Cambridge University Press
978-0-521-33601-7 - Duality and Modern Economics
Richard Cornes
Excerpt
More information

PART I

Some background

CHAPTER 1

Some formal preliminaries: An informal treatment

The theme of this book is both simple and attractive. It is that the optimizing models that are the daily bread and butter of both theoretical and applied economists can be formulated in a number of different ways and that the investment of a little time and effort in becoming familiar with alternative formulations, together with a little ingenuity in exploiting some simple implications of optimality, can tremendously simplify their formal analysis. The simplicity, and indeed elegance, that can thereby be achieved allows us much more direct insight into the economic intuition that lies behind comparative static and other properties than is afforded by the old-fashioned and rather mechanical procedures that have dominated textbooks in the past. For example, the standard approach to consumer theory once was simply to maximize utility, defined over quantities consumed, subject to the budget constraint. The analysis then proceeded through the tedious and unenlightening manipulation of first- and second-order optimality conditions, through a dense thicket of bordered Hessians, and on to statements about demand behavior whose intuitive explanation, though often simple, was in no way reflected in the tortuous mathematics. We now know that there are at least three alternative ways of stating the consumer's problem and that it pays to think about which one to use before commencing a piece of analysis. One's choice of starting point will be influenced by the nature of the questions to be asked, and some careful thought at this preliminary stage can be tremendously helpful and timesaving. Surprisingly, the exploitation of these ideas requires little more in the way of mathematics than is presumed by the old-fashioned approach. It merely requires a somewhat less mechanical and more creative use of the existing tools of constrained optimization.

I assume that the reader of this book is already familiar with the simple models of consumer behavior and producer behavior set out by, for example, Hirshleifer (1988) and McCloskey (1982), together with their geometric representations using such standard devices as indifference and isoquant maps, production frontiers, and so on.

The mathematical prerequisites involve no more than familiarity with the differential calculus and its application to classical optimization problems in which it is assumed that any constraint explicitly listed does in fact hold with equality at the optimum, and quantities consumed or produced are positive. The reader should be able to use the Lagrangian approach to derive the first-order necessary conditions that characterize an optimum and should have some appreciation of the need for, and nature of, the second-order curvature conditions that must also be satisfied in order that a point of tangency between, for example, an indifference curve and a budget constraint will truly chracterize an optimum rather than a pessimum. The algebraic expression of these mathematical ideas was expounded 50 years ago by Allen (1938) in a treatment that, though it is somewhat dated and inelegant, still has a lot to recommend it. More recent treatments have been presented by Intriligator (1971), Chiang (1984), Glaister (1984), and Weintraub (1982). In addition, I use simple matrix and vector notation in order to save space. However, I make little use of any but the most elementary operations, such as the inner product of two vectors and matrix multiplication.

Although the emphasis in this book is on the exploitation of a number of mathematical techniques to facilitate the analysis of optimizing models, the mathematics itself is not developed much beyond the prerequisites stated earlier. This is not a book on mathematics for economists. The tricks that I shall exploit arise from observing a few simple properties of optimizing problems, properties that are easy to grasp both geometrically and intuitively, obvious once understood, but that have been strangely overlooked and underexploited until recent years.

The rest of this chapter is devoted to highlighting those features of optimizing models that provide the basis for the dual tricks to be introduced in subsequent chapters. I have not attempted to be at all exhaustive, but have tried rather to concentrate on presenting a few useful ideas and the relationships between them, while continually referring to the familiar economic examples provided by the utility-maximizing consumer and profit-maximizing production plant.

In consumer theory, if the equality between marginal rates of substitution and relative prices is to characterize a utility maximum, the indifference map must curve "the right way." This leads us to the notions of quasi-concavity and quasi-convexity of a function, which are discussed in Section 1.1. Turning to the model of a profit-maximizing plant, we require the production function to curve "the right way" in order that the equality between value-marginal product and factor price will characterize a profit maximum. In this case, we require slightly stronger curvature restrictions, which lead naturally in Section 1.2 to a discussion of the notions of concavity and convexity of a function. Any function that is concave (convex) is certainly quasi-concave (quasi-convex), but the converse is not true. The relationships between these concepts are discussed in Section 1.3.

If an indifference map in commodity space, or a production frontier, has the appropriate curvature property to ensure that tangency implies optimality, this is because the modeler has made it so by assumption. Many of the functions encountered in this book, by contrast, have inherent curvature properties by virtue of their very definition. For example, the least cost that must be incurred by a consumer in order to attain a given utility level can be expressed as a concave function of the prices faced, as is shown in Chapter 3. In Chapter 5, I show that the maximum profit attained by a competitive production plant can be expressed as a convex function of the input and output prices faced. In these and other examples, such curvature properties, which turn out to have simple and useful implications for behavior, are direct results of the fact that we are dealing with maximum or minimum value functions – that is, functions that express the maximum or minimum value of an objective function in terms of the parameters of the optimizing problem. Section 1.4 discusses such functions and explores their curvature properties in the context of two particularly common classes of optimization problems.

Were it not for the clumsiness of the word, I would be tempted to call this book "Envelopeness and Modern Economics." Among the useful properties of maximum value functions is one known as the envelope property. It is this, as summarized by the envelope theorem in Section 1.5, that provides a close link between maximum value functions on the one hand and the supply and demand functions of maximizing agents on the other. This link enables us to use what we know about the former in order to draw inferences about behavioral functions. The observations that compensated consumer demand curves cannot slope upward and that profit-maximizing output supply curves cannot slope downward are among the many that can be derived in a simple and straightforward manner from the recognition of this link.

Section 1.6 takes the analysis further by considering a second-order, or generalized, envelope property. As its label suggests, this considers further properties of the second derivatives of maximum value functions with respect to the parameters of optimizing problems.

1.1 Quasi-concave and quasi-convex functions

Consider the familiar model of a utility-maximizing consumer. If the equilibrium involves strictly positive consumption of every commodity, then it is characterized by a set of first-order conditions requiring relative prices to equal marginal rates of substitution, represented geometrically as a point of tangency between the budget plane and an indifference surface, *ii*. The tangency requirement is the starting point for the standard analysis of consumer behavior that I present in later chapters. By itself, however, tangency is not sufficient to guarantee even a local, let alone a global, maximum. Only in panel (*a*) of Figure 1.1 does the tangency at E pick out a global utility maximum. In all the other cases there are other allocations, shown in the hatched areas, that are feasible and also yield higher utility than the allocation at E. In panels (*b*) and (*c*), E is not even a local maximum, whereas in (*d*), although E looks attractive by comparison with all feasible allocations in its neighborhood, it is not a global maximum. The bundle at G, however, is. Because of the linearity of the budget constraint, it is a relatively simple matter to find a set of restrictions on the curvature of indifference surfaces that are sufficient to rule out the situations depicted in panels (*b*), (*c*), and (*d*). We require the indifference curve that touches at E to lie otherwise wholly outside the feasible consumption set. If the allocation at E not only yields the maximum value of the objective function but also is the only allocation that yields that value, then it is called a strong global maximum. This certainly implies that in the neighborhood of E it should curve away from the budget set, so that panels (*b*) and (*c*) are ruled out. Farther away from E, we can allow it to curve back, but not by too much. Suppose there is a tangency at E, so that it is a local maximum. Then a sufficient condition for it to be also a strong global maximum is that the function $U(\cdot)$ be "strictly quasi-concave at E."

Before presenting this and other definitions, I should say something about the functions that appear throughout this book, and also about my conventions concerning notation. A function of n variables is simply a rule that associates with every point in some subset S of Euclidean n-dimensional space, R^n, a point on the real line, or a scalar. Points in R^n are represented by vectors, which are denoted by boldface letters. For example, the utility function $U(\mathbf{q})$ associates with each bundle of n commodities, $\mathbf{q} \equiv (q_1, q_2, \ldots, q_n)$, a real number that represents an individual's utility level. The production function $F(\boldsymbol{\ell})$ associates with each bundle of input quantities, $\boldsymbol{\ell} \equiv (\ell_1, \ell_2, \ldots, \ell_n)$, a real number that is the maximum attainable output from that bundle in the light of the prevailing

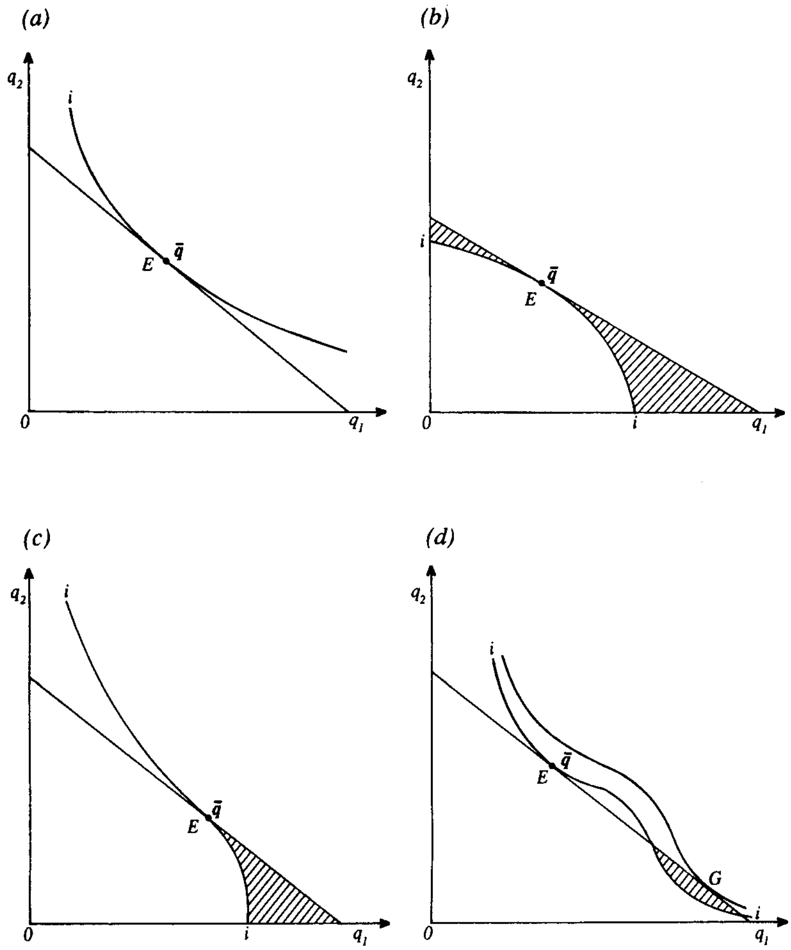

Figure 1.1

technology. The subset S of R^n over which a function is defined is called the domain of that function. I shall always assume that the domain S of any defined function $f(\cdot)$ is a convex set. The domain can often be taken as R^n itself. However, the economic interpretation of vectors as representing quantities or prices makes it natural often to confine attention to the nonnegative orthant of R^n, which is denoted by R^n_+. In so doing, we are confining attention to vectors whose elements are all either positive or zero. At many points, the argument can be simplified and certain technical problems avoided by further restricting attention to the strictly positive orthant of R^n, denoted by R^n_{++}. The statement "$\mathbf{x} \in S \subset R^n_{++}$" means that the vector $\mathbf{x}$ is an element of a set S that is itself a subset of the strictly positive orthant of Euclidean n-dimensional space, so that every element of $\mathbf{x}$ is strictly positive. For example, if $\mathbf{x}$ were a price vector, such a restriction would have the consequence of ruling out free goods. Let us now turn to the definition of strict quasi-concavity at a point:

Definition 1: Strict quasi-concavity at a point. *The function $f(\mathbf{x})$, defined on the convex set $S \subset R^n$, is strictly quasi-concave at the point $\bar{\mathbf{x}} \in S$ if for any other $\mathbf{x} \in S$ such that $f(\mathbf{x}) \geq f(\bar{\mathbf{x}})$,*

$$f(\bar{\mathbf{x}}) < f(\theta\bar{\mathbf{x}} + [1-\theta]\mathbf{x}) \quad \textit{for all } \theta \textit{ such that } 0 < \theta < 1.$$

In the context of consumer theory, this means that if an allocation $\mathbf{q}$ is at least as attractive as $\bar{\mathbf{q}}$, then any allocation on the chord joining $\mathbf{q}$ and $\bar{\mathbf{q}}$ is strictly preferred to $\bar{\mathbf{q}}$. This is sufficient to ensure that the tangency at $\bar{\mathbf{q}}$ characterizes a utility maximum.

Figure 1.2 emphasizes that strict quasi-concavity at $\bar{\mathbf{q}}$ is consistent with there being "dents" elsewhere in the indifference surface. The dent along the segment AB is not sufficiently great to mask $\bar{\mathbf{q}}$ from any other point on the surface – wherever $\mathbf{q}$ is chosen on ii, the chord lies wholly above the indifference curve. However, the presence of the dent does mean that had the budget constraint been chosen to touch ii between A and B, that tangency would not have characterized a utility maximum. Because we want a condition that guarantees a maximum at the tangency with any arbitrary budget constraint, our restriction has to be strengthened so as to rule out any dents in indifference surfaces. This leads to the global property of "strict quasi-concavity":

Definition 2: Strict quasi-concavity. *The function $f(\mathbf{x})$, defined on the convex set $S \subset R^n$, is strictly quasi-concave if for any pair of distinct points $\mathbf{x}' \in S$ and $\mathbf{x}'' \in S$ such that $f(\mathbf{x}') \geq f(\mathbf{x}'') = k$, $f(\theta\mathbf{x}' + [1-\theta]\mathbf{x}'') > k$ for all θ such that $0 < \theta < 1$.*

This is certainly true in Figure 1.1(a), but not in the remaining panels. Observe that the definition, unlike the geometry, can handle as many dimensions as desired.

In the context of consumer theory, an immediate and very convenient consequence of strict quasi-concavity is that for any budget constraint consistent with there being a point of tangency, that point is unique and is guaranteed to achieve the maximum attainable utility.

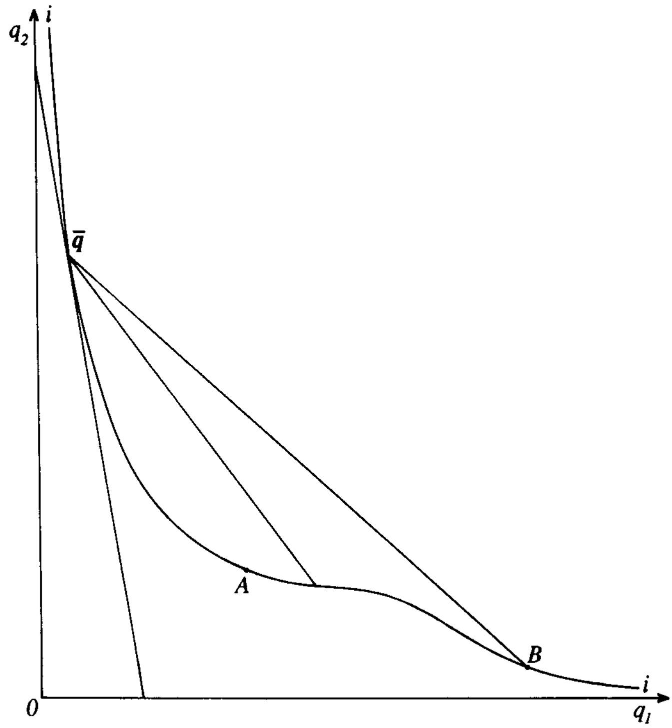

Figure 1.2

Often a slightly weaker property is assumed, and it allows indifference surfaces to have linear portions in them. This is the property of quasi-concavity:

Definition 3: Quasi-concavity. *The function $f(\mathbf{x})$, defined on the convex set $S \subset R^n$, is quasi-concave if for any pair of distinct points $\mathbf{x}' \in S$ and $\mathbf{x}'' \in S$ such that $f(\mathbf{x}') \geq f(\mathbf{x}'') = k$, $f(\theta\mathbf{x}' + [1-\theta]\mathbf{x}'') \geq k$ for all θ such that* $0 < \theta < 1$.

There is another extremely elegant definition in terms of sets. Given any arbitrary point $\hat{\mathbf{x}}$, consider the set of all points $\mathbf{x}$ such that $f(\mathbf{x}) \geq f(\hat{\mathbf{x}})$. Koopmans (1957) calls this the "no-worse-than-$\hat{\mathbf{x}}$" set. Alternatively, it is the upper contour set associated with $\hat{\mathbf{x}}$. For example, if $\hat{\mathbf{x}}$ is identified with $\bar{\mathbf{q}}$ in Figure 1.1(a), the no-worse-than-$\hat{\mathbf{x}}$ set is the set of points above and to the right of ii. Then we have the following:

Definition 3′: Quasi-concavity. *The function $f(\mathbf{x})$ is quasi-concave if the upper contour set associated with every point* $\mathbf{x}$ is convex.

The statement that preferences are assumed to be convex refers precisely to this assumption on the utility function. A quasi-concave utility function permits an infinite number of allocations consistent with utility maximization. For example, if the linear portion of an indifference curve coincides with a segment of the budget constraint, all points along the common stretch give rise to the same maximized value of utility. This in no way threatens the structure of dual arguments. However, it is convenient to assume that there is a unique maximizing allocation – that is, a strong global maximum – and strict quasi-concavity represents a convenient way to ensure that this is so. Henceforth, I shall generally assume that the maximizing vector $\hat{\mathbf{x}}$ is unique.

Quasi-convex and strictly quasi-convex functions, which are in a natural sense the opposites of quasi-concave and strictly quasi-concave functions, are less commonly encountered in economics. However, an important example is introduced in Chapter 2. Intuition suggests that the maximum attainable utility can be expressed as a function of the prices and income faced by a consumer. Such an indirect utility function turns out to be an inherently quasi-convex function of prices. Quasi-convexity can be defined in a manner similar to that used earlier. This, for example, is done by Intriligator (1971, p. 464). An alternative definition is as follows:

Definition 4: [Strict] Quasi-convexity. *A function $f(\mathbf{x})$ is [strictly] quasi-convex if its negative, $-f(\mathbf{x})$, is [strictly] quasi-concave.*

I leave it to the reader to check, for example, that this definition implies that f is quasi-convex if for any value of $\hat{\mathbf{x}}$, the set of $\mathbf{x}$ such that $f(\mathbf{x}) \leq f(\hat{\mathbf{x}})$ is convex. This, if you like, is the no-better-than-$\hat{\mathbf{x}}$ set, or the lower contour set associated with $\hat{\mathbf{x}}$.

To summarize, quasi-concavity is a property relating to the curvature of the level sets, such as indifference surfaces or isoquants, associated with particular values of a given function, such as a utility or production function. Its particular significance for economists arises from the fact that the quasi-concavity of the utility function $U(\mathbf{q})$, combined with the linearity of the individual's budget constraint, guarantees that the first-order conditions associated with a point of tangency are sufficient as well as necessary for a utility maximum. Strict quasi-concavity ensures that there is only one bundle that yields the maximum utility. Quasi-convexity is simply a mirror image of quasi-concavity.

1.2 Concave and convex functions

Now consider a simple problem drawn from producer theory, that of a competitive profit-maximizing plant. For the moment, let there be one input, ℓ, and one output, y. Their market prices are, respectively, W and P. The problem involves choosing ℓ and y to maximize profit $Py - W\ell$, which is a simple linear function of the quantities ℓ and y, subject to the technology constraint as expressed by the production function, $y = F(\ell)$. The level sets of the profit identity, or iso-profit curves, are parallel straight lines in (y, ℓ) space. The question to be asked is, When can we be certain that fulfillment of the first-order conditions, represented geometrically by a tangency between an iso-profit line and the production function, does indeed represent a true global profit maximum? Before answering this, I should make explicit an assumption concerning the production function. I shall assume that the marginal product of labor is everywhere positive, so that $F(\ell)$ is monotonic increasing. A quick check reveals that in the special case of a function of a single variable, this is sufficient to ensure that it is both strictly quasi-concave and strictly quasi-convex. Figure 1.3 shows that such properties in no way ensure that tangency implies optimality. In panel (*a*), the allocation T is indeed optimal, but in panels (*b*) and (*c*) the curvature of $F(\ell)$ implies feasible allocations that are more profitable than the allocation at the tangent T. In panel (*d*), T emerges well from local comparisons, but there remain feasible allocations in the hatched area that are more profitable. The most common restriction imposed on the curvature of $F(\cdot)$ to remove such possibilities and to ensure a unique point of tangency is the assumption that $F(\cdot)$ is strictly concave. The definition of a strictly concave function is as follows:

Definition 5: Strict concavity. *The function $f(\mathbf{x})$, defined on the convex set $S \subset R^n$, is strictly concave if for any pair of distinct points $\mathbf{x}' \in S$ and $\mathbf{x}'' \in S$, $f(\theta\mathbf{x}' + [1-\theta]\mathbf{x}'') > \theta f(\mathbf{x}') + [1-\theta] f(\mathbf{x}'')$ for all θ such that $0 < \theta < 1$.*

A strictly concave function may alternatively be characterized by noting that the tangent at any point consistently lies above the function at every other point. Often it can be helpful to exploit this characterization.

The slightly weaker concept of a concave function allows for linear surfaces:

Definition 6: Concavity. *The function $f(\mathbf{x})$, defined on the convex set $S \subset R^n$, is concave if for any pair of distinct points $\mathbf{x}' \in S$ and $\mathbf{x}'' \in S$, $f(\theta\mathbf{x}' + [1-\theta]\mathbf{x}'') \geq \theta f(\mathbf{x}') + [1-\theta] f(\mathbf{x}'')$ for all θ such that $0 \leq \theta \leq 1$.*