Thomas A. Weber

Optimal Control Theory with Applications in Economics

Optimal Control Theory with Applications in Economics

Thomas A. Weber

Foreword by A. V. Kryazhimskiy

The MIT Press
Cambridge, Massachusetts
London, England

For information about special quantity discounts, please email
special_sales@mitpress.mit.edu.

This book was set in Palatino by Westchester Book Composition. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Weber, Thomas A., 1969–
Optimal control theory with applications in economics / Thomas A. Weber; foreword by A. V. Kryazhimskiy.
p. cm.
Includes bibliographical references and index.
ISBN 978-0-262-01573-8 (hardcover : alk. paper)
1. Economics—Mathematical models. 2. Control theory. 3. Mathematical optimization. 4. Game theory. I. Title.
HB135.W433 2011
330.01′515642—dc22 2010046482

10 9 8 7 6 5 4 3 2 1

Foreword

Since the discovery, by L. S. Pontryagin, of the necessary optimality conditions for the control of dynamic systems in the 1950s, mathematical control theory has found numerous applications in engineering and in the social sciences. T. A. Weber has dedicated his book to optimal control theory and its applications in economics. Readers can find here a succinct introduction to the basic control-theoretic methods, and also clear and meaningful examples illustrating the theory.

Remarkable features of this text are rigor, scope, and brevity, combined with a well-structured hierarchical approach. The author starts with a general view on dynamical systems from the perspective of the theory of ordinary differential equations; on this basis, he proceeds to the classical optimal control theory, and he concludes the book with more recent views of game theory and mechanism design, in which optimal control plays an instrumental role.

The treatment is largely self-contained and compact; it amounts to a lucid overview, featuring much of the author's own research. The character of the problems discussed in the book promises to make the theory accessible to a wide audience. The exercises placed at the chapter endings are largely original.

I am confident that readers will appreciate the author's style and students will find this book a helpful guide on their path of discovery.

A. V. Kryazhimskiy
Steklov Institute of Mathematics, Russian Academy of Sciences
International Institute for Applied Systems Analysis

Acknowledgments

This book is based on my graduate course on Applied Optimal Control Theory taught both at Moscow State University and Stanford University. The development of this course was made possible through funding from Moscow State University (MSU), the Steklov Mathematical Institute in Moscow, and Stanford University. In particular, a course development grant from James Plummer and Channing Robertson, dean and former vice dean of the School of Engineering at Stanford University, allowed Elena Rovenskaya to spend time at Stanford to work on problem sets and solutions. The numerous discussions with her were invaluable, and I am very grateful for her contributions. Nikolai Grigorenko, the deputy head of the Optimal Control Department at MSU, was instrumental in making this possible.

I am very grateful to Evgenii Moiseev, dean of the Faculty of Computational Mathematics and Cybernetics at Lomonosov Moscow State University, for his 2007 invitation to deliver a summer course on dynamic optimization with applications in economics. My deepest gratitude also goes to Arkady Kryazhimskiy and Sergey Aseev for their encouragement to write this book and their continuous support and friendship. They have fueled my interest in optimal control since 2001, when I was fortunate enough to meet them while participating in the Young Scientists Summer Program at the International Institute for Applied Systems Analysis in Laxenburg, Austria. They also invited me to the 2008 International Conference on Differential Equations and Topology in Moscow, dedicated to the centennial anniversary of Lev Pontryagin, where the discussions about the summer course continued.

I am indebted to the students in Moscow and Stanford for the many after-class discussions where they often taught me, perhaps unknowingly, just as much as I taught them. Kenneth Gillingham provided a set of impeccable handwritten notes from the course, which helped organize the thoughts for the book. Naveed Chehrazi was an excellent course assistant who contributed numerous insightful suggestions. Elena Rovenskaya and Denis Pivovarchuk organized the practice sessions in Moscow. Stefan Behringer, Andrei Dmitruk, and three MIT Press reviewers provided useful feedback on earlier versions of the book. Markus Edvall from TOMLAB was helpful in debugging several numerical algorithms. I would like to thank Alice Cheyer for her detailed copyediting, as well as Alexey Smirnov and my assistant, Marilynn Rose, for their help with editing earlier versions of the manuscript. Jane Macdonald at the MIT Press believed in this project from the moment I told her about it at the 2009 European Meeting of the Econometric Society in Barcelona. I am very grateful for her helpful advice, encouragement, and great support throughout the publishing process.

I should also like to acknowledge my indebtedness to the great teachers in control theory and economics, whom I encountered at MIT, the University of Pennsylvania, and Stanford University, in particular Kenneth Arrow, Dimitri Bertsekas, David Cass, Richard Kihlstrom, Alexandre Kirillov, Steve Matthews, and Ilya Segal. Richard Vinter introduced me to control theory in 1994 as a wonderful research advisor at Imperial College London. In my thinking and intellectual approach I also owe very much to my advisor at MIT, Alexandre Megretski, who did not believe in books and whose genius and critical mind I admire. Paul Kleindorfer, my dissertation advisor at Wharton, taught me so much, including the fact that broad interests can be a real asset. I am very grateful for his friendship and constant support.

My special thanks go to Sergey Aseev, Eric Clemons, David Luenberger, James Sweeney, and Andrew Whinston for their friendship, as well as Ann and Wim for their sacrifice, endless love, and understanding.

Stanford, California
May 2011

Contents

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

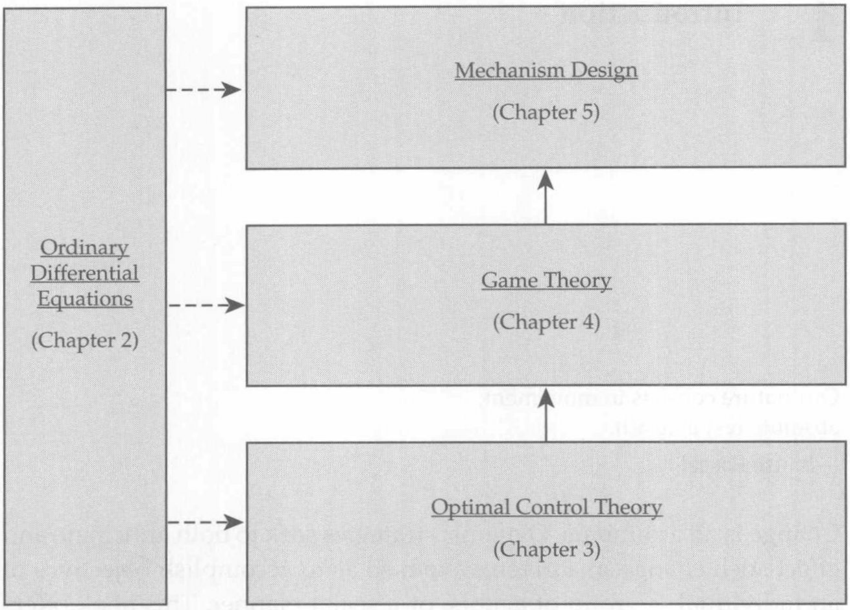

Figure 1.1
Topics covered in this book.

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 $\int_{t_0}^{T} C(t,x(t),u(t))\,dt \le B$. It can be rewritten as $\dot{y}(t) = C(t,x(t),u(t))$, $y(t_0) = 0$, $y(T) \le B$. The latter formulation falls squarely within the general optimal-control formalism developed in this book, so isoperimetric constraints do not need special consideration.

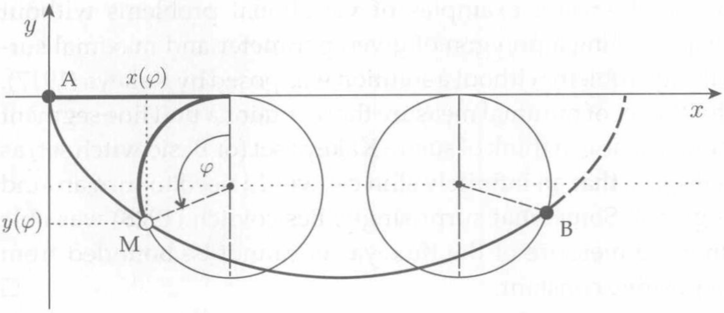

Figure 1.2
Brachistochrone connecting the points A and B in parametric form: $(x(\varphi), y(\varphi)) = (\alpha(\varphi - \sin(\varphi)), \alpha(\cos(\varphi) - 1))$, where $\varphi = \varphi(t) = \sqrt{g/\alpha}\, t$, and $g \approx 9.81$ meters per second squared is the gravitational constant. The 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).