Springer Monographs in Mathematics

Piernicola Bettiol
Richard B. Vinter

Principles of Dynamic Optimization

Springer

Springer Monographs in Mathematics

This series publishes advanced monographs giving well-written presentations of the "state-of-the-art" in fields of mathematical research that have acquired the maturity needed for such a treatment. They are sufficiently self-contained to be accessible to more than just the intimate specialists of the subject, and sufficiently comprehensive to remain valuable references for many years. Besides the current state of knowledge in its field, an SMM volume should ideally describe its relevance to and interaction with neighbouring fields of mathematics, and give pointers to future directions of research.

Piernicola Bettiol • Richard B. Vinter

Principles of Dynamic Optimization

Springer

Piernicola Bettiol
Department of Mathematics
University of Brest
Brest, France

Richard B. Vinter
Electrical & Electronic Engineering
Imperial College London
London, UK

ISSN 1439-7382 ISSN 2196-9922 (electronic)
Springer Monographs in Mathematics
ISBN 978-3-031-50088-6 ISBN 978-3-031-50089-3 (eBook)
https://doi.org/10.1007/978-3-031-50089-3

Mathematics Subject Classification: 49-02, 49K15, 49K21

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

To Francis Clarke, il miglior fabbro

Preface

What control strategy will transfer a space vehicle from one circular orbit to another in least time or, alternatively, with minimum fuel consumption? What should be the strategy for harvesting a renewable resource (a fish population, say) to maximize financial returns while satisfying sustainability constraints? In chemo-immunotherapy for cancer, what treatment regime (concentration and frequency of cytotoxic drug doses) will minimize the tumour cell population while maintaining the blood cell population above a critical level? Mitigation strategies are available to counter an epidemic, including vaccination, livestock culling and host removal; how should we deploy these strategies while minimizing the social and economic costs involved? How should a batch distillation column be operated to maximize the yield, subject to specified constraints on product purity?

There are a number of common features in these questions. First, they all concern phenomena where the relevant 'state of nature' (relating, for example, to the position of a space vehicle or the size of a diseased population) is dynamic, in the sense that it evolves with time. Second, the evolution of the state of nature, or state as we shall simply call it, is affected by the choice of a control strategy. Third, we can attach a cost to a control strategy and the evolving state to which it gives rise. The underlying problem is to choose a control strategy that minimizes the cost.

In certain cases, problems in the classical calculus of variations ('minimization of an integral functional over arcs and their derivatives') match this description. Here, the independent variable is interpreted as time, the 'state' is the value of the arc at the current time and the control its rate of change. But techniques for their solution provided by this earlier theory fail to take account of the dynamic constraints that are so often encountered today, in engineering, applied science and economics. Here, by 'dynamic constraints' we mean the mathematical relations governing future evolution of the state, which will depend on the control strategy.

Dynamic optimization is the name given to the systematic study of optimization problems with dynamic constraints. General study of optimization problems with dynamic constraints dates from the late 1950s, which saw several crucial advances, one conceptual and two technical. The conceptual advance, due by L. S. Pontryagin et al., was the realization that optimization problems where the dynamic constraint took the form of a controlled differential equation covered a wide range of engineering control problems involving mechanical systems such as space vehicles and, furthermore, was amenable to analysis. As for the two technical advances, one was Pontryagin's maximum principle, a set of necessary conditions for a control strategy to be optimal. The other was dynamic programming, a procedure initiated by R. Bellman, which reduces the search for an optimal strategy to finding the solution to a partial differential equation (the Hamilton Jacobi equation).

'Dynamic optimization' is synonymous with 'optimal control'. We have chosen the nomenclature dynamic optimization in this book, to convey the idea that optimization problems with general dynamic constraints merit study in their own right and that the field has widespread application, within and beyond engineering. We seek then to avoid the specificity of 'optimal control', a name introduced to describe a branch of control engineering, in which the control design objectives are expressed in terms minimizing a cost, rather than, say, in terms of stability and robustness requirements.

From the mid-1970s, it became apparent that progress in the study of dynamic optimization problems was being impeded by a lack of suitable analytic tools for investigating local properties of functions which are nonsmooth, i.e. not differentiable in the traditional sense. Nonsmooth functions were encountered at first attempts to put Dynamic Programming on a rigorous footing, specifically attempts to relate value functions and solutions to the Hamilton Jacobi equation. It was found that, for many dynamic optimization problems of interest, the only 'solutions' to the Hamilton Jacobi equation have discontinuous derivatives. How should we interpret these solutions? New ideas were required to answer this question since the Hamilton Jacobi equation of dynamic optimization is a nonlinear partial differential equation for which traditional interpretations of generalized solutions, based on the distributions they define, are inadequate.

Nonsmooth functions surfaced once again when efforts were made to extend the applicability of necessary conditions such as the maximum principle. A notable feature of the maximum principle (and one which distinguishes it from necessary conditions derivable using classical techniques) is that it can take account of pathwise constraints on values of the control functions. For some practical problems, the constraints on values of the control depend on the vector state variable. In flight mechanics, for example, the maximum and minimum thrust of a jet engine (a control variable) will depend on the altitude (a component of the state vector). The maximum principle in its original form is not, in general, valid for problems involving state-dependent control constraints. One way to derive necessary conditions for these problems, and others not covered by the maximum principle, is to reformulate them as generalized problems in the calculus of variations, the cost integrands for which include penalty terms to take account of the constraints. The reformulation comes at a price, however. To ensure equivalence with the original problems, it is necessary to employ penalty terms with discontinuous derivatives. So the route to necessary conditions via generalized problems in the calculus of variations can be followed only if we know how to adapt traditional necessary conditions to allow for nonsmooth cost integrands.

Two important breakthroughs occurred in the 1970s. One was the end product of a long quest for effective, local descriptions of 'non-smooth' functions, based on generalizations of the concept of 'subdifferentials' of convex functions, to larger function classes. F. H. Clarke's theory of generalized gradients, by achieving this goal, launched the field of nonsmooth analysis and provided a bridge to necessary conditions of optimality for nonsmooth variational problems (and in particular dynamic optimization problems reformulated as generalized problems in the calculus of variations). The other breakthrough, a somewhat later development, was the concept of viscosity solutions, due to M. G. Crandall and P.-L. Lions, which provides a framework for proving existence and uniqueness of generalized solutions to Hamilton Jacobi equations arising in dynamic optimization.

Nonsmooth analysis and viscosity methods were introduced to overcome obstacles in dynamic optimization. But they have come to have a significant impact on nonlinear analysis as a whole. Nonsmooth analysis provides an important new perspective: useful properties of functions, even differentiable functions, can be proved by examining related nondifferentiable functions, in the same way that trigonometric identities relating to real numbers can sometimes simply be derived by a temporary excursion into the field of complex numbers. Viscosity methods, on the other hand, provide a fruitful approach to studying generalized solutions to broad classes of nonlinear partial differential equations which extend beyond Hamilton Jacobi equations of dynamic optimization and their approximation for computational purposes. The calculus of variations (in its modern guise as dynamic optimization) continues to uphold a long tradition then, as a stimulus to research in other areas of mathematics.

The main purpose of this book is to bring together as a single comprehensive, up-to-date publication major advances in the theory dynamic optimization, with emphasis on those accomplished through the use of nonsmooth analytical techniques. Necessary conditions receive special attention. But other topics are covered as well. Material on the important topic of minimizer regularity provides a showcase for the application of nonsmooth necessary conditions to derive qualitative information about solutions to variational problems. The chapter on dynamic programming stands a little apart from other sections of the book, as it is complementary to mainstream research in the area based on viscosity methods (and which in any case is the subject matter of a number of substantial expository texts). Instead we concentrate on aspects of dynamic programming well matched to the analytic techniques of this book, notably the characterization (in terms of the Hamilton Jacobi equation) of extended-valued value functions associated with problems having endpoint and state constraints, inverse verification theorems, sensitivity relationships and links with the maximum principle.

A subsidiary purpose is to meet the needs of readers with little prior exposure to modern dynamic optimization who seek quick answers to the questions: what are the main results, what were the deficiencies of the 'classical' theory and to what extent have they been overcome? Chapter 1 provides, for their benefit, a lengthy overview, in which analytical details are suppressed and the emphasis is placed instead on communicating the underlying ideas.

To render this book self-contained, preparatory chapters are included on nonsmooth analysis, measurable multifunctions and differential inclusions. Much of this material is implicit in the books of R. T. Rockafellar and J. B. Wets [177] and Clarke et al. [85], and of J.-P. Aubin and H. Frankowska [14]. It is expected, however, that readers, whose main interest is in optimization rather than in broader application areas of nonsmooth analysis which require additional techniques, will find these chapters helpful, because of the strong focus on topics relevant to optimization.

Dynamic optimization is a large field and the choice of material for this is necessary selective. The techniques used here to derive necessary conditions of optimality are, for the most part, within a tradition of research pioneered and developed by Clarke, Ioffe, Loewen, Mordukhovich, Rockafellar, Vinter and others, based on perturbation, elimination of constraints and passage to the limit. The necessary conditions are 'state of the art', as far as this tradition is concerned. Alternative approaches, based on set separation ideas, also make an appearance, but principally for comparison purposes and historical perspective. We do not enter into the topic of higher order necessary conditions nor computational aspects of dynamic optimization.

This book is similar in structure and content to the 2000 book *Optimal Control* [194]. It brings up to date this earlier publication by, in many instances, providing new, simpler proofs of key theorems, where these have become available, and by broadening the applicability of the theory. We provide, for the first time in book form, recent improvements to necessary conditions of optimality for problems for dynamic optimization problems involving a differential inclusion constraint, referred to as the Ioffe refinement. It draws on recent research developments, unavailable at the time of the earlier publication, to provide a thorough discussion and analysis of necessary conditions in the form of Clarke's Hamiltonian inclusion. The book includes new material on necessary conditions for problems with mixed state/control constraints and on problems with free end-times, drawing on latest research in these areas. Also included is a new framework for dynamic programming treating dynamic constraints with discontinuous time dependence.

Brest, France — Piernicola Bettiol
London, UK — Richard B. Vinter

Acknowledgements

We wish to thank our colleagues at Imperial College, including Martin Clarke, David Angeli and Alessandro Astolfi, and at University of Brest, including Rainer Buckdahn, Chloé Jimenez, Vuk Milisic, Marc Quincampoix and Miloud Sadkane for creating working environments in which writing this book has been possible. Many people have influenced our thinking on the contents and presentation of this book, which is indeed the product of many years of collaboration with colleagues and graduate students in the control systems and applied analysis communities. We make special mention of the insights we have gained from Francis Clarke, whose pioneering work in nonsmooth analysis, perturbation methods and applications in nonsmooth optimization laid the foundation for the work reported here. An incomplete list of our collaborators is Zvi Artstein, Aram Arutyunov, Julien Bernis, Andrea Boccia, Frédéric Bonnans, Bernard Bonnard, Alberto Bressan, Arrigo Cellina, Giovanni Colombo, Maria do Rosário de Pinho, Asen Dontchev, Paola Falugi, Margarida Ferreira, Fernando Fontes, Helene Frankowska, Alexander Ioffe, Nathalie Khalil, Yuri Ledyaev, Philip Loewen, Carlo Mariconda, Helmut Maurer, Monica Motta, Michele Palladino, Fernando Lobo Pereiera, Franco Rampazzo, Alain Rapaport, Ralph Tyrrell Rockafellar, Javier Rosenblueth, Jérémy Rouot, Geraldo Silva, Peter Wolenski, and Harry Zheng.

Finally, and most importantly,

'I express my profound gratitude to my wife Giorgia who has provided unwavering and irreplaceable support in my projects, including this book. She and my wonderful children (Eloïse, Gabriele and Leonardo) with their presence, enthusiasm and understanding made it possible for me to dedicate time and effort to bring this book to fruition. I would like to extend my heartfelt appreciation and recognition to Richard; working alongside him has been an enlightening and inspiring experience.' (Piernicola)

'My wife, Donna, and children, Magdalena, Becky and Hannah, have given me unconditional support and encouragement in everything I have wanted to do. This book was no exception. Writing a book like this takes time; I wish to express my deepest thanks to them, especially to Donna for giving me that time, generously and with such good grace. I add my thanks to Piernicola, both for his friendship and for his indispensable part in our fruitful mathematical collaboration.' (Richard)

Contents

Notation

$\mathbb{B}$	Closed unit ball in Euclidean space
$\lvert x \rvert$	Eulidean norm of x
$a \wedge b$	$\min\{a, b\}$
$a \vee b$	$\max\{a, b\}$
$\mathbb{R}_+$	Non-negative real numbers
$d_C(x)$	Euclidean distance of x from C
$\overset{\circ}{C}$, $\operatorname{int} C$	Interior of C
∂C, $\operatorname{bdy} C$	Boundary of C
$\bar{C}$	Closure of C
$\operatorname{co} D$	Convex hull of D
$\overline{\operatorname{co}}\, D$	Closure of the convex hull of D
$N_C^P(x)$	Proximal normal cone to C at x
$\hat{N}_C(x)$	Strict normal cone to C at x
$N_C(x)$	Limiting normal cone to C at x
$T_C(x)$	Bouligand tangent cone to C at x
$\bar{T}_C(x)$	Clarke tangent cone to C at x
$\partial^P f(x)$	Proximal subdifferential of f at x
$\hat{\partial} f(x)$	Strict subdifferential of f at x
$\partial f(x)$	Limiting subdifferential of f at x
$\partial_P^\infty f(x)$	Asymptotic proximal subdifferential of f at x
$\hat{\partial}^\infty f(x)$	Asymptotic strict subdifferential of f at x
$\partial^\infty f(x)$	Asymptotic limiting subdifferential of f at x
$\operatorname{dom} f$	(Effective) domain of f
$\operatorname{Gr} F$	Graph of F
$\operatorname{epi} f$	Epigraph of f

$f^0(x; v)$	Generalized directional derivative of f at x in the direction v
$\Psi_C(x)$	Indicator function of the set C at the point x
$\nabla f(x)$	Gradient vector of f at x
$x_i \stackrel{C}{\to} x$	$x_i \to x$ and $x_i \in C \quad \forall i$
$x_i \stackrel{f}{\to} x$	$x_i \to x$ and $f(x_i) \to f(x)$
supp μ	Support of the measure μ
$\mathcal{L}$	Lebesgue subsets of $I \subset \mathbb{R}$
$\mathcal{B}^k$	Borel subsets of $\mathbb{R}^k$
$L^1(I; \mathbb{R}^n)$	Integrable functions $f : I \to \mathbb{R}^n$
$L^1(S, T)$	Integrable functions $f : [S, T] \to \mathbb{R}$
$W^{1,1}(I; \mathbb{R}^n)$	Absolutely continuous functions $f : I \to \mathbb{R}^n$
$NBV([S, T]; \mathbb{R}^n)$	Functions of bounded variation $f : [S, T] \to \mathbb{R}^n$, right continuous on (S, T)
$H, \mathcal{H}$	Hamiltonian, un-maximized Hamiltonian

Chapter 1
Overview

Abstract Dynamic optimization emerged as a distinct field of research in the late 1950's, to address new kinds of optimization problems, in aerospace, economics and other areas. The distinctive feature of these problems was an underlying dynamic constraint, typically in the form of a controlled differential equation, which placed these problems beyond the scope of earlier variational techniques. Rapid advances were made in the 1970's and 80's, with the discovery of the maximum principle and methodologies (dynamic programming) that linked optimal strategies and the Hamilton Jacobi equation. These were the main elements in what, today, is known as the classical theory of dynamic optimization. While classical dynamic optimization was adequate for many applications, deficiencies became apparent, leading to a new body of theory in the 1980's, including Clarke's nonsmooth maximum principle and generalized solutions of Hamilton-Jacobi equations, based on techniques of nonsmooth analysis.

The purpose of this overview chapter is twofold. First, it provides a self-contained exposition of the classical theory suitable for a first course in dynamic optimization (at undergraduate or graduate level). It includes motivating examples, a derivation of the classical maximum principle, optimality conditions of dynamic programming type expressed in terms of solutions to the Hamilton Jacobi equation, and extensive discussion. Second, it gives answers to the questions: what are the shortcomings of the classical theory and how are they surmounted by more recent developments? We argue that many of the deficiencies of the earlier theory arise from the lack of appropriate analytic techniques for constructing useful local approximations of non-differentiable functions and closed sets with irregular boundaries. We then cover rudiments of nonsmooth analysis, which was developed precisely for this purpose, and show how we can use it to derive new, improved optimality conditions, unshackled by the restrictive hypotheses of classical dynamic optimization. We thereby offer readers the 'big picture' in preparation for later chapters, and also to equip them better to understand the contemporary literature.

P. Bettiol, R. B. Vinter, *Principles of Dynamic Optimization*, Springer Monographs in Mathematics, https://doi.org/10.1007/978-3-031-50089-3_1

1.1 Dynamic Optimization

Dynamic optimization emerged as a distinct field of research in the 1950's, to address in a unified fashion optimization problems arising in scheduling and the control of engineering devices, beyond the reach of earlier analytical and computational techniques. This field was initially called optimal control, but this earlier name is increasingly giving way to dynamic optimization, to convey a wider range of potential application, beyond control engineering. Aerospace engineering is an important source of such problems, and the relevance of dynamic optimization to the American and Russian space programmes gave powerful initial impetus to research in this area. A simple example is:

The Maximum Orbit Transfer Problem A rocket vehicle is in a circular orbit. What is the radius of the largest possible co-planar orbit to which it can be transferred over a fixed period of time? The motion of the vehicle during the manoeuvre is governed by the rocket thrust and by the rocket thrust orientation, both of which can vary with time. See the Fig. 1.1. The variables involved are

r = radial distance of vehicle from attracting centre,
u = radial component of velocity,
v = tangential component of velocity,
m = mass of vehicle,
T_r = radial component of thrust,
T_t = tangential component of thrust.

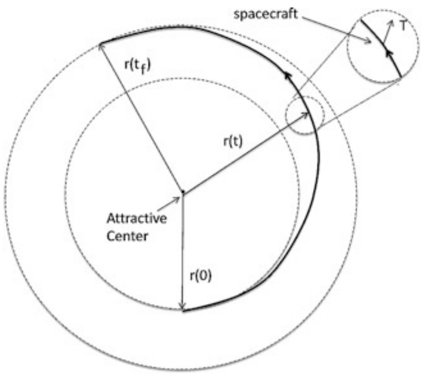

Fig. 1.1 The maximum orbit transfer problem

The constants are

$$\begin{aligned}
r_0 &= \text{initial radial distance},\\
m_0 &= \text{initial mass of vehicle},\\
\gamma_{\max} &= \text{maximum fuel consumption rate},\\
T_{\max} &= \text{maximum thrust},\\
\mu &= \text{gravitational constant of attracting centre},\\
t_f &= \text{duration of manoeuvre}.
\end{aligned}$$

A precise formulation of the problem, based on an idealized point mass model of the space vehicle, is as follows:

$$\begin{cases}
\text{Minimize } -r(t_f)\\
\text{over radial and tangential components of the thrust history,}\\
\qquad\qquad (T_r(t), T_t(t)),\ 0 \le t \le t_f, \text{ satisfying}\\
\dot{r}(t) = u,\\
\dot{u}(t) = v^2(t)/r(t) - \mu/r^2(t) + T_r(t)/m(t),\\
\dot{v}(t) = -u(t)v(t)/r(t) + T_t(t)/m(t),\\
\dot{m}(t) = -(\gamma_{\max}/T_{max})(T_r^2(t) + T_t^2(t))^{1/2},\\
(T_r^2(t) + T_t^2(t))^{1/2} \le T_{\max},\\
m(0) = m_0,\ r(0) = r_0,\ u(0) = 0,\ v(0) = \sqrt{\mu/r_0},\\
u(t_f) = 0,\ v(t_f) = \sqrt{\mu/r(t_f)}.
\end{cases}$$

Here $\dot{r}(t)$ denotes $dr(t)/dt$, etc. It is standard practice in dynamic optimization to formulate optimization problems as minimization problems. Accordingly, the problem of maximizing the radius of the terminal orbit $r(t_f)$ is replaced by the equivalent problem of minimizing the 'cost' $-r(t_f)$. Notice that knowledge of the *control function* or *strategy* $(T_r(t), T_t(t))$, $0 \le t \le t_f$ permits us to calculate the cost $-r(t_f)$: we solve the differential equations, for the specified boundary conditions at time $t = 0$, to obtain the corresponding *state trajectory* $(r(t), u(t), v(t), m(t))$, $0 \le t \le t_f$, and thence determine $-r(t_f)$. The control strategy therefore has the role of choice variable in the optimization problem. We seek a control strategy which minimizes the cost, from among the control strategies whose associated state trajectories satisfy the specified boundary conditions at time $t = t_f$.

For the following values of relevant dimensionless parameters:

$$\frac{T_{\max}/m_0}{\mu/r_0^2} = 0.1405,\quad \frac{\gamma_{\max}}{T_{\max}/\sqrt{\mu/r_0}} = 0.07487,\quad \frac{t_f}{\sqrt{r^3/\mu}} = 3.32\ ,$$

the radius of the terminal circular orbit is

$$r(t_f) \;=\; 1.5\, r_0.$$

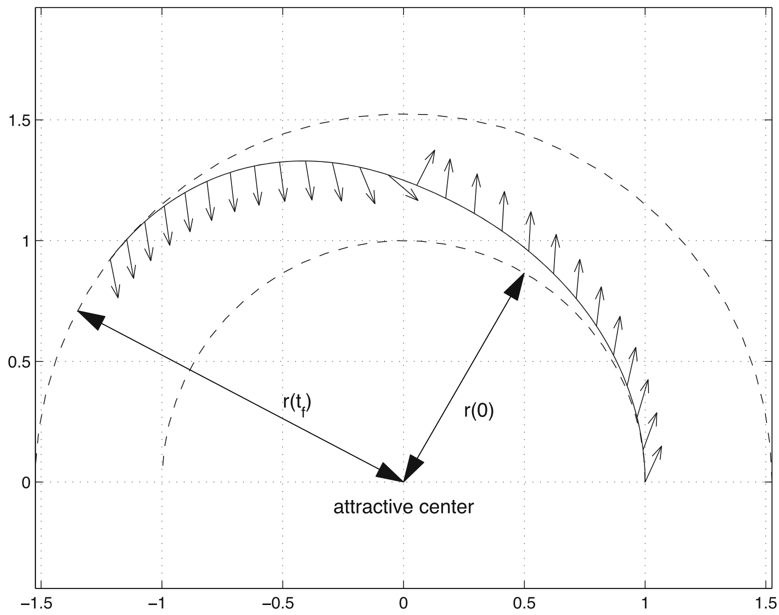

Fig. 1.2 A control strategy for the maximum orbit transfer problem

In Fig. 1.2, the arrows indicate the magnitude and orientation of the thrust at times $t = 0, 0.1t_f, 0.2t_f, \ldots, t_f$. As indicated, full thrust is maintained. The thrust is outward for (approximately) the first half of the manoeuvre and inward for the second.

Suppose, for example, that the attracting centre is the Sun, the space vehicle weighs 10,000 lb, the initial radius is 1.50 million miles (the radius of a circle approximating the Earth's orbit), the maximum thrust is 0.85 lb (i.e. a force equivalent to the gravitational force on a 0.85 lb mass on the surface of the earth, which corresponds to $T_{\max} = 3.778\,\text{N}$, the maximum rate of fuel consumption is 1.81 lb/day and the transit time is 193 days. Corresponding values of the constants are

$$\begin{aligned} &T_{\max} = 3.778\,\text{N}, && m_0 = 4.536 \times 10^3\,\text{kg}, \\ &r_0 = 1.496 \times 10^{11}\,\text{m}, && \gamma_{\max} = 0.9496 \times 10^{-5}\,\text{kg}\,\text{s}^{-1}, \\ &t_f = 1.6675 \times 10^7\,\text{s}, && \mu = 1.32733 \times 10^{20}\,\text{m}^3\,\text{s}^{-2}. \end{aligned}$$

Then the terminal radius of the orbit is 2.44 million miles. (This is the radius of a circle approximating the orbit of the planet Mars.)

Numerical methods, inspired by necessary conditions of optimality akin to the maximum principle of Chap. 7, were used to generate the above control strategy.

Optimal Control of a Growth/Consumption Model Dynamic optimization problems are encountered also in the field of economics. One example is the 'growth versus consumption' problem of neoclassical macro-economics, based on the Ramsey model of economic growth. The question here is, what balance should be struck between investment and consumption to maximize overall spending on social programmes over a fixed period time? A simple formulation of the problem is as follows.

$$\begin{cases} \text{Minimize } -\int_0^T (1-u(t))x^\alpha(t)dt \\ \text{subject to} \\ \dot{x}(t) = -ax(t) + bx^\alpha(t)u(t) \quad \text{for a.e. } t \in [0,T], \\ u(t) \in [0,1] \quad \text{for a.e. } t. \in [0,T]\,, \\ x(t) \geq 0 \;\; \text{for all } t \in [0,T]\,, \\ x(0) = x_0\,. \end{cases}$$

Here, $a > 0$, $b > 0$, $x_0 \geq 0$ and $\alpha \in (0, 1)$ are given constants and $[0, T]$ is a given interval.

It has the following interpretation: x denotes global economic output. The rate of financial return $r(x)$ from economic output x is modelled as

$$r(x) = bx^\alpha\,.$$

The term $-ax$ takes account of fixed costs reducing growth (wages, etc.).

To describe the solution to this problem, we introduce the constants

$$\hat{x} := \left(\frac{\alpha b}{a}\right)^{\frac{1}{1-\alpha}} \text{ and } \Delta := \frac{1}{a\alpha}\ln\left(\frac{1}{1-\alpha}\right)$$

and also the state feedback function $\chi : [0,T] \times (0,\infty) \to [0,1]$:

$$\chi(t,x) := \begin{cases} 0 & \text{if } x > \bar{y}(t) \\ 1 & \text{if } x < \bar{y}(t) \\ \alpha & \text{if } x = \bar{y}(t) \text{ and } t \leq T-\Delta \\ 0 & \text{if } x = \bar{y}(t) \text{ and } t > T-\Delta\,, \end{cases}$$

in which $\bar{y} : (-\infty, T] \to (0, \infty)$ is the function

$$\bar{y}(t) := \begin{cases} \hat{x} & \text{if } t \leq T-\Delta \\ \left[\frac{b}{a}(1-e^{-a\alpha(T-t)}\right]^{\frac{1}{1-\alpha}} & \text{if } t > T-\Delta\,. \end{cases}$$

Techniques of dynamic programming covered in Chap. 13 provide the following solution to this problem:

Given arbitrary initial data $(t_0, x_0) \in [0,T] \times (0,\infty)$*, the optimal output* x^**is the unique solution in the space of Lipschitz continuous functions on* $[t_0, T]$ *of the differential equation*

$$\begin{cases} \dot{x}^*(t) = -ax^*(t) + bx^{*\alpha}(t)\chi(t, x^*(t)) \ \text{ a.e. } t \in [t_0, T], \\ x(t_0) = x_0 \,. \end{cases} \tag{1.1.1}$$

The optimal proportion of financial return for investment $u^*$is unique (w.r.t. the equivalence class of almost everywhere equal functions) and is given by

$$u^*(t) = \chi(t, x^*(t)), \ \text{ for a.e. } t \in [t_0, T] \,.$$

Notice that the solution above is expressed in *state feedback* form; that is, the optimal control u^* is expressed as a function of the current state. For any given initial state and time t_0, the optimal state expressed as a function of time, i.e. in *open loop form*, is the solution to the 'closed loop' state equation (1.1.1) for the given initial state and time t_0. We then obtain the optimal control as a function of time (open loop form) by plugging the optimal state trajectory into the state feedback function. Notice that the feedback form captures, within a single relation, the optimal strategies for every initial state x_0 and time t_0.

Intuition would suggest that if, at the start of the time interval, economic output is low, the optimal control should have a first phase of maximum investment during which economic output builds up to some critical value, followed by a second phase of intermediate investment over which economic output is maintained and, finally, a third phase over which there is no investment because the remaining time is too small for the benefits of investment to show through. This is indeed the optimal control, with the qualification that, if the initial output is high, the optimal control is pure consumption in the first phase. There are also values of the initial investment and T such that there is no first phase or no first and second phase. Analysis is required, of course, to determine precisely the times separating the phases, the critical value of output and the proportion of financial return for investment required to maintain it; also to identify the situations when there are fewer than three phases. Optimal state trajectories, for various choices of initial data, are illustrated in Fig. 1.3.

Optimal Control in Anti-Cancer Treatment

We illustrate applications of dynamic optimization in medicine. Chemotherapy is a treatment aimed at destroying cancer cells by means of a cocktail of drugs,

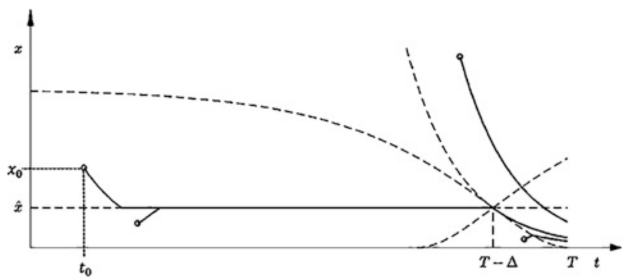

Fig. 1.3 Optimal trajectories for the consumption/growth problem

administered either at specific times or continuously. It is typically part of a complex overarching treatment plan, in which chemotherapy is following up by procedures, surgical or drug-based, for inhibiting renewed tumour growth.

Traditionally, chemotherapy treatments have been based on the maximum tolerated dose paradigm. But a side effect of chemotherapy, a ‘two-edged sword’, is damage to normal cells. Modern day treatments aim to improve outcomes by balancing destruction of cancer cells and suppression of side effects. Empirical design of treatment plans based on clinical trials is time consuming and extremely expensive. Mathematical models of the underlying pharmaco-dynamic processes involved have an important role, because they can be used to simulate on the computer the effects of different treatment strategies, simply and at low cost. Dynamic optimization is the appropriate tool for designing optimal treatment strategies based on these models.

The following formulation of treatment planning as a dynamic optimization problem is taken from [188]. The underlying dynamic model involves the time-varying state variable components c_1, c_2, n and w and the control variable u:

$$
\begin{aligned}
c_1 &= \text{concentration of administered anti-cancer drug in plasma,} \\
c_2 &= \text{active drug concentration at the tumour cellular level,} \\
n &= \text{number of tumour cells,} \\
w &= \text{number of white blood cells (WBCs),} \\
u &= \text{drug dosage.}
\end{aligned}
$$

The evolution of the state variable components for some control strategy $u(t)$, $0 \leq t \leq T$, is governed by the differential equations over the fixed time interval $[0, T]$

$$
\left.
\begin{aligned}
\dot{c}_1(t) &= -(k_1 + k_2)c_1(t) + \left(\tfrac{1}{V_1}\right) u(t) \\
\dot{c}_2(t) &= k_{12}\left(\tfrac{V_1}{V_2}\right) c_1(t) - k_2 c_2(t) \\
\dot{n}(t) &= \Lambda \psi(n(t)) - K \max\{c_2(t) - C_{\min}, 0\} \\
\dot{w}(t) &= r_c - V w(t) - \mu w(t) c_1(t)
\end{aligned}
\right\} \qquad (1.1.2)
$$

in which ψ is the function $\psi(n) := n \log_e \left(\frac{\theta}{n}\right)$.

The initial conditions on state variable components are

$$c_1(0) = 0, c_2(0) = 0, n(0) = n_0 \text{ and } w(0) = w_0.$$

The first differential equation relates the administered drug concentration to the drug dosage. The second relates the active drug concentration to the administered drug concentration. The third is a Gompertz-type differential equation governing tumour growth with an exogenous term to account for the suppressive effects of the active drug concentration. The fourth determines how the WBC population, whose decrease reflects chemotherapy toxicity, responds to the administered drug concentration.

The chosen values of parameters in the model are as in Table 1.1.

Table 1.1 Values of the parameters in the model

Par.	Description	Value
V_1	Volume of distribution in first compartment	25 litres
V_2	Volume of distribution in second compartment	15 litres
k_1	Process of drug elimination from plasma compartment	1.6 day^{-1}
k_{12}	Link process between two compartments	0.4 day^{-1}
θ	Largest tumour	10^{12}
n_0	Initial size of tumour at $t = 0$	30×10^9 cells
Λ	Gompertz growth parameter for tumour	3×10^{-3} day^{-1}
$C_{\min}$	Threshold below which no tumour cells are killed	0.0001 gml^{-1}
K	Rate of cell killing	30 g^{-1} litres day^{-1}
μ	Delayed toxicity of drug concentration on WBCs	80 g^{-1} litres day^{-1}
w_0	Initial physiology level of WBCs at $t = 0$	8×10^9 litres $^{-1}$
V	Nominal turnover constant	0.15 day^{-1}
r_c	Rate of WBCs production	0.2×10^9 litre^{-1} day^{-1}
$C_{\max}$	Maximum allowable drug concentration	0.01 gml^{-1}
W_D	Absolute leukopenia level	2×10^9 litres^{-1}
T	Terminal time	40 days

The control problem is to minimize a weighted sum of the tumour volume and the total amount of drug, subject to upper and lower bounds at each time on drug toxicity and white blood cell population respectively, both of which affect the patient's health.

$$\begin{cases} \text{Minimize } \int_0^T (\alpha n(t) + u(t))dt, \\ \text{over control strategies } u : [0, T] \to \mathbb{R} \\ \qquad\qquad\qquad\qquad \text{and state trajectories } (c_1, c_2, n, w) \\ \text{satisfying} \\ u(t) \in [0, 1] \text{ for } t \in [0, T], \\ c_1(t) \leq C_{\max} \text{ for } t \in [0, T], \\ w(t) \geq W_D \text{ for } t \in [0, T]. \end{cases}$$

The upper and lower bounds, $C_{\max}$ and W_D, are as given in Table 1.1.

In [188] a combination of analytical and computational techniques are employed to determine a control strategy $\bar{u}$ which satisfies necessary conditions of optimality, when the weighting factor is chosen to be $\alpha = (3/5) \times 10^{-10}$. The control $\bar{u}$ gives rise to a quite complicated, 5-subarc state trajectory structure, involving two short bang-bang pulses and three subarcs, in each of which $u(t)$ takes a constant value.

By neglecting the bang-bang pulses, we arrive at a simpler, and therefore more practical, drug treatment strategy, with only slightly increased cost. See Fig. 1.4. According to this strategy the dosage is held constant at a higher level over an initial period, reduced to a lower level for a subsequent period and finally reduced to 0 for the final period:

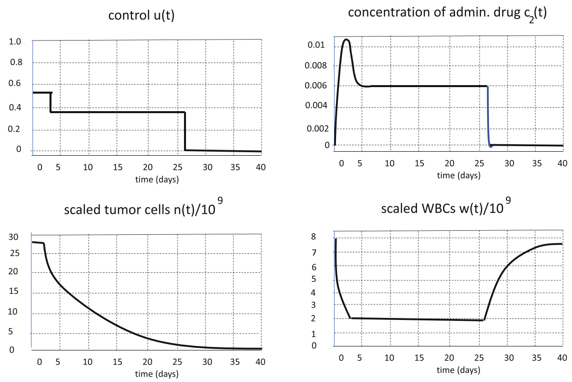

Fig. 1.4 Optimized drug strategy

$$\bar{u}(t) = \begin{cases} u_1 & \text{for } 0 \le t < t_1 \\ u_2 & \text{for } t_1 \le t < t_2 \\ 0 & \text{for } t_2 \le t < 40. \end{cases}$$

Here, $t_1 = 2.2786$ days, $t_2 = 25.855$ days, $u_1 = 0.60254$ and $u_2 = 0.33737$.

1.2 The Calculus of Variations

From a mathematical perspective, dynamic optimization is an outgrowth of the calculus of variations (in one independent variable) that takes account of new kinds of constraints (differential equation constraints, pathwise constraints on control functions 'parameterizing' the differential equations, etc.) encountered in advanced engineering design and dynamic decision making. A number of key developments in dynamic optimization have resulted from marrying old ideas from the calculus of variations and modern analytical techniques. For purposes both of setting dynamic optimization in its historical context and of illuminating later developments in dynamic optimization, we pause to review relevant material from the classical calculus of variations.

The basic problem in the calculus of variations is that of finding an arc $\bar{x}$ which minimizes the value of an integral functional

$$J(x) = \int_S^T L(t, x(t), \dot{x}(t))dt$$

over some class of arcs satisfying the boundary condition

$$x(S) = x_0 \text{ and } x(T) = x_1.$$

Here $[S, T]$ is a given interval, $L : [S, T] \times \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ is a given function, and x_0 and x_1 are given points in $\mathbb{R}^n$.

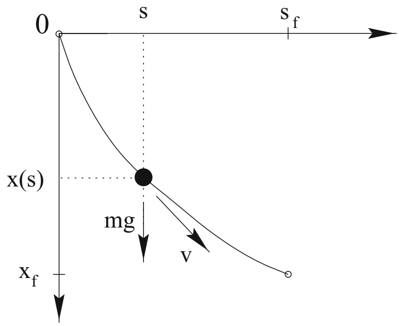

Fig. 1.5 The Brachistochrone problem

The Brachistochrone Problem An early example of such a problem was the *brachistochrone problem* circulated by Johann Bernoulli in the late seventeenth century. Positive numbers s_f and x_f are given. A frictionless bead, initially located at the point (0, 0), slides along a wire under the force of gravity. The wire, which is located in a fixed vertical plane, joins the points (0, 0) and (s_f, x_f). What should the shape of the wire be, in order that the bead arrives at its destination, the point (s_f, x_f), in minimum time? See Fig. 1.5.

There are a number of possible formulations of this problem. We now describe one of them. Denote by s and x the horizontal and vertical distances of a point on the path of the bead (vertical distances are measured downward). We restrict attention to wires describable as the graph of a suitably regular function $x(s)$, $0 \leq s \leq s_f$. For any such function x, the speed $v(s)$ is related to the downward displacement $x(s)$, when the horizontal displacement is s, according to

$$mgx(s) = \frac{1}{2}mv^2(s) \tag{1.2.1}$$

('loss of potential energy equals gain of kinetic energy'). For any $s \in [0, s_f]$, we denote by $t(s)$ the time elapsed when the position of the bead is $(s, x(s))$. If it is assumed that speed v is positive valued, the functions t and v are related by the equation

$$v(s)\,\frac{dt}{ds}(s) = \sqrt{1 + |\frac{dx}{ds}(s)|^2}\,, \qquad \text{for } t \in [0, s_f]\,.$$

Denote by t_f the transit time: $t_f = t(s_f)$. The change of independent variable $t(s) = \int_0^s v^{-1}(s')\,\sqrt{1 + |dx(s')/ds)|^2}\,ds'$ now gives the following formula for t_f:

$$t_f = \int_0^{t_f} dt = \int_0^{s_f} \frac{\sqrt{1 + |dx(s)/ds|^2}}{v(s)} ds.$$

Using (1.2.1) to eliminate $v(s)$, we arrive at a formula for the transit time:

$$J(x) = \int_0^{s_f} L(s, x(s), \dot{x}(s))ds,$$

in which

$$L(s, x, w) := \frac{\sqrt{1 + |w|^2}}{\sqrt{2gx}}.$$

The problem is to minimize $J(x)$ over some class of arcs x satisfying

$$x(0) = 0 \text{ and } x(s_f) = x_f.$$

This is an example of the basic problem of the calculus of variations, in which $(S, x_0) = (0, 0)$ and $(T, x_1) = (s_f, x_f)$. Suppose that we seek a minimizer in the class of absolutely continuous arcs. It can be shown that the minimum time t^* and the minimizing arc $(x(t), s(t))$, $0 \leq t \leq t^*$ (expressed in parametric form with independent variable time t) are given by the formulae

$$x(t) = a\left(1 - \cos\sqrt{\frac{g}{a}}\,t\right) \quad \text{and} \quad s(t) = a\left(\sqrt{\frac{g}{a}}\,t - \sin\sqrt{\frac{g}{a}}\,t\right).$$

Here, a and t^* are constants which uniquely satisfy the conditions

$$\begin{aligned} x(t^*) &= x_f, \\ s(t^*) &= t_f, \\ 0 \leq \sqrt{\frac{g}{a}}\,t^* &\leq 2\pi \ . \end{aligned}$$

The minimizing curve is a cycloid, with infinite slope at the point of departure: it coincides with the locus of a point on the circumference of a disc of radius a, which rolls without slipping along a line of length t_f.

Problems of this kind, the minimization of integral functionals, may perhaps have initially attracted attention as individual curiosities. But throughout the eighteenth and nineteenth centuries their significance became increasingly evident, as the list lengthened of laws of physics which identified states of nature with minimizing curves and surfaces. Some examples of rules of the minimum are as follows:

Fermat's Principle in Optics The path of a light ray achieves a local minimum of the transit times over paths between specified end-points which visit the relevant reflecting and refracting boundaries. The principle predicts Snell's Laws of Reflection and Refraction, and the curved paths of light rays in inhomogeneous media. See Fig. 1.6.

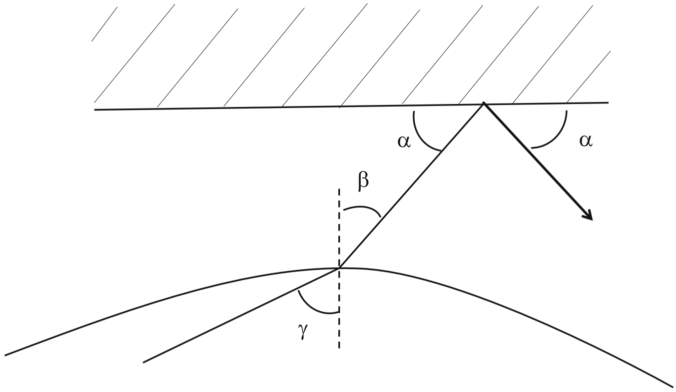

Fig. 1.6 Fermat's principle predicts Snell's laws

Dirichlet's Principle Take a bounded, open set $\Omega \subset \mathbb{R}^2$ with boundary $\partial\Omega$, in which a static two-dimensional electric field is distributed. Denote by $V(x)$ the voltage at point $x \in \Omega$. Then $V(x)$ satisfies Poisson's equation

$$\begin{aligned} \Delta V(x) &= 0 \;\text{ for }\; x \in \Omega \\ V(x) &= \bar{V}(x) \;\text{ for }\; x \in \partial\Omega . \end{aligned}$$

Here, $\bar{V} : \partial\Omega \to \mathbb{R}$ is a given function, which supplies the boundary data.

Dirichlet's principle characterizes the solution to this partial differential equation as the solution of a minimization problem

$$\begin{cases} \text{Minimize } \int_\Omega \nabla V(x) \cdot \nabla V(x) dx \\ \text{over surfaces } V \text{ satisfying } V(x) = \bar{V}(x) \text{ on } \partial\Omega . \end{cases}$$

This optimization problem involves finding a *surface* which minimizes a given integral functional. See Fig. 1.7.

Dirichlet's principle and its generalizations are important in many respects. They are powerful tools for the study of existence and regularity of solutions to boundary value problems. Furthermore, they point the way to Galerkin methods for computing solutions to partial differential equations, such as Poisson's equation: the solution is approximated by the minimizer of the Dirichlet integral above over some finite dimensional subspace $\mathcal{S}_N$ of the domain of the original optimization problem, spanned by a finite collection of 'basis' functions $\{\phi_i\}_{i=1}^N$,

$$\mathcal{S}_N = \{\sum_{i=1}^N \alpha_i \phi_i(x) : \alpha \in \mathbb{R}^N\} .$$

The widely used finite element methods are modern implementations of Galerkin's method.