1 Introduction

Our nature consists in movement;
absolute rest is death.
—Blaise Pascal

Change is all around us. Dynamic strategies seek to both anticipate and effect such change in a given system so as to accomplish objectives of an individual, a group of agents, or a social planner. This book offers an introduction to continuous-time systems and methods for solving dynamic optimization problems at three different levels: single-person decision making, games, and mechanism design. The theory is illustrated with examples from economics. Figure 1.1 provides an overview of the book's hierarchical approach.

The first and lowest level, single-person decision making, concerns the choices made by an individual decision maker who takes the evolution of a system into account when trying to maximize an objective functional over feasible dynamic policies. An example would be an economic agent who is concerned with choosing a rate of spending for a given amount of capital, each unit of which can either accumulate interest over time or be used to buy consumption goods such as food, clothing, and luxury items.

The second level, games, addresses the question of finding predictions for the behavior and properties of dynamic systems that are influenced by a group of decision makers. In this context the decision makers (players) take each other's policies into account when choosing their own actions. The possible outcomes of the game among different players, say, in terms of the players' equilibrium payoffs and equilibrium actions, depend on which precise concept of equilibrium is applied. Nash (1950) proposed an equilibrium such that players' policies do not give any player an incentive to deviate from his own chosen policy, given

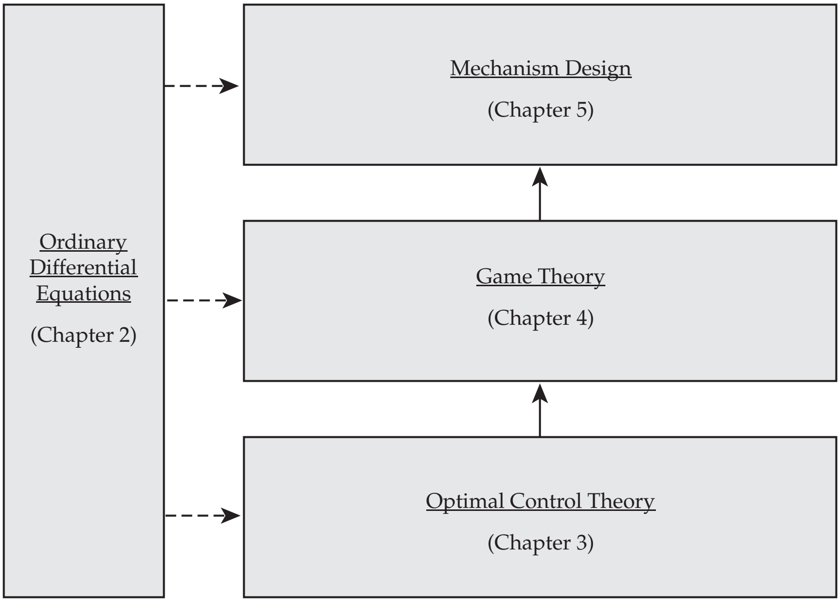

Figure 1.1
Topics covered in this book.

the other players' choices are fixed to the equilibrium policies. A classic example is an economy with a group of firms choosing production outputs so as to maximize their respective profits.

The third and highest level of analysis considered here is mechanism design, which is concerned with a designer's creation of an environment in which players (including the designer) can interact so as to maximize the designer's objective functional. Leading examples are the design of nonlinear pricing schemes in the presence of asymmetric information, and the design of markets. Arguably, this level of analysis is isomorphic to the first level, since the players' strategic interaction may be folded into the designer's optimization problem.

The dynamics of the system in which the optimization takes place are described in continuous time, using ordinary differential equations. The theory of ordinary differential equations can therefore be considered the backbone of the theory developed in this book.

1.1 Outline

Ordinary Differential Equations (ODEs) Chapter 2 reviews basic concepts in the theory of ODEs. One-dimensional linear first-order ODEs can be solved explicitly using the Cauchy formula. The key insight from the construction of this formula (via variation of an integration constant) is that the solution to a linear initial value problem of the form

$$\dot{x} + g(t)x = h(t), \qquad x(t_0) = x_0,$$

for a given tuple of initial data (t_0, x_0) can be represented as the superposition of a homogeneous solution (obtained when $h = 0$) and a particular solution to the original ODE (but without concern for the initial condition). Systems of linear first-order ODEs,

$$\dot{x} = A(t)x + b(t), \tag{1.1}$$

with an independent variable of the form $x = (x_1, \ldots, x_n)$ and an initial condition $x(t_0) = x_0$ can be solved if a fundamental matrix $\Phi(t, t_0)$ as the solution of a homogeneous equation is available. Higher-order ODEs (containing higher-order derivatives) can generally be reduced to first-order ODEs. This allows limiting the discussion to (nonlinear) first-order ODEs of the form

$$\dot{x} = f(t, x), \tag{1.2}$$

for $t \geq t_0$. Equilibrium points, that is, points $\bar{x}$ at which a system does not move because $f(t, \bar{x}) = 0$, are of central importance in understanding a continuous-time dynamic model. The stability of such points is usually investigated using the method developed by Lyapunov, which is based on the principle that if system trajectories $x(t)$ in the neighborhood of an equilibrium point are such that a certain real-valued function $V(t, x(t))$ is nonincreasing (along the trajectories) and bounded from below by its value at the equilibrium point, then the system is stable. If this function is actually decreasing along system trajectories, then these trajectories must converge to an equilibrium point. The intuition for this finding is that the Lyapunov function V can be viewed as energy of the system that cannot increase over time. This notion of energy, or, in the context of economic problems, of value or welfare, recurs throughout the book.

Optimal Control Theory Given a description of a system in the form of ODEs, and an objective functional $J(u)$ as a function of a dynamic policy or control u, together with a set of constraints (such as initial conditions or control constraints), a decision maker may want to solve an *optimal control problem* of the form

$$J(u) = \int_{t_0}^{T} h(t, x(t), u(t))\, dt \;\longrightarrow\; \max_{u(\cdot)}, \tag{1.3}$$

subject to $\dot{x}(t) = f(t, x(t), u(t))$, $x(t_0) = x_0$, and $u \in \mathcal{U}$, for all $t \in [t_0, T]$. Chapter 3 introduces the notion of a controllable system, which is a system that can be moved using available controls from one state to another. Then it takes up the construction of solutions (in the form of state-control trajectories $(x^*(t), u^*(t))$, $t \in [t_0, T]$) to such optimal control problems: necessary and sufficient optimality conditions are discussed, notably the Pontryagin maximum principle (PMP) and the Hamilton-Jacobi-Bellman (HJB) equation. Certain technical difficulties notwithstanding, it is possible to view the PMP and the HJB equation as two complementary approaches to obtain an understanding of the solution of optimal control problems. In fact, the HJB equation relies on the existence of a continuously differentiable value function $V(t, x)$, which describes the decision maker's optimal payoff, with the optimal control problem initialized at time t and the system in the state x. This function, somewhat similar to a Lyapunov function in the theory of ODEs, can be interpreted in terms of the value of the system for a decision maker. The necessary conditions in the PMP can be informally derived from the HJB equation, essentially by restricting attention to a neighborhood of the optimal trajectory.

Game Theory When more than one individual can make payoff-relevant decisions, game theory is used to determine predictions about the outcome of the strategic interactions. To abstract from the complexities of optimal control theory, chapter 4 introduces the fundamental concepts of game theory for simple discrete-time models, along the lines of the classical exposition of game theory in economics. Once all the elements, including the notion of a Nash equilibrium and its various refinements, for instance, via subgame perfection, are in place, attention turns to differential games. A critical question that arises in dynamic games is whether the players can trust each other's equilibrium strategies, in the sense that they are credible even after the game has started. A player may, after a while, find it best to deviate from a

Nash equilibrium that relies on a “noncredible threat.” The latter consists of an action which, as a contingency, discourages other players from deviating but is not actually beneficial should they decide to ignore the threat. More generally, in a Nash equilibrium that is not subgame-perfect, players lack the ability to commit to certain threatening actions (thus, noncredible threats), leading to “time inconsistencies.”

Mechanism Design A simple economic mechanism, discussed in chapter 5, is a collection of a message space and an allocation function. The latter is a mapping from possible messages (elements of the message space) to available allocations. For example, a mechanism could consist of the (generally nonlinear) pricing schedule for bandwidth delivered by a network service provider. A mechanism designer, who is often referred to as the principal, initially announces the mechanism, after which the agent sends a message to the principal, who determines the outcome for both participants by evaluating the allocation function. More general mechanisms, such as an auction, can include several agents playing a game that is implied by the mechanism.

Optimal control theory becomes useful in the design of a static mechanism because of an information asymmetry between the principal and the various agents participating in the mechanism. Assuming for simplicity that there is only a single agent, and that this agent possesses private information that is encapsulated in a one-dimensional type variable θ in a type space $\Theta = [\underline{\theta}, \bar{\theta}]$, it is possible to write the principal's mechanism design problem as an optimal control problem.

1.2 Prerequisites

The material in this book is reasonably self-contained. It is recommended that the reader have acquired some basic knowledge of dynamic systems, for example, in a course on linear systems. In addition, the reader should possess a firm foundation in calculus, since the language of calculus is used throughout the book without necessarily specifying all the details or the arguments if they can be considered standard material in an introductory course on calculus (or analysis).

1.3 A Brief History of Optimal Control

Origins The human quest for finding extrema dates back to antiquity. Around 300 B.C., Euclid of Alexandria found that the minimal distance between two points A and B in a plane is described by the straight line $\overline{AB}$, showing in his *Elements* (Bk I, Prop. 20) that any two sides of a triangle together are greater than the third side (see, e.g., Byrne 1847, 20). This is notwithstanding the fact that nobody has actually ever seen a straight line. As Plato wrote in his Allegory of the Cave[1] (*Republic*, Bk VII, ca. 360 B.C.), perceived reality is limited by our senses (Jowett 1881). Plato's theory of forms held that ideas (or forms) can be experienced only as shadows, that is, imperfect images (W. D. Ross 1951). While Euclid's insight into the optimality of a straight line may be regarded merely as a variational inequality, he also addressed the problem of finding extrema subject to constraints by showing in his *Elements* (Bk VI, Prop. 27) that "of all the rectangles contained by the segments of a given straight line, the greatest is the square which is described on half the line" (Byrne 1847, 254). This is generally considered the earliest solved maximization problem in mathematics (Cantor 1907, 266) because

$$\frac{a}{2} \in \arg\max_{x \in \mathbb{R}} \{x(a-x)\},$$

for any $a > 0$. Another early maximization problem, closely related to the development of optimal control, is recounted by Virgil in his *Aeneid* (ca. 20 B.C.). It involves queen Dido, the founder of Carthage (located in modern-day Tunisia), who negotiated to buy as much land as she could enclose using a bull's hide. To solve her isoperimetric problem, that is, to find the largest area with a given perimeter, she cut the hide into a long strip and laid it out in a circle. Zenodorus, a Greek mathematician, studied Dido's problem in his book *On Isoperimetric Figures* and showed that a circle is greater than any regular polygon of equal contour (Thomas 1941, 2:387–395). Steiner (1842) provided five different proofs that any figure of maximal area with a given perimeter in the plane must be a circle. He omitted to show that there actually *exists* a solution to the isoperimetric problem. Such a proof was given later by Weierstrass (1879/1927).[2]

Remark 1.1 (Existence of Solutions) Demonstrating the existence of a solution to a variational problem is in many cases both important and nontrivial. Perron (1913) commented specifically on the gap left by Steiner in the solution of the isoperimetric problem regarding existence,

1. In the Allegory of the Cave, prisoners in a cave are restricted to a view of the real world (which exists behind them) solely via shadows on a wall in front of them.
2. Weierstrass's numerous contributions to the calculus of variations, notably on the existence of solutions and on sufficient optimality conditions, are summarized in his extensive lectures on *Variationsrechnung*, published posthumously based on students' notes.

and he provided several examples of variational problems without solutions (e.g., finding a polygon of given perimeter and maximal surface). A striking problem without a solution was posed by Kakeya (1917). He asked for the set of minimal measure that contains a unit line segment in all directions. One can think of such a Kakeya set (or Besicovitch set) as the minimal space that an infinitely slim car would need to turn around in a parking spot. Somewhat surprisingly, Besicovitch (1928) was able to prove that the measure of the Kakeya set cannot be bounded from below by a positive constant. □

The isoperimetric constraint appears naturally in economics as a budget constraint, which was recognized by Frisi in his written-in commentary on Verri's (1771) notion that a political economy shall be trying to maximize production subject to the available labor supply (Robertson 1949). Such budget-constrained problems are natural in economics.[3] For example, Sethi (1977) determined a firm's optimal intertemporal advertising policy based on a well-known model by Nerlove and Arrow (1962), subject to a constraint on overall expenditure over a finite time horizon.

Calculus of Variations The infinitesimal calculus (or later just *calculus*) was developed independently by Newton and Leibniz in the 1670s. Newton formulated the modern notion of a derivative (which he termed *fluxion*) in his *De Quadratura Curvarum*, published as an appendix to his treatise on *Opticks* in 1704 (Cajori 1919, 17–36). In 1684, Leibniz published his notions of derivative and integral in the *Acta Eruditorum*, a journal that he had co-founded several years earlier and that enjoyed a significant circulation in continental Europe. With the tools of calculus in place, the time was ripe for the calculus of variations, the birth of which can be traced to the June 1696 issue of the *Acta Eruditorum*. There, Johann Bernoulli challenged his contemporaries to determine the path from point A to point B in a vertical plane that minimizes the time for a mass point M to travel under the influence of gravity between A and B. This problem of finding a *brachistochrone* (figure 1.2) was posed

3. To be specific, let $C(t, x, u)$ be a nonnegative-valued cost function and $B > 0$ a given budget. Then along a trajectory $(x(t), u(t))$, $t \in [t_0, T]$, a typical *isoperimetric constraint* is of the form $\int_{t_0}^{T} C(t, x(t), u(t))\, dt \leq B$. It can be rewritten as $\dot{y}(t) = C(t, x(t), u(t))$, $y(t_0) = 0$, $y(T) \leq B$. The latter formulation falls squarely within the general optimal-control formalism developed in this book, so isoperimetric constraints do not need special consideration.

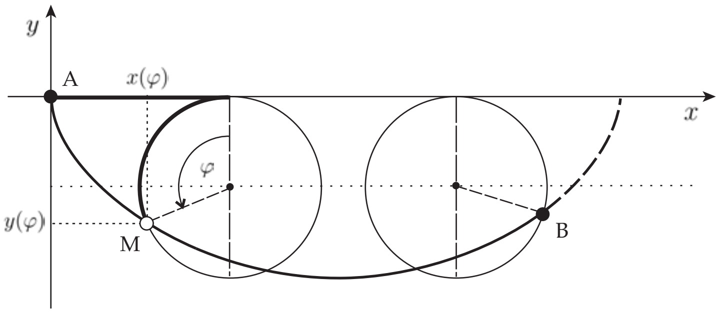

Figure 1.2
Brachistochrone connecting the points A and B in parametric form: $(x(\varphi), y(\varphi)) = (\alpha(\varphi - \sin(\varphi)), \alpha(\cos(\varphi) - 1))$, where $\varphi = \varphi(t) = \sqrt{g/\alpha}\, t$, and $g \approx 9.81$ meters per second squared is the gravitational constant. The parameter α and the optimal time $t = T^*$ are determined by the endpoint condition $(x(\varphi(T^*)), y(\varphi(T^*))) = B$.

earlier (but not solved) by Galilei (1638).[4] In addition to his own solution, Johann Bernoulli obtained four others, by his brother Jakob Bernoulli, Leibniz, de l'Hôpital, and Newton (an anonymous entry). The last was recognized immediately by Johann *ex ungue leonem* ("one knows the lion by his claw").

Euler (1744) investigated the more general problem of finding extrema of the functional

$$J = \int_0^T L(t, x(t), \dot{x}(t))\, dt, \tag{1.4}$$

subject to suitable boundary conditions on the function $x(\cdot)$. He derived what is now called the Euler equation (see equation (1.5)) as a necessary optimality condition used to this day to construct solutions to variational problems. In his 1744 treatise on variational methods, Euler did not create a name for his complex of methods and referred to variational calculus simply as the isoperimetric method. This changed with a 1755 letter from Lagrange to Euler informing the latter of his δ-calculus, with δ denoting variations (Goldstine 1980, 110–114). The name "calculus of variations" was officially born in 1756, when the minutes of meeting

4. Huygens (1673) discovered that a body which is bound to fall following a cycloid curve oscillates with a periodicity that is independent of the starting point on the curve, so he termed this curve *tautochrone*. The brachistochrone is also a cycloid and thus identical to the tautochrone, which led Johann Bernoulli to remark that "nature always acts in the simplest possible way" (Willems 1996).

no. 441 of the Berlin Academy on September 16 note that Euler read "Elementa calculi variationum" (Hildebrandt 1989).

Remark 1.2 (Extremal Principles) Heron of Alexandria explained the equality of angles in the reflection of light by the principle that nature must take the shortest path, for "[i]f Nature did not wish to lead our sight in vain, she would incline it so as to make equal angles" (Thomas 1941, 2:497). Olympiodorus the younger, in a commentary (ca. 565) on Aristotle's *Meteora*, wrote, "[T]his would be agreed by all . . . Nature does nothing in vain nor labours in vain" (Thomas 1941, 2:497).

In the same spirit, Fermat in 1662 used the *principle of least time* (now known as Fermat's principle) to derive the law of refraction for light (Goldstine 1980, 1–6). More generally, Maupertuis (1744) formulated the *principle of least action*, that in natural phenomena a quantity called action (denoting energy × time) is to be minimized (cf. also Euler 1744). The calculus of variations helped formulate more such extremal principles, for instance, *d'Alembert's principle*, which states that along any virtual displacement the sum of the differences between the forces and the time derivatives of the moments vanishes. It was this principle that Lagrange (1788/1811) chose over Maupertuis's principle in his *Mécanique Analytique* to firmly establish the use of differential equations to describe the evolution of dynamic systems. Hamilton (1834) subsequently established that the law of motion on a time interval $[t_0, T]$ can be derived as extremal of the functional in equation (1.4) (*principle of stationary action*), where L is the difference between kinetic energy and potential energy. Euler's equation in this variational problem is also known as the *Euler-Lagrange equation*,

$$\frac{d}{dt}\frac{\partial L(t,x(t),\dot{x}(t))}{\partial \dot{x}} - \frac{\partial L(t,x(t),\dot{x}(t))}{\partial x} = 0, \tag{1.5}$$

for all $t \in [t_0, T]$. With the Hamiltonian function $H(t,x,\dot{x},\psi) = \langle \psi, \dot{x} \rangle - L(t,x,\dot{x})$, where $\psi = \partial L/\partial \dot{x}$ is an adjoint variable, one can show that (1.5) is in fact equivalent to the *Hamiltonian system*,[5]

---

5. To see this, note first that (1.6) holds by definition and that irrespective of the initial conditions,

$$0 = \frac{dH}{dt} - \frac{dH}{dt} = \frac{\partial H}{\partial t} + \left\langle \frac{\partial H}{\partial x}, \dot{x} \right\rangle + \left\langle \frac{\partial H}{\partial \psi}, \dot{\psi} \right\rangle - \left(\langle \dot{\psi}, \dot{x} \rangle + \langle \psi, \ddot{x} \rangle - \frac{\partial L}{\partial t} - \left\langle \frac{\partial L}{\partial x}, \dot{x} \right\rangle - \left\langle \frac{\partial L}{\partial \dot{x}}, \ddot{x} \right\rangle \right),$$

whence, using $\psi = \partial L/\partial \dot{x}$ and $\dot{x} = \partial H/\partial \psi$, we obtain

$$0 = \frac{\partial H}{\partial t} + \frac{\partial L}{\partial t} + \left\langle \frac{\partial H}{\partial x} + \frac{\partial L}{\partial x}, \dot{x} \right\rangle + \left\langle \frac{\partial H}{\partial \psi}, \dot{\psi} \right\rangle - \langle \dot{\psi}, \dot{x} \rangle = \left\langle \frac{\partial H}{\partial x} + \frac{\partial L}{\partial x}, \dot{x} \right\rangle.$$

Thus, $\partial H/\partial x = -\partial L/\partial x$, so the Euler-Lagrange equation (1.5) immediately yields (1.7).

$$\dot{x}(t) = \frac{\partial H(t, x(t), \dot{x}(t), \psi(t))}{\partial \psi}, \tag{1.6}$$

$$\dot{\psi}(t) = -\frac{\partial H(t, x(t), \dot{x}(t), \psi(t))}{\partial x}, \tag{1.7}$$

for all $t \in [t_0, T]$. To integrate the Hamiltonian system, given some initial data (t_0, x_0), Jacobi (1884, 143–157) proposed to introduce an action function,

$$V(t, x) = \int_{t_0}^{t} L(s, x(s), \dot{x}(s))\, ds,$$

on an extremal trajectory, which satisfies (1.6)–(1.7) on $[t_0, t]$ and connects the initial point (t_0, x_0) to the point (t, x). One can now show (see, e.g., Arnold 1989, 254–255) that

$$\frac{dV(t, x(t))}{dt} = \frac{\partial V(t, x(t))}{\partial t} + \frac{\partial V(t, x(t))}{\partial x} = \langle \psi(t), \dot{x}(t) \rangle - H(t, x(t), \dot{x}(t), \psi(t)),$$

so that $H = -\partial V / \partial t$ and $\psi = \partial V / \partial x$, and therefore the *Hamilton-Jacobi equation*,

$$-\frac{\partial V(t, x(t))}{\partial t} = H(t, x(t), \dot{x}(t), \frac{\partial V(t, x(t))}{\partial x}), \tag{1.8}$$

holds along an extremal trajectory. This result is central for the construction of sufficient as well as necessary conditions for solutions to optimal control problems (see chapter 3). Extremal principles also play a role in economics. For example, in a Walrasian exchange economy, prices and demands will adjust so as to maximize a welfare functional. □

Remark 1.3 (Problems with Several Independent Variables) Lagrange (1760) raised the problem of finding a surface of minimal measure, given an intersection-free closed curve. The Euler-Lagrange equation for this problem expresses the fact that the mean curvature of the surface must vanish everywhere. This problem is generally referred to as *Plateau's problem*, even though Plateau was born almost half a century after Lagrange had formulated it originally. (Plateau conducted extended experiments with soap films leading him to discover several laws that were later proved rigorously by others.) Plateau's problem was solved independently by Douglas (1931) and Radó (1930). For historical details see, for instance, Fomenko (1990) and Struwe (1989). This book considers only problems where the independent variable is one-dimensional, so all systems can be described using ordinary (instead of partial) differential equations. □

In an article about beauty in problems of science the economist Paul Samuelson (1970) highlighted several problems in the calculus of variations, such as the brachistochrone problem, and connected those insights to important advances in economics. For example, Ramsey (1928) formulated an influential theory of saving in an economy that determines an optimal growth path using the calculus of variations. The Ramsey model, which forms the basis of the theory of economic growth, was further developed by Cass (1965) and Koopmans (1965).[6]

Feedback Control Before considering the notion of a control system, one can first define a *system* as a set of connected elements, where the connection is an arbitrary relation among them. The complement of this set is the *environment* of the system. If an element of the system is not connected to any other element of the system, then it may be viewed as part of the environment. When attempting to model a real-world system, one faces an age-old trade-off between veracity and usefulness. In the fourteenth century William of Occam formulated the *law of parsimony* (also known as Occam's razor), *entia non sunt multiplicanda sine necessitate*, to express the postulate that "entities are not to be multiplied without necessity" (Russell 1961, 453).[7] The trade-off between usefulness and veracity of a system model has been rediscovered many times, for instance, by Leonardo da Vinci ("simplicity it is the ultimate sophistication") and by Albert Einstein ("make everything as simple as possible, but not simpler").[8]

A *control system* is a system with an input (or *control*) $u(t)$ that can be influenced by human intervention. If the *state* $x(t)$ of the system can also be observed, then the state can be used by a *feedback law* $u(t) = \mu(t, x(t))$ to adjust the input, which leads to a *feedback control system* (figure 1.3).

There is a rich history of feedback control systems in technology, dating back at least to Ktesibios's float regulator in the third century B.C. for a water clock, similar to a modern flush toilet (Mayr 1970). Wedges

6. For more details on the modern theory of economic growth, see, e.g., Acemoglu (2009), Aghion and Howitt (2009), and Weitzman (2003).
7. For a formalization of Occam's razor, see Pearl (2000, 45–48).
8. Some "anti-razors" warn of oversimplification, e.g., Leibniz's *principle of plenitude* ("everything that can happen will happen") or Kant's insight that "[t]he variety of entities is not to be diminished rashly" (1781, 656).

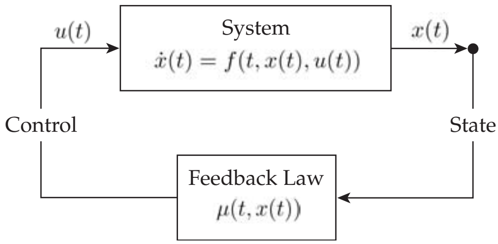

Figure 1.3
Feedback control system.

were inserted in the water flow to control the speed at which a floating device would rise to measure the time. In 1788, Watt patented the design of the centrifugal governor for regulating the speed of a rotary steam engine, which is one of the most famous early feedback control systems. Rotating flyballs, flung apart by centrifugal force, would throttle the engine and regulate its speed. A key difference between the Ktesibios's and Watt's machines is that the former does not use feedback to determine the control input (the number and position of the wedges), which is therefore referred to as *open-loop* control. Watt's flyball mechanism, on the other hand, uses the state of the system (engine rotations) to determine the throttle position that then influences the engine rotations, which is referred to as *closed-loop* (or feedback) control. Wiener (1950, 61) noted that "feedback is a method of controlling a system by reinserting into it the results of its past performance." He suggested the term *cybernetics* (from the Greek word $\kappa\upsilon\beta\epsilon\rho\nu\eta\tau\eta\varsigma$—governor) for the study of control and communication systems (Wiener 1948, 11–12).[9] Maxwell (1868) analyzed the stability of Watt's centrifugal governor by linearizing the system equation and showing that it is stable, provided its eigenvalues have strictly negative real parts. Routh (1877) worked out a numerical algorithm to determine when a characteristic equation (or equivalently, a system matrix) has stable roots. Hurwitz (1895) solved this problem independently, and to this day a stable system matrix A in equation (1.1) carries his name (see lemma 2.2). The stability of nonlinear systems of the form (1.2) was advanced by the seminal work of Lyapunov (1892), which showed that if an energy function $V(t, x)$ could be found such that it is bounded from below and decreasing along any

9. The term was suggested more than a hundred years earlier for the control of sociopolitical systems by Ampère (1843, 140–141).

system trajectory $x(t)$, $t \geq t_0$, then the system is (asymptotically) stable, that is, the system is such that any trajectory that starts close to an equilibrium state converges to that equilibrium state. In variational problems the energy function $V(t,x)$ is typically referred to as a value function and plays an integral role for establishing optimality conditions, such as the Hamilton-Jacobi equation (1.8), or more generally, the Hamilton-Jacobi-Bellman equation (3.16).

In 1892, Poincaré published the first in a three-volume treatise on celestial mechanics containing many path-breaking advances in the theory of dynamic systems, such as integral invariants, Poincaré maps, the recurrence theorem, and the first description of chaotic motion. In passing, he laid the foundation for a geometric and qualitative analysis of dynamic systems, carried forward, among others, by Arnold (1988). An important alternative to system stability in the sense of asymptotic convergence to equilibrium points is the possibility of a limit cycle. Based on Poincaré's work between 1881 and 1885,[10] Bendixson (1901) established conditions under which a trajectory of a two-dimensional system constitutes a limit cycle (see proposition 2.13); as a by-product, this result implies that chaotic system behavior can arise only if the state-space dimension is at least 3. The theory of stability in feedback control systems has proved useful for the description of real-world phenomena. For example, Lotka (1920) and Volterra (1926) proposed a model for the dynamics of a biological predator-prey system that features limit cycles (see example 2.8).

In technological applications (e.g., when stabilizing an airplane) it is often sufficient to linearize the system equation and minimize a cost that is quadratic in the magnitude of the control and quadratic in the deviations of the system state from a reference state (or tracking trajectory)[11] in order to produce an effective controller. The popularity of this linear-quadratic approach is due to its simple closed-form solvability. Kalman and Bucy (1961) showed that the approach can also be very effective in dealing with (Gaussian) noise incorporating a state-estimation component, resulting in a continuous-time version of the Kalman filter, which was first developed by Rudolf Kalman for discrete-time systems. To deal with control constraints in a noisy

10. The relevant series of articles was published in the *Journal de Mathématiques*, reprinted in Poincaré (1928, 3–222); see also Barrow-Green (1997).

11. A linear-quadratic regulator is obtained by solving an optimal control problem of the form (1.3), with linear system function $f(t,x,u) = Ax + Bu$ and quadratic payoff function $h(t,x,u) = -x'Rx - u'Su$ (with R, S positive definite matrices); see example 3.3.

environment, the linear-quadratic approach has been used in receding-horizon control (or model predictive control), where a system is periodically reoptimized over the same fixed-length horizon.[12] More recently, this approach has been applied in financial engineering, for example, portfolio optimization (Primbs 2007).

Optimal Control In the 1950s the classical calculus of variations underwent a transformation driven by two major advances. Both advances were fueled by the desire to find optimal control interventions for given feedback control systems, in the sense that the optimal control trajectory $u^*(t)$, $t \in [t_0, T]$, would maximize an objective functional $J(u)$ by solving a problem of the form (1.3). The first advance, by Richard Bellman, was to incorporate a control function into the Hamilton-Jacobi variational equation, leading to the Hamilton-Jacobi-Bellman equation,[13]

$$-V_t(t,x) = \max_{u \in \mathcal{U}} \{h(t,x,u) + \langle V_x(t,x), f(t,x,u)\rangle\}, \tag{1.9}$$

which, when satisfied on the rectangle $[t_0, T] \times \mathcal{X}$ (where the state space $\mathcal{X}$ contains all the states), together with the endpoint condition $V(T,x) \equiv 0$, serves as a sufficient condition for optimality. The optimal feedback law $\mu(t,x)$ is obtained as the optimal value for u on the right-hand side of (1.9), so the optimal state trajectory $x^*(t)$, $t \in [t_0, T]$, solves the initial value problem (IVP)

$$\dot{x} = f(t,x,\mu(t,x)), \qquad x(t_0) = x_0,$$

which yields the optimal control

$$u^*(t) = \mu(t, x^*(t)),$$

for all $t \in [t_0, T]$. This approach to solving optimal control problems by trying to construct the value function is referred to as dynamic programming (Bellman 1957).[14] The second advance, by Lev Pontryagin and his students, is related to the lack of differentiability of the value function $V(t,x)$ in (1.9), even for the simplest problems (see, e.g., Pontryagin et al. 1962, 23–43, 69–73) together with the difficulties of actually solving the partial differential equation (1.9) when the value function is

---

12. Receding-horizon control has also been applied to the control of nonlinear systems, be they discrete-time (Keerthi and Gilbert 1988) or continuous-time (Mayne and Michalska 1990).
13. Subscripts denote partial derivatives.
14. The idea of dynamic programming precedes Bellman's work: for example, von Neumann and Morgenstern (1944, ch. 15) used backward induction to solve sequential decision problems in perfect-information games.

differentiable. Pontryagin (1962), together with his students, provided a rigorous proof for a set of necessary optimality conditions for optimal control problems of the form (1.3). As shown in section 3.3, the conditions of the Pontryagin maximum principle (in its most basic version) can be obtained, at least heuristically, from the Hamilton-Jacobi-Bellman equation. A rigorous proof of the maximum principle usually takes a different approach, using needle variations introduced by Weierstrass (1879/1927). As Pontryagin et al. (1962) pointed out,

> The method of dynamic programming was developed for the needs of optimal control processes which are of a much more general character than those which are describable by systems of differential equations. Therefore, the method of dynamic programming carries a more universal character than the maximum principle. However, in contrast to the latter, this method does not have the rigorous logical basis in all those cases where it may be successfully made use of as a valuable heuristic tool. (69)

In line with these comments, the Hamilton-Jacobi-Bellman equation is often used in settings that are more complex than those considered in this book, for instance for the optimal control of stochastic systems. The problem with the differentiability of the value function was addressed by Francis Clarke by extending the notion of derivative, leading to the concept of nonsmooth analysis (Clarke 1983; Clarke et al. 1998).[15] From a practical point of view, that is, to solve actual real-world problems, nonsmooth analysis is still in need of exploration. In contrast to this, an abundance of optimal control problems have been solved using the maximum principle and its various extensions to problems with state-control constraints, pure state constraints, and infinite time horizons. For example, Arrow (1968) and Arrow and Kurz (1970a) provided an early overview of optimal control theory in models of economic growth.

1.4 Notes

An overview of the history and content of mathematics as a discipline can be found in Aleksandrov et al. (1969) and Campbell and Higgins (1984). Blåsjö (2005) illuminates the background of the isoperimetric problem. The historical development of the calculus of variations is summarized by Goldstine (1980) and Hildebrandt and Tromba (1985). For a history of technological feedback control systems, see Mayr (1970).

---

15. Vinter (2000) provided an account of optimal control theory in the setting of nonsmooth analysis.