar p \cdot y \leqslant f(\bar p).$$

Since f is proper and convex, its epigraph is nonempty and convex. Since f is homogeneous, its epigraph is a cone. And since f is lower semicontinuous, its epigraph is closed. Now the point $\big(\bar p, f(\bar p) - \varepsilon\big)$ does not belong to $\operatorname{epi} f$, which is nonempty, closed, and convex, so by the Strong Separating Hyperplane Theorem 8.3.1 we can separate it from the epigraph. That is, there exists some $(x, \lambda) \in \boldsymbol{R}^{\mathrm{m}+1}$ satisfying

$$(x, \lambda) \cdot \big(\bar p, f(\bar p) - \varepsilon\big) < (x, \lambda) \cdot (p, \beta) \quad \text{for all } (p, \beta) \in \operatorname{epi} f. \tag{1}$$

Expanding this and evaluating at $\beta = f(p)$, gives

$$\bar p \cdot x + \lambda f(\bar p) - \lambda\varepsilon < 0 \leqslant p \cdot x + \lambda f(p) \quad \text{for all } p. \tag{2}$$

Evaluating the right-hand side at $p = \bar p$ implies $-\lambda\varepsilon < 0$, so

$$\lambda > 0.$$

Taking just the left-hand inequality gives

$$\bar p \cdot x < \lambda\varepsilon - \lambda f(\bar p),$$

so dividing both sides by $-\lambda < 0$ reverses the inequality, and setting $y = -(1/\lambda)x$ gives

$$\bar{p} \cdot y > f(\bar{p}) - \varepsilon.$$

It remains to show that y belongs to C_f. Since $x = -\lambda y$, the right-hand side of (2) says that

$$0 \leqslant -\lambda p \cdot y + \lambda f(p) \quad \text{for all } p,$$

which, since $\lambda > 0$, implies $f(p) \geqslant p \cdot y$ for all p. In other words, $y \in C_f$. This does two things. It shows that $f = \pi_{C_f}$ and also that C_f is nonempty. ∎

The technique of separating a point from an epigraph will appear later on in connection with regular convex functions and subgradients.

9.2 Sublinear functions

9.2.1 Definition *A function f from a convex cone C in a real vector space into $\boldsymbol{R}^\sharp$ is*

positively homogeneous of degree 1 *if for every vector $x \in C$ and every real $\lambda > 0$,*

$$f(\lambda x) = \lambda f(x).$$

*We usually shorten this by saying simply that f is **homogeneous**.*

subadditive *if for all vectors x and y in C,*

$$f(x + y) \leqslant f(x) + f(y).$$

superadditive *if for all vectors x and y in C,*

$$f(x + y) \geqslant f(x) + f(y).$$

sublinear *if it is both homogeneous and subadditive.*

9.2.2 Remark

- There are other notions of homogeneity. More generally, f is positively homogeneous of degree k if for every $\lambda > 0$, we have $f(\lambda x) = \lambda^k f(x)$. If I ever mean anything other than of homogeneity of degree one, I will make it explicit.
- By these definitions we ought to say that f is **superlinear** if it is both homogeneous and superadditive, but I've never heard the term.
- Note the definition of homogeneity restricts attention to $\lambda > 0$, not $\lambda \geqslant 0$. This avoids the question of deciding how to interpret $0 \cdot \infty$ for extended real valued functions. (If $\lambda > 0$, then $\lambda \cdot \infty = \infty$.)

- Note that for a homogeneous function defined at 0, we have $f(0) = f(\lambda 0) = \lambda f(0)$ for any $\lambda > 0$, so $f(0) = 0$.
- So if f is homogenous on the punctured cone $C \setminus \{0\}$, where C is a true cone, it can be extended to be homogenous on all of C simply by setting $f(0) = 0$.
- Note that f defined by $f(0) = 0$, and $f(x) = \infty$ for $x \neq 0$ is homogeneous.
- If f is a homogeneous function on a cone C, we can extend it to be homogeneous on the entire vector space by setting $f(x) = \infty$ for any nonzero x not in C.

9.2.3 Exercise A homogeneous function is subadditive if and only it is convex. It is superadditive if and only if it is concave.

The epigraph of a sublinear function is a convex cone. The hypograph of a homogeneous concave function is a convex cone. □

9.3 Gauge functions

See Arrow and Hahn [2] for applications.

9.3.1 Definition *The* ***gauge function****, or more simply the* ***gauge****,* p_A *of a subset* A *of a vector space is defined by*

$$p_A(x) = \inf\{\alpha > 0 : x \in \alpha A\},$$

where, you may recall, $\inf \varnothing = \infty$.

Note that the gauge of A is always nonnegative.

9.3.2 Example

- The most important example of a gauge is a norm. If U is the closed unit ball in a normed space,

$$U = \{x \in X : \|x\| \leqslant 1\},$$

then

$$p_U(x) = \|x\|.$$

- The gauge of $\boldsymbol{R}$ is always zero. So is the gauge of the integers. Thus two distinct sets may have the same gauge.

□

We shall see in a moment (Lemma 9.3.6) that gauges are the *nonnegative* sublinear functions.

9.3.3 Definition

- *A set A is* ***star-shaped about zero*** *if $x \in A$ implies that the line segment $[0, x]$ is a subset of A.*

- *A set A is* ***absorbing*** *if it is star-shaped about zero and its gauge p_A is everywhere finite. This means that for every $x \in X$, there is some λ_x so that for $0 < \lambda \leqslant \lambda_x$, we have $\lambda x \in A$.*

- *A set A is* ***circled****, or* ***balanced****, if for each $x \in A$ the line segment $[-x, x]$ lies in A.*

9.3.4 Remark

- A balanced set is star-shaped about zero.

- Circled sets have the property that the gauge p_A satisfies $p_A(x) = p_A(-x)$.

- The unit ball in any normed space is convex, absorbing, and circled.

9.3.5 Definition

- *A* ***seminorm*** *is a subadditive function $p \colon X \to \boldsymbol{R}$ (not $\boldsymbol{R}^\sharp$!) on a vector space satisfying*
$$p(\alpha x) = |\alpha| p(x)$$
for all $\alpha \in \boldsymbol{R}$ and all $x \in X$.

- *A seminorm p that satisfies $p(x) = 0$ if and only if $x = 0$ is called a* ***norm****.*

9.3.6 Lemma *For a nonnegative extended real function $f \colon X \to \boldsymbol{R}^\sharp$ on a vector space, we have the following.*

1. *f is homogeneous if and only if it is a gauge function, in which case it is the gauge of the set*
$$U_f = \{x \in X : f(x) \leqslant 1\}.$$

2. *f is sublinear if and only if it is the gauge of U_f and U_f is convex.*

3. *f is a seminorm if and only if it is the gauge of U_f and U_f is circled, convex, and absorbing. The set U_f is called the* ***unit ball*** *of the seminorm.*

Proof: First observe that for any set A its gauge has the property that

$$x \in \alpha A \implies p_A(x) \leqslant \alpha.$$

Next observe that if A is star-shaped about zero, the converse is true, so

$$x \in \alpha A \iff p_A(x) \leqslant \alpha. \tag{3}$$

Now note that if f is nonnegative and homogeneous, then the set

$$U_f = \{x \in X : f(x) \leqslant 1\}$$

is star-shaped about zero.

(1) We want to show that if f is nonnegative and homogeneous, then it is the gauge p_{U_f} of the star-shaped set U_f. From (3) and the definition of U_f we have

$$p_{U_f}(x) \leqslant \alpha \iff x \in \alpha U_f \iff f(x) \leqslant \alpha,$$

so $f = p_{U_f}$

(2) We want to show that if f is nonnegative and homogeneous, then f is subadditive if and only U_f is convex. We already know that $f = p_{U_f}$.

So first assume that U_f is convex. Let $\alpha, \beta > 0$ satisfy $x \in \alpha U_f$ and $y \in \beta U_f$. Then $x + y \in \alpha U_f + \beta C = (\alpha + \beta)C$, so $p_{U_f}(x + y) \leqslant \alpha + \beta$. Taking infima yields $p_{U_f}(x + y) \leqslant p_{U_f}(x) + p_{U_f}(y)$, so $p_{U_f} = f$ is subadditive.

For the converse, assume that f is subadditive. We need to show that U_f is convex. So let $x, y \in U_f$ and let $0 < \lambda < 1$. If f is subadditive, then

$$f((1-\lambda)x + \lambda y) \leqslant f\Big((1-\lambda)x\Big) + f(\lambda y) = (1-\lambda)f(x) + \lambda f(y) \leqslant 1,$$

where the first inequality is subadditivity, the equality is homogeneity, and the last inequality is just $x, y \in U_f$. But this just asserts that $(1-\lambda)x + \lambda y \in U_f$ as desired.

(3) The proof is very similar to the above and is left as an exercise. ∎

9.4 Gauge functions and support functions

Profit functions are sublinear (Theorem 9.1.3), and by Lemma 9.3.6 every nonnegative sublinear function is a gauge. Thus every nonnegative profit function is also a gauge. How doe we guarantee that π_A is nonnegative? One simple way is to require that $0 \in A$. Then $\sup x \in A p \cdot x \geqslant p \cdot \underbrace{0}_{\in A} = 0$.

In fact, if $A \subset \boldsymbol{R}^{\mathrm{m}}$ is a closed convex set that contains 0, then by Lemma 9.3.6 its profit function π_A is the gauge of

$$\{p \in \boldsymbol{R}^{\mathrm{m}} : \pi_A(p) \leqslant 1\} = \Big\{p : (\forall x \in A)\ [p \cdot x \leqslant 1]\Big\}.$$

This suggest the following definition.[2]

9.4.1 Definition *Given a nonempty set A in $\boldsymbol{R}^{\mathrm{m}}$, define the **polar** of A, denoted $A^{\bigcirc}$ by*

$$A^{\bigcirc} = \Big\{p \in \boldsymbol{R}^{\mathrm{m}} : (\forall x \in A)\ [p \cdot x \leqslant 1]\Big\}.$$

[2] This definition is what Roko and I called the **one-sided polar** in [1, p. 215].

9.4.2 Exercise The polar $A^{\bigcirc}$ is a closed convex set that contains 0. □

9.4.3 Exercise (The polar of cone) If C is cone, then $C^{\bigcirc} = C^*$, where $C^* = \left\{p : (\forall x \in C)\ [p \cdot x \leqslant 0]\right\}$ is the dual cone of C. □

9.4.4 Bipolar Theorem *For a nonempty set A,*

$$A^{\bigcirc\bigcirc} = \overline{\mathrm{co}}(A \cup \{0\}).$$

Consequently, if C is a closed convex set that contains 0, then $C = C^{\bigcirc\bigcirc}$.

Proof: By definition

$$A^{\bigcirc\bigcirc} = \left\{x : \left(\forall p \in A^{\bigcirc}\right)\ [p \cdot x \leqslant 1]\right\}.$$

Now if $x \in A$ and $p \in A^{\bigcirc}$ we have $p \cdot x \leqslant 1$, so $A \subset A^{\bigcirc\bigcirc}$. Thus by Exercise 9.4.2, we have

$$\overline{\mathrm{co}}(A \cup \{0\}) \subset A^{\bigcirc\bigcirc}.$$

For the reverse inclusion, we shall prove that $y \notin \overline{\mathrm{co}}(A \cup \{0\}) \implies y \notin A^{\bigcirc\bigcirc}$. So assume $y \notin \overline{\mathrm{co}}(A \cup \{0\})$. By the Strong Separating Hyperplane Theorem 8.3.1 there exists a nonzero p and a real α with

$$p \cdot y > \alpha > p \cdot x \tag{$\star$}$$

for all $x \in \overline{\mathrm{co}}(A \cup \{0\})$. Taking $x = 0$ we see that, $\alpha > 0$, so we can multiply the inequality ($\star$) by $1/\alpha$. Setting $p' = (1/\alpha)p$, we have $p' \cdot y > 1 > p' \cdot x$ for all $x \in A$. So $p' \in A^{\bigcirc}$. But $p' \cdot y > 1$, so $y \notin A^{\bigcirc\bigcirc}$. This completes the proof. ∎

9.4.5 Proposition *Let C be a closed convex set that contains 0. Then:*

1. $\pi_C = p_{C^\circ}$.
2. $p_C = \pi_{C^\circ}$.

Proof: This is not really hard. I just have a hard dealing with suprema. I'm never sure what's obvious and what isn't so I will probably spell this out in too much detail.

Tighten this up. Worry about infinite values.

1. By Exercise 9.4.2, $C^{\bigcirc}$ is convex and contains 0. Then for any p, the set $\{\alpha : p \in \alpha C^{\bigcirc}\}$ is an *interval* with lower bound $p_{C^\circ}(p)$. So,

$$\begin{aligned} \beta > p_{C^\circ}(p) &\implies p \in \beta C^{\bigcirc}, \\ &\implies (\forall x \in C)\ [p \cdot x \leqslant \beta] \\ &\implies \quad \pi_C(p) \leqslant \beta. \end{aligned}$$

That is, $\beta > p_{C^\circ}(p) \implies \beta \geqslant \pi_C(p)$. So

$$\pi_C(p) \leqslant p_{C^\circ}(p).$$

For the reverse inequality,

$$\begin{aligned} \beta > \pi_C(p) &\implies (\forall x \in C)\ [\, p \cdot x \leqslant \pi_C(p) < \beta \,] \\ &\implies p \in \beta C^{\bigcirc} \\ &\implies \beta \geqslant p_{C^\circ}(p). \end{aligned}$$

Thus

$$\pi_C(p) \geqslant p_{C^\circ}(p).$$

These two inequalities imply that

$$\pi_C = p_{C^\circ}.$$

2. From part (1) applied to $C^{\bigcirc}$ we see that $\pi_{C^\circ} = p_{C^{\circ\circ}} = p_C$, where the last equality is the Bipolar Theorem.

∎

References

[1] C. D. Aliprantis and K. C. Border. 2006. *Infinite dimensional analysis: A hitchhiker's guide*, 3d. ed. Berlin: Springer–Verlag.

[2] K. J. Arrow and F. H. Hahn. 1971. *General competitive analysis*. San Francisco: Holden–Day.

[3] W. Fenchel. 1953. Convex cones, sets, and functions. Lecture notes, Princeton University, Department of Mathematics. From notes taken by D. W. Blackett, Spring 1951.

[4] A. Mas-Colell, M. D. Whinston, and J. R. Green. 1995. *Microeconomic theory*. Oxford: Oxford University Press.