

ONE

Essential Elements of Continuous Time Dynamic Optimization

In order to motivate the following introductory material on dynamic optimization problems, it will be advantageous to draw heavily on your knowledge of static optimization theory. To that end, we begin by recalling the definition of the prototype unconstrained static optimization problem, namely,

$$\phi(\alpha) \stackrel{\text{def}}{=} \max_{\mathbf{x} \in \Re^N} f(\mathbf{x}; \alpha), \tag{1}$$

where $\mathbf{x} \in \Re^N$ is a vector of *decision* or *choice variables*, $\alpha \in \Re^A$ is a vector of *parameters*, $f(\cdot)$ is the twice continuously differentiable *objective function*, that is, $f(\cdot) \in C^{(2)}$, and $\phi(\cdot)$ is the *indirect* or *maximized objective function*. This is terminology you should be more or less familiar with from prior courses.

Because we will deal repeatedly with vectors and matrices as well as the derivatives of scalar- and vector-valued functions in this book, we pause momentarily to establish three notational conventions that we shall adhere to throughout. First, all vectors are treated as column vectors. To denote a row vector, we therefore employ the transpose operator, denoted by the symbol $'$. Thus $\mathbf{x} \in \Re^N$ is taken to be an N -element column vector, whereas $\mathbf{x}'$ is an N -element row vector. Note also that vectors appear in **boldface** type.

Second, if $\mathbf{g}(\cdot) : \Re^N \rightarrow \Re^M$ is a $C^{(1)}$ vector-valued function, thereby implying that $\mathbf{g}(\cdot) \stackrel{\text{def}}{=} (g^1(\cdot), g^2(\cdot), \dots, g^M(\cdot))'$, then at any $\mathbf{x} \in \Re^N$, we define the $M \times N$ Jacobian matrix of $\mathbf{g}(\cdot)$ by

$$\underbrace{\mathbf{g}_{\mathbf{x}}(\mathbf{x})}_{M \times N} \stackrel{\text{def}}{=} \begin{bmatrix} g^1_{x_1}(\mathbf{x}) & g^1_{x_2}(\mathbf{x}) & \cdots & g^1_{x_N}(\mathbf{x}) \\ g^2_{x_1}(\mathbf{x}) & g^2_{x_2}(\mathbf{x}) & \cdots & g^2_{x_N}(\mathbf{x}) \\ \vdots & \vdots & \ddots & \vdots \\ g^M_{x_1}(\mathbf{x}) & g^M_{x_2}(\mathbf{x}) & \cdots & g^M_{x_N}(\mathbf{x}) \end{bmatrix}, \tag{2}$$

where $g^m_{x_n}(\mathbf{x})$ is the partial derivative of $g^m(\cdot)$ with respect to x_n evaluated at the point $(\mathbf{x})$. It is also the element in the m th row and n th column of $\mathbf{g}_{\mathbf{x}}(\mathbf{x})$. This definition implies that if $M = 1$, so that $\mathbf{g}(\cdot) : \Re^N \rightarrow \Re$ is now a scalar-valued function, then

$g_{\mathbf{x}}(\mathbf{x}) = (g_{x_1}(\mathbf{x}), g_{x_2}(\mathbf{x}), \dots, g_{x_N}(\mathbf{x}))$ is a row vector, or equivalently, a $1 \times N$ matrix. This means that the derivative of a scalar-valued function with respect to a column vector is a row vector. As an extension of this notation, if we now assume that $g(\cdot) : \Re^{N+A} \rightarrow \Re^M$ is a $C^{(1)}$ function whose arguments are the vectors $\mathbf{x} \in \Re^N$ and $\alpha \in \Re^A$, then $g_{\mathbf{x}}(\mathbf{x}; \alpha)$ is the $M \times N$ Jacobian matrix given in Eq. (2), whereas $g_{\alpha}(\mathbf{x}; \alpha)$ is an $M \times A$ Jacobian matrix defined similarly.

Third, if $g(\cdot) : \Re^{N+A} \rightarrow \Re$ is a $C^{(2)}$ scalar-valued function whose arguments are the vectors $\mathbf{x} \in \Re^N$ and $\alpha \in \Re^A$, then there are four Hessian matrices that can be defined based on $g(\cdot)$ because of the two different sets of variables that it depends on, scilicet,

$$\underbrace{g_{\mathbf{x}\mathbf{x}}(\mathbf{x}; \alpha)}_{N \times N} \stackrel{\text{def}}{=} \begin{bmatrix} g_{x_1 x_1}(\mathbf{x}; \alpha) & g_{x_1 x_2}(\mathbf{x}; \alpha) & \cdots & g_{x_1 x_N}(\mathbf{x}; \alpha) \\ g_{x_2 x_1}(\mathbf{x}; \alpha) & g_{x_2 x_2}(\mathbf{x}; \alpha) & \cdots & g_{x_2 x_N}(\mathbf{x}; \alpha) \\ \vdots & \vdots & \ddots & \vdots \\ g_{x_N x_1}(\mathbf{x}; \alpha) & g_{x_N x_2}(\mathbf{x}; \alpha) & \cdots & g_{x_N x_N}(\mathbf{x}; \alpha) \end{bmatrix},$$

$$\underbrace{g_{\mathbf{x}\alpha}(\mathbf{x}; \alpha)}_{N \times A} \stackrel{\text{def}}{=} \begin{bmatrix} g_{x_1 \alpha_1}(\mathbf{x}; \alpha) & g_{x_1 \alpha_2}(\mathbf{x}; \alpha) & \cdots & g_{x_1 \alpha_A}(\mathbf{x}; \alpha) \\ g_{x_2 \alpha_1}(\mathbf{x}; \alpha) & g_{x_2 \alpha_2}(\mathbf{x}; \alpha) & \cdots & g_{x_2 \alpha_A}(\mathbf{x}; \alpha) \\ \vdots & \vdots & \ddots & \vdots \\ g_{x_N \alpha_1}(\mathbf{x}; \alpha) & g_{x_N \alpha_2}(\mathbf{x}; \alpha) & \cdots & g_{x_N \alpha_A}(\mathbf{x}; \alpha) \end{bmatrix},$$

$$\underbrace{g_{\alpha\alpha}(\mathbf{x}; \alpha)}_{A \times A} \stackrel{\text{def}}{=} \begin{bmatrix} g_{\alpha_1 \alpha_1}(\mathbf{x}; \alpha) & g_{\alpha_1 \alpha_2}(\mathbf{x}; \alpha) & \cdots & g_{\alpha_1 \alpha_A}(\mathbf{x}; \alpha) \\ g_{\alpha_2 \alpha_1}(\mathbf{x}; \alpha) & g_{\alpha_2 \alpha_2}(\mathbf{x}; \alpha) & \cdots & g_{\alpha_2 \alpha_A}(\mathbf{x}; \alpha) \\ \vdots & \vdots & \ddots & \vdots \\ g_{\alpha_A \alpha_1}(\mathbf{x}; \alpha) & g_{\alpha_A \alpha_2}(\mathbf{x}; \alpha) & \cdots & g_{\alpha_A \alpha_A}(\mathbf{x}; \alpha) \end{bmatrix},$$

and the $A \times N$ matrix $g_{\alpha\mathbf{x}}(\mathbf{x}; \alpha)$, the computation of which we leave for a mental exercise. We remark in passing that there is a matrix version of the invariance of the second-order partial derivatives to the order of differentiation, and this too is left for a mental exercise.

With the notational matters settled, let's now return to the unconstrained static optimization problem defined in Eq. (1). Assume that an optimal solution exists to problem (1), say $\mathbf{x} = \mathbf{x}^*(\alpha)$. Typically, we would find the solution by simultaneously solving the first-order necessary conditions (FONCs) of problem (1), which are given by

$$f_{\mathbf{x}}(\mathbf{x}; \alpha) = \mathbf{0}'_N$$

in vector notation, where $\mathbf{0}_N$ is the null (column) vector in $\Re^N$, or by

$$f_{x_n}(\mathbf{x}; \alpha) = 0, \quad n = 1, 2, \dots, N$$

in index notation. Unless the objective function $f(\cdot)$ happened to be of a particularly simple functional form, an explicit solution for $\mathbf{x} = \mathbf{x}^*(\alpha)$ is more often than not rare. If, however, we assume that the second-order sufficient condition (SOSC) holds

at $\mathbf{x} = \mathbf{x}^*(\alpha)$, that is,

$$\mathbf{h}' f_{\mathbf{xx}}(\mathbf{x}^*(\alpha); \alpha) \mathbf{h} < 0 \forall \mathbf{h} \in \mathbb{R}^N, \quad \mathbf{h} \neq \mathbf{0}_N,$$

or

$$\sum_{i=1}^N \sum_{j=1}^N f_{x_i x_j}(\mathbf{x}^*(\alpha); \alpha) h_i h_j < 0, \quad \text{not all } h_i = 0,$$

then we can apply the implicit function theorem to the FONCs to solve for the optimal choice vector $\mathbf{x} = \mathbf{x}^*(\alpha)$ in principle. To see why this is so, recall that the Jacobian matrix of the FONCs is given by the $N \times N$ matrix $f_{\mathbf{xx}}(\mathbf{x}; \alpha)$, which is identical to the Hessian matrix of the objective function. Moreover, the SOSC implies that the Hessian determinant of the FONCs is nonvanishing at the optimal solution, that is, that $|f_{\mathbf{xx}}(\mathbf{x}; \alpha)| \neq 0$ when evaluated at $\mathbf{x} = \mathbf{x}^*(\alpha)$. Because this is equivalent to the nonvanishing of the Jacobian determinant of the FONCs, the implicit function theorem may be applied to the FONCs to solve, in principle, for the optimal choice vector $\mathbf{x} = \mathbf{x}^*(\alpha)$. Again, this line of reasoning should be familiar to you from prior courses in microeconomic theory.

In light of the dynamic problems that will occupy us in this book, the most important aspect of the above discussion concerning problem (1) is that for a given value of the parameters, say, $\alpha = \alpha^\circ$, we typically solve for a particular value for each of the decision or choice variables, say, $x_n^\circ = x_n^*(\alpha^\circ)$, $n = 1, 2, \dots, N$. We do this by solving a set of algebraic equations, that is to say, the FONCs. If the parameter vector is different, say, $\alpha = \alpha^1$, then usually a different value of the choice variables is implied, say, $x_n^1 = x_n^*(\alpha^1)$, $n = 1, 2, \dots, N$. By their very nature, therefore, static optimization problems ask the decision maker to pick out a particular value of the decision variables given the parameters of the problem. For $A = N = 1$, Figure 1.1 depicts this situation graphically.

To add an economic spin to all of this, recall the prototype profit-maximizing model of the price-taking firm:

$$\pi^*(p, w_1, w_2) \stackrel{\text{def}}{=} \max_{x_1, x_2} \{pF(x_1, x_2) - w_1 x_1 - w_2 x_2\},$$

where $F(\cdot)$ is the twice continuously differentiable production function, x_1 and x_2 are the decision variables representing the inputs of the firm, w_1 and w_2 are the market prices of the inputs, p is the output price, and $\pi^*(\cdot)$ is the indirect profit function. Given a particular set of prices, say, $(p, w_1, w_2) = (p^\circ, w_1^\circ, w_2^\circ)$, the firm seeks to determine the values of the inputs that maximize its profit, say, $x_n^\circ = x_n^*(p^\circ, w_1^\circ, w_2^\circ)$, $n = 1, 2$. If such optimal values exist, then they are found, in principle, by simultaneously solving the FONCs, given by

$$p^\circ F_{x_1}(x_1, x_2) - w_1^\circ = 0,$$

$$p^\circ F_{x_2}(x_1, x_2) - w_2^\circ = 0.$$




Figure 1.1

Note that these are *algebraic equations* that in general must be solved simultaneously for the optimal values of the inputs. To repeat, the most important point to take away from this discussion is that the choice of the optimal input combination is made just one time: there is no planning for the future, nor are there future decisions to be made in this problem. This is exactly as the static framework of the problem dictates.

In contrast, given the parameters of a dynamic optimization problem, its solution is a *sequence* of optimal decisions in discrete time, or a *time path* or *curve* of optimal decisions in continuous time, over the relevant *planning period* or *planning horizon*, not just one particular value for each of the decision variables like the solution to a static optimization problem. The optimal time path or curve is, by definition, the one that optimizes some type of objective function. The type of objective function in dynamic problems, however, is quite different from that in static problems. Because the solution to a continuous time dynamic optimization problem is a time path or curve, it appears to be reasonable and even natural for the objective function to place a value on the decision variables at each point in time of the planning horizon and to add up the resulting values over the relevant planning period, akin to what is done when one computes the present value of some stream of net benefits that are received over time.

To better motivate the form of the objective function in dynamic problems, consider Figure 1.2. Here three typical time paths of a function $x(\cdot)$, or curves $x(t)$ associated with the function $x(\cdot)$, are displayed along with the resulting value of the objective function associated with each time path $J[x(\cdot)]$, the latter of which we refer to as a *path value*. We have denoted the independent variable by the letter t and refer to it as time, as this is the natural interpretation of the independent variable in intertemporal problems in economics. Notice that all time paths or curves begin




Figure 1.2

at time $t = t_0$ at the point $x = x_0$ and end at time $t = t_1$ at the point $x = x_1$, all four of which are given or fixed, thereby requiring that the paths being compared begin and end at the same position and time. The typical problem in dynamic optimization seeks to find a time path or curve $x(t)$, or equivalently a function $x(\cdot)$, that, say, maximizes the objective function $J[\cdot]$. Thus to each time path or curve $x_i(t)$, $i = a, b, c$, or function $x_i(\cdot)$, $i = a, b, c$, there is a corresponding value of the objective function $J[x_i(\cdot)]$, $i = a, b, c$.

The relationship between paths $x(t)$ or functions $x(\cdot)$ and the resulting value $J[x(\cdot)]$ is quite different from that encountered in typical lower division mathematics courses. It represents a mapping from paths or curves to real numbers, or equivalently, from functions to real numbers, and therefore is not a mapping from real numbers to real numbers as in the case of functions. Such a mapping from paths or curves to path values, or from functions to real numbers, is what Figure 1.2 depicts and is called a *functional*. The general notation we shall employ for such a mapping from functions to real numbers is $J[x(\cdot)]$, which has been employed above. This notation emphasizes that the functional $J[\cdot]$ depends on the *function* $x(\cdot)$, or equivalently, on the *entire curve* $x(t)$. Moreover, it highlights the fact that it is a change in the position of the entire path or curve $x(t)$, that is, the *variation* in the path or curve $x(t)$, rather than the change in t , that results in a change in the path value or functional $J[\cdot]$. Thus, a dynamic optimization problem in continuous time seeks to find a path or curve $x(t)$, or equivalently, a function $x(\cdot)$, that optimizes an objective functional $J[\cdot]$.

Next we consider in more detail the form of an archetype objective functional $J[\cdot]$. Because the optimal solution to a continuous time dynamic optimization problem is a path or curve $x(t)$, as noted above, associated with the path is its slope $\dot{x}(t) \stackrel{\text{def}}{=} dx(t)/dt$ at each point in time t in the planning horizon, assuming, of course, that the path $x(t)$ is smooth enough so that $\dot{x}(t)$ is defined. Suppose, moreover, that

there exists a function, say, $F(\cdot)$, that assigns or imputes a value to the path and its associated derivative at each point in time in the planning horizon, the latter represented by the closed interval $[t_0, t_1]$, $0 < t_0 < t_1$. The imputed value of the path at each point in the planning horizon, therefore, depends on the moment of time t the decision is made and the value of the decision or choice variable at that time $x(t)$, as well as on the slope at that time $\dot{x}(t)$. Hence we have $F(t, x(t), \dot{x}(t))$ as the value of the function that imputes a value to the path $x(t)$ with slope $\dot{x}(t)$ at time t . Because the path $x(t)$ must necessarily travel through an interval of time, namely, the planning horizon, its total value as represented by the functional $J[x(\cdot)]$ is given by the “sum” of all the imputed values $F(t, x(t), \dot{x}(t))$ for each t in the planning horizon $[t_0, t_1]$. Moreover, because we are operating in continuous time, the appropriate notion of summation is represented by a definite integral over the closed interval $[t_0, t_1]$. Thus, the value of the functional we wish to optimize is given by

$$J[x(\cdot)] \stackrel{\text{def}}{=} \int_{t_0}^{t_1} F(t, x(t), \dot{x}(t)) dt. \tag{3}$$

In economics, $J[x(\cdot)]$ often represents the present value of net benefits from pursuing the policy $x(t)$, with instantaneous net benefits given by $F(t, x(t), \dot{x}(t))$. In general, one calls $x(t)$ the state of the system, or value of the state variable, at time t , and $\dot{x}(t)$ the rate of change of the system, or velocity, at time t . Typically, the *explicit* appearance of t as an argument of $F(\cdot)$ is a result of discounting in economic problems. Equation (3) represents the prototype form of the objective functional for *calculus of variations* problems, which are but one class of continuous time dynamic optimization problems. More general objective functionals will be introduced shortly, when we commence with the study of optimal control theory. But for now, the present discussion and motivation are sufficient, for the idea of a mapping from paths or curves to the real line and the form of the objective functional is what one must come away with.

Before moving on to some additional motivational material, a few remarks about the form of the objective functional in Eq. (3) are warranted. First, the functional $J[\cdot]$ is *not*, in general, the area under the curve $x(t)$ between the points $t = t_0$ and $t = t_1$, the latter of which is represented by the integral $\int_{t_0}^{t_1} x(t) dt$. Thus, in optimizing $J[\cdot]$, we are *not* making decisions to optimize the area under the curve $x(t)$. Rather, we are picking paths $x(t)$ such that the “sum” of the values imputed to the path $x(t)$ and its derivative $\dot{x}(t)$ at each point in time in the planning horizon is optimized. Second, the form of $J[\cdot]$ given in Eq. (3) is not the most general form of the objective functional for calculus of variations problems, but it is the most common or canonical in economic theory, as examples to be introduced latter will confirm. In motivating the form of $J[\cdot]$, for example, there is no particular reason why $F(\cdot)$ would not, in general, depend on the second or higher derivatives of $x(t)$. What dictates which derivatives of the path $x(t)$ that are relevant to the form of $J[\cdot]$

Cambridge University Press

0521603684 - Foundations of Dynamic Economic Analysis: Optimal Control Theory and Applications

Michael R. Caputo

Excerpt

[More information](#)

is the particular economic phenomena under study, and as we will shortly see, most of the problems of relevance to economists dictate the presence of $\dot{x}(t)$ in $F(\cdot)$, but rarely higher derivatives.

The question you may now be wrestling with is: How does one know when to construct a dynamic model, as opposed to a static model, to study the economic events under investigation? The answer is most easily explained and motivated within the context of a simple example. Imagine an individual on an isolated island on which fresh water is available in an unlimited supply, and for which essentially no effort is required to collect the water. You may picture a brook or stream of fresh water passing next to the individual's hut. The only source of food is fish, which does require the expenditure of effort on the part of this individual. Clearly this individual can survive on this island, and the problem that this person faces is how much fish to catch each day for consumption.

Initially, let's assume that the harvested fish are impossible to store for any length of time, either because of the extreme heat of the environment or because of the lack of materials necessary to build a suitable storage facility. As this individual gets up on any day, the decision to be made is how much fish to catch *for this day only*, as storage has been ruled out. Any fish caught beyond the amount to be consumed that day simply rots and is wasted. Because fishing is costly to the individual, the catch on any day will not exceed the individual's consumption per day. Notice that the decision of how much fish to catch this day is independent of fish caught on previous days or fish expected to be caught on future days, as the lack of storage prevents any carryover of the fish. This lack of storage breaks any link between past decisions and the present decision, and any link between the present decision and future decisions. In other words, the absence of any durable asset or the inability of this person to store any of the asset (fish) renders current choices or actions independent of those made in the past, or those to be made in the future. For example, even if this person caught twice as many fish as could be consumed in a day, this would not relieve the individual of fishing the following day because the fish simply spoil, leaving zero edible fish for tomorrow. Thus the decision of how much fish to catch on any day is dependent only on the circumstances or environmental conditions of that day. As the reader may have guessed, this is exactly what dictates the decision problem faced by this individual as simply a sequence of static choice problems, each day's decision being independent of past and future decisions, and identical to that to be made on any other day, save for differing environmental conditions.

The reader may wonder why this situation is not a dynamic optimization problem since a sequence of optimal decisions must be made through time. It is not the sequence of decisions or the introduction of time per se that defines a dynamic choice problem but the link between past, present, and future decisions that makes a problem dynamic. In the scenario above, the lack of storage breaks this link, reducing the problem to a sequence of independent static optimization problems. So to have a dynamic optimization problem, *there must be some systematic link between past, present, and future decisions.*

Cambridge University Press

0521603684 - Foundations of Dynamic Economic Analysis: Optimal Control Theory and Applications

Michael R. Caputo

Excerpt

[More information](#)

Now let's assume that fish caught on any day in excess of that day's consumption can be stored. Because we are concerned here about the structure of a problem that makes it dynamic, the actual period of time in which the fish can be stored or preserved is not important; the fact that fish caught on one day can be stored for future consumption is the important idea. Just as in the previous scenario, when this individual wakes up on any given day, a decision about how much fish to catch that day must be made. What is different, however, is that a stock of fish may exist in storage from the previous day's catch, and this stock must be taken into account in today's harvesting decision. Thus, the assumption of storage (or a durable good) provides a direct link between past decisions and the current decision, a link that was absent when storage was ruled out. Likewise, the decision to catch fish (or not) today impacts the amount of fish in storage for future consumption, and therefore impacts future decisions about catching fish. Storage provides a link between current decisions and future decisions as well. It is exactly this intertemporal linking of decisions that makes this second scenario a dynamic choice problem: decisions made in the past affect the current choice, which, in turn, affects future choices.

Although it may appear that in this simple example, there is only one variable, *scilicet*, fish, there are actually two: catching fish and fish in storage. The storage of fish is not a variable that is controlled directly by the individual; it responds to the amount caught, amount eaten, and time elapsed between catches. In macroeconomic terminology, the amount of fish in storage is a stock variable, or a state variable in the language of optimal control theory; that is, it is defined at a point in time, not over a period or length of time. The act of fishing, on the other hand, is defined over a period of time (a flow variable) and is directly under the control of the individual. In the language of optimal control theory, the catch rate of the fish is the control variable.

With the essence of a continuous time dynamic optimization problem now conveyed, let's turn to the motivation and basic mathematical setup of an optimal control problem. We therefore elect to proceed directly to optimal control theory rather than first formally introducing the calculus of variations and *then* optimal control theory.

Optimal control theory is based on a new way of viewing and formulating calculus of variations problems, and thus enables one to see them in a different light. In particular, optimal control theory often brings the economic intuition and content of a continuous time dynamic optimization problem to the surface more readily than does the calculus of variations, thereby enhancing one's economic understanding of the problem. It is this change of vista that makes optimal control theory a powerful tool for solving dynamic economic problems, for the calculus of variations can solve all problems that can be solved with optimal control theory, though not necessarily as easily, in which case both theories yield equivalent results. In fact, some textbooks, such as Hadley and Kemp (1971), develop the calculus of variations in its full generality and then use the results to prove those in optimal control theory.

The focus in optimal control theory is on some system. In economic problems, this may be the economy, an individual, or a firm. As is usual in intertemporal problems, we are interested in optimizing, in some specified sense, the behavior

of the system through time. It is assumed that the manner in which the system changes through time can be described by specifying the time behavior of certain variables, say $\mathbf{x}(t) \in \Re^N$, called *state variables*, where t is the independent variable that we will almost always refer to as time. In general, the vector-valued function $\mathbf{x}(\cdot) : \Re \rightarrow \Re^N$ is assumed to be a *piecewise smooth* function of time with not more than a finite number of corners. This means that the component functions $x_n(\cdot)$, $n = 1, 2, \dots, N$, are continuous but that the derivative functions $\dot{x}_n(\cdot)$, $n = 1, 2, \dots, N$, are *piecewise continuous* in the sense that $\dot{x}_n(\cdot)$ has at most a finite number of discontinuities on each finite interval with finite jumps (i.e., one-sided limits) at each point of discontinuity. In economic problems, a capital stock, a stock of money or any asset for that matter (e.g., wealth), a stock of fish, the amount of some mineral in the ground, the stock of water in an aquifer, the number of chairs in a classroom, or even the distribution function of a random variable may represent a state variable. Generally, the state variables are defined at a given point in time, as the aforementioned examples indicate. This is why state variables are often referred to as stocks by economists.

Before pressing on, it is prudent at this juncture to pause momentarily and give a precise definition of a piecewise continuous function and a piecewise smooth function, for such functions will be encountered with some regularity in optimal control theory. To that end, we have the following definition.

Definition 1.1: A function $\phi(\cdot)$ is said to be *piecewise continuous* on an interval $\alpha \leq t \leq \beta$ if the interval can be partitioned by a finite number of points $\alpha = t_0 < t_1 < \dots < t_K = \beta$ so that

1. $\phi(\cdot)$ is continuous on each open subinterval $t_{k-1} < t < t_k$, $k = 1, 2, \dots, K$, and
2. $\phi(\cdot)$ approaches a finite limit as the end points of each subinterval are approached from within the subinterval.

In other words, a function $\phi(\cdot)$ is piecewise continuous on an interval $\alpha \leq t \leq \beta$ if it is continuous there except for a finite number of jump discontinuities. An example of a piecewise continuous function is shown in Figure 1.3. Given this definition, it is now a relatively simple matter to define a piecewise smooth function.

Definition 1.2: A function $\Phi(\cdot)$ is said to be *piecewise smooth* on an interval $\alpha \leq t \leq \beta$ if its derivative function $\dot{\Phi}(\cdot)$ is piecewise continuous on the interval $\alpha \leq t \leq \beta$.

This definition therefore implies that the derivative of a piecewise smooth function is piecewise continuous, and that the integral of a piecewise continuous function is piecewise smooth.

It is also assumed that there exists another class of variables known as *control variables*, say $\mathbf{u}(t) \in \Re^M$, where $M \neq N$ in general. The control variables may undergo jump changes and are therefore only restricted to be *piecewise continuous* in general, that is, the function $\mathbf{u}(\cdot) : \Re \rightarrow \Re^M$ is assumed to be a *piecewise continuous*




Figure 1.3

function of time. In economic problems, control variables are usually represented by flow variables that are typically defined over an interval of time. Archetypal control variables in economic problems include the investment rate in an asset, the consumption rate of a good, and the harvest rate of a resource. Control variables are those variables that the decision maker has explicit control of in the optimal control problem, as the name implies. Control variables are thus the analogues of the choice or decision variables in static optimization theory. The control variables will not, in general, be allowed to take on arbitrary values. Generally, it is assumed that for each t in the planning horizon, the control variables are restricted to vary in a fixed and prespecified set $U \subseteq \mathbb{R}^M$, called the *control set* or *control region*. It is typically required that $\mathbf{u}(t) \in U$ for the entire planning horizon. The control set U may be any fixed set in $\mathbb{R}^M$, say, an open or closed set, but is not restricted in any special way. Of particular importance is the case in which U is a closed set in $\mathbb{R}^M$. In this case, the control variables are allowed to take values at the boundary of the control set, a situation that the classical calculus of variations cannot handle so easily. In typical economic problems, the control set may be represented by nonnegativity restrictions on the control variables or by fixed inequality constraints that bound the control variables. Finally, it is assumed that every variable of interest can be classified as a state variable or a control variable.

In view of the restrictions placed on the two classes of variables, we may think of the control variables as governing not the values of the state variables, but their rate of change. More specifically, it will be assumed that the dependence of the state variables on the control variables can be described by a first-order differential equation system, namely,

$$\dot{\mathbf{x}}(t) = \mathbf{g}(t, \mathbf{x}(t), \mathbf{u}(t))$$

in vector notation, or

$$\dot{x}_n(t) = g^n(t, \mathbf{x}(t), \mathbf{u}(t)), \quad n = 1, 2, \dots, N$$