Duality

for The New Palgrave Dictionary of Economics, 2nd ed.

Lawrence E. Blume

Headwords: CONVEXITY, DUALITY, LAGRANGE MULTIPLIERS, PARETO EFFICIENCY, QUASI-CONCAVITY

1 Introduction

The word 'duality' is often used to invoke a contrast between two related concepts, as when the informal, peasant, or agricultural sector of an economy is labeled as dual to the formal, or profit-maximizing sector. In microeconomic analysis, however, 'duality' refers to connections between quantities and prices which arise as a consequence of the hypotheses of optimization and convexity. Connected to this duality are the relationship between utility and expenditure functions (and profit and production functions), primal and dual linear programs, shadow prices, and a variety of other economic concepts. In most textbooks, the duality between, say, utility and expenditure functions, arises from a sleight-of-hand with the first order conditions for optimization. These dual relationships, however, are not naturally a product of the calculus; they are rooted in convex analysis and, in particular, in different ways of describing a convex set. This entry will lay out some basic duality theory from the point of view of convex analysis, as a remedy for the microeconomic theory textbooks the reader may have suffered.

2 Mathematical Background

Duality in microeconomics is properly understood as a consequence of convexity assumptions, such as laws of diminishing marginal returns. In microeconomic models, many sets of interest are closed convex sets. The mathematics here is surveyed in the CONVEX PROGRAMMING entry. The urtext for this material is Rockafellar (1970).

Closed convex sets can be described in two ways; by listing their elements, the 'primal' description of the set, and by listing the closed half-spaces that contain it. A closed (upper) half-space in $\mathbf{R^n}$ is a set of the form $h_{pa} = \{x : p \cdot x \geq a\}$, where p is another n-dimensional vector, a is a number and $p \cdot x$ is the inner product. The vector p is the *normal vector* to the half-spaces h_{pa}. Geometrically speaking, this is the set of points lying on or above the line $p \cdot x = a$. The famous separation theorem of convex sets implies that every closed convex set is the intersection of the half-spaces containing it.

Suppose that C is a closed convex set, and that p is a vector in $\mathbf{R^n}$. How do we find all the numbers a such that $C \subset h_{pa}$? If there is an $x \in C$ such that $p \cdot x < a$, then a is too big. So the natural candidate is $w = \inf_{x \in C} p \cdot x$. If $a > w$ there will be an $x \in C$ such that $p \cdot x < a$ on the other hand, if $a < w$, then $p \cdot x > a$ for all $x \in C$. So the half-spaces h_{pa} for $a \leq w$ are the closed half-spaces containing C.

This construction can be applied to functions: A concave function on $\mathbf{R^n}$ is an $[-\infty, \infty)$-valued function f such that the *hypograph* of f, the set $\operatorname{hypo} f = \{(x, a) \in \mathbf{R^{n+1}} : a \leq f(x)\}$, is convex. If $\operatorname{hypo} f$ is closed, f is said to be *upper semi-continuous* (usc). The *domain* $\operatorname{dom} f$ of concave f is the set of vectors in $\mathbf{R^n}$ for which f is finite-valued. Concave (and convex) functions are very well-behaved on the *relative interiors* of their effective domains. The relative interior of a convex set is the interior relative to the smallest affine set containing it (see CONVEX PROGRAMMING), and on $\operatorname{ri} \operatorname{dom} f$, f (concave or convex) is continuous.

Suppose that f is usc. The minimal level a such that $h_{(p,-1)a}$, the hyperplane in $\mathbf{R^{n+1}}$ with normal vector $(p, -1)$, contains $\operatorname{hypo} f$ is $f^*(p) \equiv \inf_x p \cdot x - f(x)$. Why the normal vector $(p, -1)$? Because the graph of the affine function $x \mapsto f^*(p) + px$ is a tangent line to f; the graph of f lies everywhere beneath it, and no other line with the same slope and a smaller intercept has this property. The function $f^*(p)$ is the *Fenchel transform* or *conjugate* of f, and is traditionally denoted f^*. The construction of the preceding paragraph can be done just this way: The *concave indicator function* of a convex set C is the function $\delta_C(x)$ which is 0 on C and $-\infty$ otherwise, and $\delta^*_C(p) = \inf_{x \in C} p \cdot x$. For any function f, not necessarily usc or concave, the Fenchel transform f^* is usc and concave. If f is in fact both usc and concave, then $f^{**} = f$. This fact is known as the conjugate duality theorem.

Convex functions with range $(-\infty, \infty]$ are treated identically. The function f is convex if and only if $-f$ is concave, but the definitions are handled slightly differently in order to preserve the intuition just described. The set $\operatorname{epi} f = \{x, a) : a > f(x)\}$, and $f^*(p) = \sup_x p \cdot x - f(x)$. The convex Fenchel transform is defined differently: $f^*(p) = \sup_x p \cdot x - f(x)$. The convex indicator function of a convex set C is the function $\delta^C(x)$ which is 0 on C and $+\infty$ otherwise; its (convex) conjugate is $\delta^{C*}(p) = \sup_x p \cdot x$. These facts are discussed in the CONVEX PROGRAMMING entry.

If concave functions have tangent lines, then they must have something like gradients. A vector p is a *subgradient* of f at x if $f(x) + p \cdot (y - x) \leq f(y)$. If f has a unique subgradient at x, then f is differentiable at x and $p = \nabla f(x)$, and conversely. But the subgradient need not be unique: The set $\partial f(x)$ of subgradients at x is the *subdifferential* of f at x. The *domain* of f, $\operatorname{dom} f$, is the set of x such that $f(x) > -\infty$. The subdifferential is non-empty for all x in its *relative interior*. It follows from the definition of concavity (and is proved in CONVEX OPTIMIZATION that the subdifferential correspondence is *monotonic*: if $p \in \partial f(x)$ and $q \in \partial f(y)$, then $(p - q) \cdot (x - y) \leq 0$. If f is convex, then the inequality is reversed, and $(p - q) \cdot (x - y) \geq 0$. Finally, suppose f is usc and concave. Then so is its conjugate f^*, and their subdifferentials have an inverse relationship: $p \in \partial f(x)$ if and only if $x \in \partial f^*(p)$.

3 Cost, Profit and Production

In the theory of the firm, profit functions and cost functions are alternative ways of describing the firms' technology choices. A technology is described by a set of vectors F in $\mathbf{R}^{\mathbf{N}}$. Each vector $z \in F$ is an input-output vector. We adopt the convention that negative coefficients correspond to input quantities and positive quantities correspond to outputs. Suppose that the first L goods are inputs and the last $M = N - L$ are outputs, so that $F \subset \mathbf{R}^{\mathbf{L}}_{-} \times \mathbf{R}^{\mathbf{M}}_{+}$.It is convenient to assume free disposal, so that if $(x, y) \in F$, and both $x' \leq x$ and $y' \leq y$ (more input and less output), then $(x', y') \in F$. Two important dual representations of the technology are the cost and profit functions. The profit function is $\pi(p, w) = \sup_{(x,y) \in F} p \cdot y + w \cdot x$ for $p \in$ and $w \in \mathbf{R}^{\mathbf{L}}$, which is the conjugate of the convex indicator function of F. The cost function too can be obtained through conjugacy. The set $F(y) = \{x : (x, y) \in F\}$ is the set of all input bundles that produce y. Then $C(y, w) = -\sup_{x \in F(y)} w \cdot x$, that is, $C(y, \cdot) = -\delta^{F(y)*}$.

Immediately the properties of the Fenchel transform imply that $\pi(p, w)$ is convex in its arguments and $C(y, w)$ is concave in w, the profit function is lsc and the cost function is usc. (This implies that both functions are continuous on the relative interior of their effective domains.) Cost and profit functions are also linear homogeneous. Doubling all prices doubles both costs and revenues. Cost is also monotonic. If $w'_l \leq w_l$ for every input l, then $C(y, w') \leq C(y, w)$ and if $w'_l < w_l$ for all l, then $C(y, w') < C(y, w)$.

The point of duality is that if the technology is closed and convex, then cost profit functions each characterize the technology F. The conjugate duality theorem (see CONVEX PROGRAMMING) implies that $\pi^*(x, y) \equiv \delta^{F**}(x, y) = \delta^F(x, y)$, the convex indicator function of F:

$$\sup_{(p,w) \in \mathbf{R^N}} p \cdot x + w \cdot y - \pi(p, w) = \begin{cases} 0 & \text{if } (x, y) \in F, \\ +\infty & \text{otherwise.} \end{cases}$$

If F is closed and convex, then each $F(y)$ is convex. If F is closed then $F(y)$ will also be closed. Then $\delta^{F(y)}$ is concave and usc, so $\sup_{w \in \mathbf{R}^\mathbf{L}_+} w \cdot x + C(y, w) = \sup_{w \in \mathbf{R}^\mathbf{L}_+} w \cdot x \delta^{F(y)*}(w) = \delta^{F(y)}(x)$.

Hotelling's lemma is a famous result of duality theory. It says that the net supply function of good i is the derivative of the profit function with respect to the price of good i. The usual proof is via the envelope theorem: The marginal change in profits from a change in price p is the quantity of good i times the change in the price plus the price of all goods times the changes in their respective quantities. But the quantity changes are second-order because the quantities solve the profit maximization first order conditions, that price times the marginal change in quantities in technologically feasible directions is 0. Every advanced microeconomics text proves this. A result like this is true whenever the technology is convex, even if the technology is not smooth.

The convex version of Hotelling's lemma is a consequence of the inversion property of subdifferentials for concave and convex f; that $p \in \partial f(x)$ if and only if $x \in \partial f^*(p)$. See CONVEX PROGRAMMING for a brief discussion.

Hotelling's Lemma: $(x, y) \in \partial \pi(p, w)$ if and only if (x, y) is profit-maximizing at prices (p, w).

Hotelling's lemma is quickly argued. If $(x, y) \in \partial \pi(p, w) = \partial \delta^{F*}(p, w)$, then $(p, w) \in \partial \delta^{F**}(x, y) = \partial \delta^F(x, y)$. Then $\delta^F(x, y) + (p, w) \cdot \big((x', y') - (x, y)\big) \leq \delta^F(x', y')$ for all (x', y'). This implies that $x \in F$ and furthermore that $(p, w) \cdot \big((x', y') - (x, y)\big) \leq 0$ for all $(x, y) \in F$, in other words, that (x, y) is profit-maximizing at prices (p, w). Conversely, suppose that (x, y) is profit maximizing at prices (p, w). Then (p, w) satisfies the subgradient inequality of δ^F at (x, y), and so $(p, w) \in \partial \delta^F$. Consequently, $(x, y) \in \partial \delta^{F*}(p, w) \equiv \partial \pi(p, w)$.

The textbook treatment of duality observes that, if net supply is the first derivative of the profit function, then the own-price derivative of net supply must be the second own-partial derivative of profit with respect to price, and convexity of the profit function implies that this partial derivative should be positive, so net supply is increasing in price. The same fact follows in the convex framework from the monotonicity properties of the subgradients. Suppose that (w, p) and (w', p') are two price vectors, and suppose that (x, y) and (x', y') are two profit-maximizing production plans corresponding to the two price vectors. Then $(w - w', p - p')(x - x', y - y') \geq 0$. If the two price vectors are identical for all prices but, say, $p_k \neq p'_k$, then $(p_k - p'_k)(y_k - y'_k) \geq 0$, and net supply is non-decreasing in price. As with net supplies, some comparative statics of conditional factor demand with respect to input price changes follows from the monotonicity property of subgradients.

Another implication of profit function convexity and (twice continuous) differentiability is symmetry of the derivatives of net supply:

$$\frac{\partial y_k}{\partial p_l} = \frac{\partial^2 \pi}{\partial p_k \partial p_l} = \frac{\partial^2 \pi}{\partial p_l \partial p_k} = \frac{\partial y_l}{\partial p_k}$$

The convex analysis version of this is that for any finite sequences of goods $i, \ldots, l$,

$$p_i \cdot (y_j - y_i) + p_j \cdot (y_k - y_j) + \cdots + p_l \cdot (y_i - y_l) \leq 0.$$

This requirement, which has a corresponding expression in terms of differences in prices, is called *cyclic monotonicity*. All subdifferential correspondences are cyclicly monotone. The connection with symmetry is not obvious, but it helps to know that Rockafellar (1974) leaves as an exercise (and so do we) that cyclic monotonicity is a property of a linear transformation corresponding to an $n \times n$ matrix M if and only if M is symmetric and positive semi-definite. Monotonicity is cyclic monotonicity for sequences of length 2.

The other famous result in duality theory for production is Shephard's lemma, which does for cost functions what Hotelling's lemma does for profit functions: Conditional input demands are the derivatives of the cost functions. This is demonstrated in the same way, since the cost function and the indicator function for the set of inputs from which y is produceable are both convex and have closed hypographs.

4 Utility and Expenditure Functions

A quasi-concave utility function U defined on the commodity space $\mathbf{R}^{\mathbf{n}}_{+}$ has upper contour sets, the sets R_u of consumptions bundles which have utility at least u, which are convex. If u is usc, these sets are closed as well.

The *expenditure function* gives for each utility level u and price vector p the minimum cost of realizing utility u at prices p: $e(p, u) = \inf\{p \cdot x : u(x) \geq u\}$. If the infimum is actually realized at a consumption bundle x, then x is the *Hicksian* or *compensated real income demand*.

In terms of convex analysis, $e(p,u)$ is the conjugate of the concave indicator function $\phi_u(x)$ of the set $R(u) = \{x : U(x) \geq u\}$, that is, $e(p,u) \equiv \phi_u^*(p)$. Thus $e(p,u)$ will be usc and concave in p for each u. The expenditure function is also linearly homogeneous in prices. If prices double, then the least cost of achieving u will double as well.

The duality of utility and expenditure functions is that each can be derived from the other; they are alternative characterizations of preference. Since the concave indicator function $\phi_u(x)$ is closed and convex, $e(\,\cdot\,,u)^* = \phi_u(x)$. For fixed u, the Fenchel transform of the expenditure function is the concave indicator function of $R(u)$; $\inf_p p \cdot x - e(p,u)$ is 0 if $U(x) \geq u$ and $-\infty$ otherwise. If $x \in R(u)$ then the cost of x at any price p can be no less than the minimum cost necessary to achieve utility u. The gap between the cost of x and the cost of utility level u is is made by taking ever smaller prices, and so its minimum is 0. Suppose that x is not in $R(u)$. The separation theorem for convex sets says there is a price p such that $p \cdot x < \inf_{y \in R(u)} p \cdot y$; there is a price at which x is cheaper than the cost of u. Now, by taking ever larger multiples of p, the magnitude of the gap can be made arbitrarily large, and so the value of the conjugate is $-\infty$. Thus the conjugate is the concave indicator function of $R(u)$.

Among the most useful consequence of the duality between utility and expenditure functions is the relationship between derivatives of the expenditure function and the Hicksian, or compensated demand. Hicksian demand. The compensated demand at prices p and utility u are those consumption bundles in $R(u)$ which minimize expenditure at prices p. This result is just Shephard's lemma for expenditure functions:

Hicks-Compensated Demand: Consumption bundle x is a Hick's compensated consumption bundle at prices p if and only if $x \in \partial_p e(p,u)$. Furthermore, if x is demanded at prices p and utility u, and y is demanded at prices q and the same utility u, then $(p-q) \cdot (x-y) \leq 0$.

The downward-sloping property just restates the monotonicity property of the subdifferential correspondence. For the special case of changes in a single price, the statement is that demand is non-increasing in its own price.

5 Equilibrium and Optimality

The equivalence between Pareto optima and competitive equilibria can also be viewed as an expression of duality. When preferences have concave utility representations, *quasi-equilibrium* emerges from Lagrangean duality. Quasi-equilibrium entails feasibility, profit maximization, and expenditure minimization rather than utility maximization. That is, each traders consumption allocation is expenditure minimizing for the level of utility it achieves. The now traditional route of Arrow (1952) and Debreu (1951) to the second welfare theorem first demonstrates that a Pareto-optimal allocation can be regarded as a quasi-equilibrium for an appropriate set of prices. Under some additional conditions, the quasi-equilibrium is in fact an competitive equilibrium, wherein utility maximization on an appropriate budget set replaces expenditure minimization. Our concern here is with the first step on this path.

Suppose that each of I individuals has preferences represented by a concave utility function on $\mathbf{R}^{\mathbf{N}}_{+}$, and that production is represented, as in section 3, by a closed and convex set F of feasible production plans. Suppose that $0 \in F$ (it is possible to produce nothing) and that the aggregate endowment e is strictly positive. Assume too that there is free disposal in production. Every Pareto optimum is the maximum of a Bergson-Samuelson social welfare function of the form $\sum_i \lambda_i u_i$ defined on the set of all consumption allocations. An allocation is a vector (x, y) where $x \in \mathbf{R}^{\mathbf{NI}}_{+}$ is a consumption allocation, a consumption bundle for each individual, and y is a production plan. The allocation is *feasible* if $y \in F$ and $y + e - \sum_i x_i \geq 0$. A Lagrangean for this convex program is

$$L(x,y,p) = \begin{cases} \sum_i u_i(x_i) + p \cdot (y + e - \sum_i x_i) & \text{if } x \in \mathbf{R}^{\mathbf{NI}}_{+},\, y \in F \text{ and } p \in \mathbf{R}^{\mathbf{L}}_{+}, \\ +\infty & \text{if } x \in \mathbf{R}^{\mathbf{NI}}_{+},\, y \in F \text{ and } p \notin \mathbf{R}^{\mathbf{L}}_{+}, \\ -\infty & \text{otherwise,} \end{cases}$$

where p is the vector of Lagrange multipliers for the L goods constraints.

The possibility of 0 production and the strict positivity of the aggregate endowment guarantee that the set of feasible solutions satisfies Slater's condition, and so a saddle point (x^*, y^*, p^*) exists; that is, $\sup_{x,y} L(x, y, p^*) \leq L(x^*, y^*, p^*) \leq L(x^*, y^*, p)$ for all $x \in \mathbf{R}^{\mathbf{NI}}$, $y \in F$ and $p \in \mathbf{R}^{\mathbf{L}}$. Then (x^*, y^*) is Pareto optimal and p^* solves the dual problem $\min_p \sup_{x,y} L(x, y, p)$. The interpretation of (x^*, y^*, p^*) as a quasi-equilibrium comes from examining the dual problem. The dual problem can be rewritten as

$$\begin{aligned} \inf_{p \in \mathbf{R}^{\mathbf{L}}_{+}} \sup_{x \in \mathbf{R}^{\mathbf{NI}}_{+}, y \in F} L(x,y,p) &= \inf_{p \in \mathbf{R}^{\mathbf{L}}_{+}} \sup_{x \in \mathbf{R}^{\mathbf{NI}}_{+}, y \in F} \sum_i u_i(x_i) + p \cdot (y + e - \sum_i x_i) \\ &= \inf_{p \in \mathbf{R}^{\mathbf{L}}_{+}} \sum_i \sup_{x_i \in \mathbf{R}^{\mathbf{L}}_{+}} \{\lambda_i u_i(x_i) - p \cdot x_i\} + \sup_{y \in F} p \cdot y. \end{aligned} \tag{1}$$

In the dual problem, the Lagrange multipliers can be thought of as goods prices. The second welfare theorem interprets the optimal allocation as an equilibrium allocation using the Lagrange multipliers as equilibrium prices. To see this, look at the second line of (1). At prices p, a production plan is chosen from y to maximize profits $p \cdot y$, so the value of this term is $\pi(p)$. Each consumer is asked to solve

$$\begin{aligned}\max_i \lambda_i u_i(x_0) - p \cdot x &= -\min p \cdot x - \lambda_i u_i(x_i) \\ &= \lambda_i u_i^* - \min p \cdot x - \lambda_i \big(u_i(x_i) - u_i^*\big)\end{aligned}$$

where $u_i^* = u_i(x_i^*)$. The term being minimized is the Lagrangean for the problem of expenditure minimization, and so x_i^* is the Hicksian demand for consumer i at prices p and utility level $u_i^* = u_i(x_i^*)$. Finally, the optimal allocation is feasible, and so (x^*, y^*, p^*) is a quasi-equilibrium.

Given the observation about expenditure minimization, the saddle value of the Lagrangean is

$$\sum_i \lambda_i u_i^* - e_i(p^*, u_i^*) + \pi(p^*)$$

The planner chooses prices to minimize net surplus, which is the sum of profits from production and the excess of total Bergson-Samuelson welfare less the cost of the consumption allocation.

6 Historical Notes

Duality ideas appeared very early in the marginal revolution. Antonelli, for instance, introduced the indirect utility function in 1886. The modern literature begins with Hotelling (1932), who provided us with Hotelling's lemma and cyclic monotonicity. Shephard (1953) was the first modern treatment of duality, making use of notions such as the support function and the separating hyperplane theorem.

The results on consumer and producer theory are surveyed more extensively in Diewert (1981), who also provides a guide to the early literature. In its focus on Fenchel duality, this review has not even touched on the duality between direct and indirect aggregators, such as utility and indirect utility, and topics that would naturally accompany this subject such as Roy's identity. Again, this is admirably surveyed in Diewert (1981).

References

Arrow, K. J. (1952): "An Extension of the Basic Theorems of Classical Welfare Economics," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, ed. by J. Neyman, pp. 507–32.

DEBREU, G. (1951): “The Coefficient of Resource Utilization,” *Econometrica*, 19, 273–292.

DIEWERT, W. E. (1981): “The Measurement of Deadweight Loss Revisited,” *Econometrica*, 49(5), 1225–1244.

HOTELLING, H. (1932): “Edgeworth’s Taxation Paradox and the Nature of Demand and Supply,” *Journal of Political Economy*, 40, 577–616.

ROCKAFELLAR, R. T. (1970): *Convex Analysis*. Princeton University Press, Princeton NJ.

——— (1974): *Conjugate Duality and Optimization*. Society of Industrial and Applied Mathematics, Philadelphia.

SHEPHARD, R. W. (1953): *Cost and Production Functions*. Princeton University Press, Princeton NJ.