

Dynamic Programming Models And Applications Eric V Denardo

Mathematical programming: an overview; solving linear programs; sensitivity analysis; duality in linear programming; mathematical programming in practice; integration of strategic and tactical planning in the aluminum industry; planning the mission and composition of the U.S. merchant Marine fleet; network models; integer programming; design of a naval tender job shop; dynamic programming; large-scale systems; nonlinear programming; a system for bank portfolio planning; vectors and matrices; linear programming in matrix form; a labeling algorithm for the maximum-flow network problem.

Optimization is the act of obtaining the "best" result under given circumstances. In design, construction, and maintenance of any engineering system, engineers must make technological and managerial decisions to minimize either the effort or cost required or to maximize benefits. There is no single method available for solving all optimization problems efficiently. Several optimization methods have been developed for different types of problems. The optimum-seeking methods are mathematical programming techniques (specifically, nonlinear programming techniques). Nonlinear Optimization: Models and Applications presents the concepts in several ways to foster understanding. Geometric interpretation: is used to re-enforce the concepts and to foster understanding of the mathematical procedures. The student sees that many problems can be analyzed, and approximate solutions found before analytical solutions techniques are applied. Numerical approximations: early on, the student is exposed to numerical techniques. These numerical procedures are algorithmic and iterative. Worksheets are provided in Excel, MATLAB®, and Maple™ to facilitate the procedure. Algorithms: all algorithms are provided with a step-by-step format. Examples follow the summary to illustrate its use and application. Nonlinear Optimization: Models and Applications: Emphasizes process and interpretation throughout Presents a general classification of optimization problems Addresses situations that lead to models illustrating many types of optimization problems Emphasizes model formulations Addresses a special class of problems that can be solved using only elementary calculus Emphasizes model solution and model sensitivity analysis About the author: William P. Fox is an emeritus professor in the Department of Defense Analysis at the Naval Postgraduate School. He received his Ph.D. at Clemson University and has taught at the United States Military Academy and at Francis Marion University where he was the chair of mathematics. He has written many publications, including over 20 books and over 150 journal articles. Currently, he is an adjunct professor in the Department of Mathematics at the College of William and Mary. He is the emeritus director of both the High School Mathematical Contest in Modeling and the Mathematical Contest in Modeling.

Features step-by-step examples based on actual data and connects fundamental mathematical modeling skills and decision making concepts to everyday applicability Featuring key linear programming, matrix, and probability concepts, "Finite Mathematics: Models and Applications" emphasizes cross-disciplinary applications that relate mathematics to everyday life. The book provides a unique combination of practical mathematical applications to illustrate the wide use of mathematics in fields ranging from business, economics, finance, management, operations research, and the life and social sciences. In order to emphasize the main concepts of each chapter, "Finite Mathematics: Models and Applications" features plentiful pedagogical elements throughout such as special exercises, end notes, hints, select solutions, biographies of key mathematicians, boxed key principles, a glossary of important terms and topics, and an overview

of use of technology. The book encourages the modeling of linear programs and their solutions and uses common computer software programs such as LINDO. In addition to extensive chapters on probability and statistics, principles and applications of matrices are included as well as topics for enrichment such as the Monte Carlo method, game theory, kinship matrices, and dynamic programming. Supplemented with online instructional support materials, the book features coverage including: Algebra Skills Mathematics of Finance Matrix Algebra Geometric Solutions Simplex Methods Application Models Set and Probability Relationships Random Variables and Probability Distributions Markov Chains Mathematical Statistics Enrichment in Finite Mathematics An ideal textbook, " Finite Mathematics: Models and Applications "is intended for students in fields from entrepreneurial and economic to environmental and social science, including many in the arts and humanities. Carla C. Morris, PhD, is Assistant Professor of Mathematics in the Associate in Arts Program at the University of Delaware. A member of The Institute for Operations Research and the Management Sciences and the Mathematical Association of America, Dr. Morris teaches courses ranging from college algebra to calculus and statistics. Robert M. Stark, PhD, is Professor Emeritus in the Departments of Mathematical Sciences and Civil and Environmental Engineering at the University of Delaware. Dr. Stark's teaching and research interests include applied probability, mathematical optimization, operations research, and mathematics education.

In this book, we study theoretical and practical aspects of computing methods for mathematical modelling of nonlinear systems. A number of computing techniques are considered, such as methods of operator approximation with any given accuracy; operator interpolation techniques including a non-Lagrange interpolation; methods of system representation subject to constraints associated with concepts of causality, memory and stationarity; methods of system representation with an accuracy that is the best within a given class of models; methods of covariance matrix estimation; methods for low-rank matrix approximations; hybrid methods based on a combination of iterative procedures and best operator approximation; and methods for information compression and filtering under condition that a filter model should satisfy restrictions associated with causality and different types of memory. As a result, the book represents a blend of new methods in general computational analysis, and specific, but also generic, techniques for study of systems theory and its particular branches, such as optimal filtering and information compression. - Best operator approximation, - Non-Lagrange interpolation, - Generic Karhunen-Loeve transform - Generalised low-rank matrix approximation - Optimal data compression - Optimal nonlinear filtering

Algorithms and Applications

Applied Mathematical Programming

Lecture Notes of the SVOR/ASRO Tutorial Thun, Switzerland, October 14–16, 1992

Dynamic Programming and Its Applications

Grey Programming and Its Applications to Water Resources Management

Nonlinear and Dynamic Programming

Dynamic programming is a powerful method for solving optimization problems, but has a number of drawbacks that limit its use to solving problems of very low dimension. To overcome these limitations, author Rein Luus suggested using it in an iterative fashion. Although this method required vast computer resources, modifications to his original schem

Graduate-level text provides complete and rigorous expositions of economic models analyzed primarily from the point of view of their mathematical properties, followed by relevant mathematical reviews. Part I covers optimizing

theory; Parts II and III survey static and dynamic economic models; and Part IV contains the mathematical reviews, which range from linear algebra to point-to-set mappings.

Permanently increasing requirements in power supply necessitate efficient control of electric power systems. An emerging subject of importance is optimization. Papers on modelling aspects of unit commitment and optimal power flow provide the introduction to power systems control and to its associated problem statement. Due to the nature of the underlying optimization problems recent developments in advanced and well established mathematical programming methodologies are presented, illustrating in which way dynamic, separable, continuous and stochastic features might be exploited. In completing the various methodologies a number of presentations have stated experiences with optimization packages currently used for unit commitment and optimal power flow calculations. This work represents a state-of-the-art of mathematical programming methodologies, unit commitment, optimal power flow and their applications in power system control.

Integer Programming: Theory, Applications, and Computations provides information pertinent to the theory, applications, and computations of integer programming. This book presents the computational advantages of the various techniques of integer programming. Organized into eight chapters, this book begins with an overview of the general categorization of integer applications and explains the three fundamental techniques of integer programming. This text then explores the concept of implicit enumeration, which is general in a sense that it is applicable to any well-defined binary program. Other chapters consider the branch-and-bound methods, the cutting-plane method, and its closely related asymptotic problem. This book discusses as well several specialized algorithms for certain well-known integer models and provides an alternative approach to the solution of the integer problem. The final chapter deals with a number of observations about the formulations and executions of integer programming models. This book is a valuable resource for industrial engineers and research workers.

Dynamic Programming and Its Application to Optimal Control

Applications to Agriculture and Natural Resources

Proceedings of the International Conference on Dynamic Programming and Its Applications, University of British Columbia, Vancouver, British Columbia, Canada, April 14-16, 1977

Forward-Looking Decision Making

Operations Research

A complete resource to Approximate Dynamic Programming (ADP), including on-line simulation code Provides a tutorial that readers can use to start implementing the learning algorithms provided in the book Includes ideas, directions, and recent results on current research issues and addresses applications where ADP has been successfully implemented The contributors are leading researchers in the field

The new edition of an introductory text that teaches students the art of computational problem solving, covering topics ranging from simple algorithms to information visualization. This book introduces students with little or no prior programming experience to the art of computational

problem solving using Python and various Python libraries, including PyLab. It provides students with skills that will enable them to make productive use of computational techniques, including some of the tools and techniques of data science for using computation to model and interpret data. The book is based on an MIT course (which became the most popular course offered through MIT's OpenCourseWare) and was developed for use not only in a conventional classroom but in a massive open online course (MOOC). This new edition has been updated for Python 3, reorganized to make it easier to use for courses that cover only a subset of the material, and offers additional material including five new chapters. Students are introduced to Python and the basics of programming in the context of such computational concepts and techniques as exhaustive enumeration, bisection search, and efficient approximation algorithms. Although it covers such traditional topics as computational complexity and simple algorithms, the book focuses on a wide range of topics not found in most introductory texts, including information visualization, simulations to model randomness, computational techniques to understand data, and statistical techniques that inform (and misinform) as well as two related but relatively advanced topics: optimization problems and dynamic programming. This edition offers expanded material on statistics and machine learning and new chapters on Frequentist and Bayesian statistics.

Humans interact with and are part of the mysterious processes of nature. Inevitably they have to discover how to manage the environment for their long-term survival and benefit. To do this successfully means learning something about the dynamics of natural processes, and then using the knowledge to work with the forces of nature for some desired outcome. These are intriguing and challenging tasks. This book describes a technique which has much to offer in attempting to achieve the latter task. A knowledge of dynamic programming is useful for anyone interested in the optimal management of agricultural and natural resources for two reasons. First, resource management problems are often problems of dynamic optimization. The dynamic programming approach offers insights into the economics of dynamic optimization which can be explained much more simply than can other approaches. Conditions for the optimal management of a resource can be derived using the logic of dynamic programming, taking as a starting point the usual economic definition of the value of a resource which is optimally managed through time. This is set out in Chapter I for a general resource problem with the minimum of mathematics. The results are related to the discrete maximum principle of control theory. In subsequent chapters dynamic programming arguments are used to derive optimality conditions for particular resources.

A collection of articles which provide examples that demonstrate the application of dynamic programming to a wide variety of decision problems in agriculture.

Applications Of Dynamic Programming To Agricultural Decision Problems
Finite Mathematics

Foundations and Principles, Second Edition
Handbook of Learning and Approximate Dynamic Programming
Dynamic Economics
With Application to Understanding Data

This book provides a complete and comprehensive reference/guide to Pyomo (Python Optimization Modeling Objects) for both beginning and advanced modelers, including students at the undergraduate and graduate levels, academic researchers, and practitioners. The text illustrates the breadth of the modeling and analysis capabilities that are supported by the software and support of complex real-world applications. Pyomo is an open source software package for formulating and solving large-scale optimization and operations research problems. The text begins with a tutorial on simple linear and integer programming models. A detailed reference of Pyomo's modeling components is illustrated with extensive examples, including a discussion of how to load data from data sources like spreadsheets and databases. Chapters describing advanced modeling capabilities for nonlinear and stochastic optimization are also included. The Pyomo software provides familiar modeling features within Python, a powerful dynamic programming language that has a very clear, readable syntax and intuitive object orientation. Pyomo includes Python classes for defining sparse sets, parameters, and variables, which can be used to formulate algebraic expressions that define objectives and constraints. Moreover, Pyomo can be used from a command-line interface and within Python's interactive command environment, which makes it easy to create Pyomo models, apply a variety of optimizers, and examine solutions. The software supports a different modeling approach than commercial AML (Algebraic Modeling Languages) tools, and is designed for flexibility, extensibility, portability, and maintainability but also maintains the central ideas in modern AMLs. "Foundations and trends in stochastic systems 1:1 (2006)"

This comprehensive study of dynamic programming applied to numerical solution of optimization problems. It will interest aerodynamic, control, and industrial engineers, numerical analysts, and computer specialists, applied mathematicians, economists, and operations and systems analysts. Originally published in 1962. The Princeton Legacy Library uses the latest print-on-demand technology to again make available previously out-of-print books from the distinguished backlist of Princeton University Press. These editions preserve the original texts of these important books while presenting them in durable paperback and hardcover editions. The goal of the Princeton Legacy Library is to vastly increase access to the rich scholarly heritage found in the thousands of books published by Princeton University Press since its founding in 1905.

Economic Modeling and Inference takes econometrics to a new level by demonstrating how to combine modern economic theory with the latest statistical inference methods to get the most out of economic data. This graduate-level textbook draws applications from both microeconomics and macroeconomics, paying special attention to financial and labor economics, with an emphasis throughout on what

observations can tell us about stochastic dynamic models of rational optimizing behavior and equilibrium. Bent Jesper Christensen and Nicholas Kiefer show how parameters often thought estimable in applications are not identified even in simple dynamic programming models, and they investigate the roles of extensions, including measurement error, imperfect control, and random utility shocks for inference. When all implications of optimization and equilibrium are imposed in the empirical procedures, the resulting estimation problems are often nonstandard, with the estimators exhibiting nonregular asymptotic behavior such as short-ranked covariance, superconsistency, and non-Gaussianity. Christensen and Kiefer explore these properties in detail, covering areas including job search models of the labor market, asset pricing, option pricing, marketing, and retirement planning. Ideal for researchers and practitioners as well as students, *Economic Modeling and Inference* uses real-world data to illustrate how to derive the best results using a combination of theory and cutting-edge econometric techniques. Covers identification and estimation of dynamic programming models Treats sources of error--measurement error, random utility, and imperfect control Features financial applications including asset pricing, option pricing, and optimal hedging Describes labor applications including job search, equilibrium search, and retirement Illustrates the wide applicability of the approach using micro, macro, and marketing examples

An Introduction

Monotonicity in Markov Reward and Decision Chains

Pyomo - Optimization Modeling in Python

Approximate Dynamic Programming

Lectures on Stochastic Programming

Stochastic Dynamic Programming and the Control of Queueing Systems

This overview provides a single-volume treatment of key algorithms and theories. Begins with the derivation of optimality conditions and discussions of convex programming, duality, generalized convexity, and analysis of selected nonlinear programs, and then explores techniques for numerical solutions and unconstrained optimization methods. 1976 edition. Includes 58 figures and 7 tables.

Dynamic Programming Models and Applications Courier Corporation

*Incorporating a number of the author's recent ideas and examples, *Dynamic Programming: Foundations and Principles, Second Edition* presents a comprehensive and rigorous treatment of dynamic programming. The author emphasizes the crucial role that modeling plays in understanding this area. He also shows how Dijkstra's algorithm is an excellent example of a dynamic programming algorithm, despite the impression given by the computer science literature. New to the Second Edition Expanded discussions of sequential decision models and the role of the state variable in modeling A new chapter on forward dynamic programming models A new chapter on the Push method that gives a dynamic programming perspective on Dijkstra's algorithm for the shortest path problem A new appendix on the Corridor method Taking into account recent developments in dynamic programming, this edition continues to provide a systematic, formal outline of Bellman's approach to dynamic programming. It looks at dynamic programming as a problem-solving methodology, identifying its constituent components and explaining its theoretical basis for tackling problems.*

Optimization problems involving stochastic models occur in almost all areas of science and engineering, such as telecommunications, medicine, and finance. Their existence compels a need for rigorous ways of formulating, analyzing, and solving such problems. This book focuses on optimization problems involving uncertain parameters and covers the theoretical foundations and recent advances in areas where stochastic

models are available. Readers will find coverage of the basic concepts of modeling these problems, including recourse actions and the nonanticipativity principle. The book also includes the theory of two-stage and multistage stochastic programming problems; the current state of the theory on chance (probabilistic) constraints, including the structure of the problems, optimality theory, and duality; and statistical inference in and risk-averse approaches to stochastic programming.

MERCHOP : a Dynamic Programming Model for Estimating the Harvest Value of Unthinned Loblolly and Slash Pine Plantations

Economic Modeling and Inference

Reinforcement Learning and Optimal Control

Introduction to Stochastic Dynamic Programming

Solving the Curses of Dimensionality

Applicability and Limits

Introduction to sequential decision processes covers use of dynamic programming in studying models of resource allocation, methods for approximating solutions of control problems in continuous time, production control, more. 1982 edition.

This book considers large and challenging multistage decision problems, which can be solved in principle by dynamic programming (DP), but their exact solution is computationally intractable. We discuss solution methods that rely on approximations to produce suboptimal policies with adequate performance. These methods are collectively known by several essentially equivalent names: reinforcement learning, approximate dynamic programming, neuro-dynamic programming. They have been at the forefront of research for the last 25 years, and they underlie, among others, the recent impressive successes of self-learning in the context of games such as chess and Go. Our subject has benefited greatly from the interplay of ideas from optimal control and from artificial intelligence, as it relates to reinforcement learning and simulation-based neural network methods. One of the aims of the book is to explore the common boundary between these two fields and to form a bridge that is accessible by workers with background in either field. Another aim is to organize coherently the broad mosaic of methods that have proved successful in practice while having a solid theoretical and/or logical foundation. This may help researchers and practitioners to find their way through the maze of competing ideas that constitute the current state of the art. This book relates to several of our other books: Neuro-Dynamic Programming (Athena Scientific, 1996), Dynamic Programming and Optimal Control (4th edition, Athena Scientific, 2017), Abstract Dynamic Programming (2nd edition, Athena Scientific, 2018), and Nonlinear Programming (Athena Scientific, 2016). However, the mathematical style of this book is somewhat different. While we provide a rigorous, albeit short, mathematical account of the theory of finite and infinite horizon dynamic programming, and some fundamental approximation methods, we rely more on intuitive explanations and less on proof-based insights. Moreover, our mathematical requirements are quite modest: calculus, a minimal use of matrix-vector algebra, and elementary probability (mathematically complicated arguments involving laws of large numbers and stochastic convergence are bypassed in favor of intuitive explanations). The book illustrates the methodology with many examples and illustrations, and uses a gradual expository approach, which proceeds along four directions: (a) From exact DP to approximate DP: We first discuss exact DP algorithms, explain why they may be difficult to implement, and then use them as the basis for approximations. (b) From finite horizon to infinite horizon problems: We first discuss finite horizon exact and approximate DP methodologies, which are intuitive and mathematically simple, and then progress to infinite horizon problems. (c) From deterministic to stochastic models: We often discuss separately deterministic and stochastic problems, since deterministic problems are simpler and offer special advantages for some of our methods. (d) From model-based to model-free implementations: We first discuss model-based implementations, and then we identify schemes that can be appropriately modified to

work with a simulator. The book is related and supplemented by the companion research monograph Rollout, Policy Iteration, and Distributed Reinforcement Learning (Athena Scientific, 2020), which focuses more closely on several topics related to rollout, approximate policy iteration, multiagent problems, discrete and Bayesian optimization, and distributed computation, which are either discussed in less detail or not covered at all in the present book. The author's website contains class notes, and a series of videolectures and slides from a 2021 course at ASU, which address a selection of topics from both books.

A path-breaking account of Markov decision processes-theory and computation This book's clear presentation of theory, numerous chapter-end problems, and development of a unified method for the computation of optimal policies in both discrete and continuous time make it an excellent course text for graduate students and advanced undergraduates. Its comprehensive coverage of important recent advances in stochastic dynamic programming makes it a valuable working resource for operations research professionals, management scientists, engineers, and others. Stochastic Dynamic Programming and the Control of Queueing Systems presents the theory of optimization under the finite horizon, infinite horizon discounted, and average cost criteria. It then shows how optimal rules of operation (policies) for each criterion may be numerically determined. A great wealth of examples from the application area of the control of queueing systems is presented. Nine numerical programs for the computation of optimal policies are fully explicated. The Pascal source code for the programs is available for viewing and downloading on the Wiley Web site at www.wiley.com/products/subject/mathematics. The site contains a link to the author's own Web site and is also a place where readers may discuss developments on the programs or other aspects of the material. The source files are also available via ftp at <ftp://ftp.wiley.com/public/scitechmed/stochastic> Stochastic Dynamic Programming and the Control of Queueing Systems features: * Path-breaking advances in Markov decision process techniques, brought together for the first time in book form * A theorem/proof format (proofs may be omitted without loss of continuity) * Development of a unified method for the computation of optimal rules of system operation * Numerous examples drawn mainly from the control of queueing systems * Detailed discussions of nine numerical programs * Helpful chapter-end problems * Appendices with complete treatment of background material Introduction to Stochastic Dynamic Programming presents the basic theory and examines the scope of applications of stochastic dynamic programming. The book begins with a chapter on various finite-stage models, illustrating the wide range of applications of stochastic dynamic programming. Subsequent chapters study infinite-stage models: discounting future returns, minimizing nonnegative costs, maximizing nonnegative returns, and maximizing the long-run average return. Each of these chapters first considers whether an optimal policy need exist—providing counterexamples where appropriate—and then presents methods for obtaining such policies when they do. In addition, general areas of application are presented. The final two chapters are concerned with more specialized models. These include stochastic scheduling models and a type of process known as a multiproject bandit. The mathematical prerequisites for this text are relatively few. No prior knowledge of dynamic programming is assumed and only a moderate familiarity with probability—including the use of conditional expectation—is necessary.

Optimization in Planning and Operation of Electric Power Systems

Algorithms for Handling CVaR-constraints in Dynamic Stochastic Programming Models with Applications to Finance

Modeling and Theory

Models and Applications

Theory, Applications, and Computations

A Computational Tool

Dynamic Programming and Its Applications provides information pertinent to the theory and application of dynamic programming. This book presents the development and future directions for dynamic programming. Organized into four parts encompassing 23 chapters, this book begins with an overview of recurrence conditions for countable state Markov decision problems, which ensure that the optimal average reward exists and satisfies the functional equation of dynamic programming. This text then provides an extensive analysis of the theory of successive approximation for Markov decision problems. Other chapters consider the computational methods for deterministic, finite horizon problems, and present a unified and insightful presentation of several foundational questions. This book discusses as well the relationship between policy iteration and Newton's method. The final chapter deals with the main factors severely limiting the application of dynamic programming in practice. This book is a valuable resource for growth theorists, economists, biologists, mathematicians, and applied management scientists.

Dynamic programming is a method of solving multi-stage problems in which decisions at one stage become the conditions governing the succeeding stages. It can be applied to the management of water reservoirs, allowing them to be operated more efficiently. Originally published in 2007, this is one of the few books dedicated solely to dynamic programming techniques used in reservoir management. It presents the applicability of these techniques and their limits on the operational analysis of reservoir systems. The dynamic programming models presented in this book have been applied to reservoir systems all over the world, helping the reader to appreciate the applicability and limits of these models. The book also includes a model for the operation of a reservoir during an emergency situation. This volume will be a valuable reference to researchers in hydrology, water resources and engineering, as well as professionals in reservoir management.

This is the 3rd edition of a research monograph providing a synthesis of old research on the foundations of dynamic programming (DP), with the modern theory of approximate DP and new research on semicontractive models. It aims at a unified and economical development of the core theory and algorithms of total cost sequential decision problems, based

on the strong connections of the subject with fixed point theory. The analysis focuses on the abstract mapping that underlies DP and defines the mathematical character of the associated problem. The discussion centers on two fundamental properties that this mapping may have: monotonicity and (weighted sup-norm) contraction. It turns out that the nature of the analytical and algorithmic DP theory is determined primarily by the presence or absence of these two properties, and the rest of the problem's structure is largely inconsequential. New research is focused on two areas: 1) The ramifications of these properties in the context of algorithms for approximate DP, and 2) The new class of semicontractive models, exemplified by stochastic shortest path problems, where some but not all policies are contractive. The 3rd edition is very similar to the 2nd edition, except for the addition of a new chapter (Chapter 5), which deals with abstract DP models for sequential minimax problems and zero-sum games. The book is an excellent supplement to several of our books: *Neuro-Dynamic Programming* (Athena Scientific, 1996), *Dynamic Programming and Optimal Control* (Athena Scientific, 2017), *Reinforcement Learning and Optimal Control* (Athena Scientific, 2019), and *Rollout, Policy Iteration, and Distributed Reinforcement Learning* (Athena Scientific, 2020).

Individuals and families make key decisions that impact many aspects of financial stability and determine the future of the economy. These decisions involve balancing current sacrifice against future benefits. People have to decide how much to invest in health care, exercise, their diet, and insurance. They must decide how much debt to take on, and how much to save. And they make choices about jobs that determine employment and unemployment levels. Forward-Looking Decision Making is about modeling this individual or family-based decision making using an optimizing dynamic programming model. Robert Hall first reviews ideas about dynamic programs and introduces new ideas about numerical solutions and the representation of solved models as Markov processes. He surveys recent research on the parameters of preferences--the intertemporal elasticity of substitution, the Frisch elasticity of labor supply, and the Frisch cross-elasticity. He then examines dynamic programming models applied to health spending, long-term care insurance,

employment, entrepreneurial risk-taking, and consumer debt. Linking theory with data and applying them to real-world problems, Forward-Looking Decision Making uses dynamic optimization programming models to shed light on individual behaviors and their economic implications.

Theory and Applications

3rd Edition

Dynamic Programming Under Parametric Uncertainty with Applications in Cyber Security and Project Management

Dynamic Programming

Scalable Approximate Dynamic Programming Models with Applications in Air Transportation

Introduction to Computation and Programming Using Python, second edition

"Christensen and Kiefer's excellent book shows how careful dynamic theory and econometrics go hand in hand, opening up new vistas in the areas of search theory, finance, and macroeconomics."--Tom Sargent, New York University and the Hoover Institution

"There is no other book that mixes dynamic economic theory, statistical inference, and real quantitative applications like this one. Christensen and Kiefer will challenge the top tier of students and take them to the research frontier."--Robert Lucas, University of Chicago

"Dynamic programming is an organizing framework that has enabled economists to integrate economic theory with empirical analysis. Few textbooks reflect the integrated nature of contemporary research, but Christensen and Kiefer reveal the power of the dynamic programming approach in a wide variety of applications from job search to portfolio choice. Their new book will be invaluable to students who wish to participate in this exciting enterprise."--John Y. Campbell, Harvard University

"The authors do a splendid job of showing how to use stochastic dynamic optimization techniques to generate the implied distributions of observables needed for estimation. There are many interesting and useful examples included in the book, ranging from applications of the theory of job search to those of asset pricing theory. This book should be a reference for anyone interested in using dynamic economic models to make inferences about the world we observe."--Dale Mortensen, Aarhus University, Denmark, and Northwestern University

"An extremely ambitious and thought-provoking book, one that combines state-of-the-art economic theory

with sophisticated econometric techniques. The dynamic programming framework brings together important results and recent developments in a unique, unified way. The book is sure to inspire many PhD students and empirically oriented researchers for years to come."--Tim Bollerslev, Duke University "I have been looking for a book like this for quite a while. *Economic Modeling and Inference* is written for those who want to do applied work and actually apply this to real-life data or run simulations. This much-needed book fills a void. It is certainly a significant contribution to the field."--Yaw Nyarko, New York University "Economic Modeling and Inference blends economic theory and statistical inference in a seamless fashion. Every dynamic decision model is discussed with an eye for it to be fit with economic data. Every econometric inference tool is developed for the purpose of testing economic decision models. This book is long overdue. It will influence and benefit young economists for generations to come."--Mark Y. An, Fannie Mae

An integrated approach to the empirical application of dynamic optimization programming models, for students and researchers. This book is an effective, concise text for students and researchers that combines the tools of dynamic programming with numerical techniques and simulation-based econometric methods. Doing so, it bridges the traditional gap between theoretical and empirical research and offers an integrated framework for studying applied problems in macroeconomics and microeconomics. In part I the authors first review the formal theory of dynamic optimization; they then present the numerical tools and econometric techniques necessary to evaluate the theoretical models. In language accessible to a reader with a limited background in econometrics, they explain most of the methods used in applied dynamic research today, from the estimation of probability in a coin flip to a complicated nonlinear stochastic structural model. These econometric techniques provide the final link between the dynamic programming problem and data. Part II is devoted to the application of dynamic programming to specific areas of applied economics, including the study of business cycles, consumption, and investment behavior. In each instance the authors present the specific optimization problem as a dynamic programming problem, characterize the optimal policy functions, estimate

the parameters, and use models for policy evaluation. The original contribution of *Dynamic Economics: Quantitative Methods and Applications* lies in the integrated approach to the empirical application of dynamic optimization programming models. This integration shows that empirical applications actually complement the underlying theory of optimization, while dynamic programming problems provide needed structure for estimation and policy evaluation. This book is intended to provide an introductory text of *Nonlinear and Dynamic Programming* for students of managerial economics and operations research. The author also hopes that engineers, business executives, managers, and others responsible for planning of industrial operations may find it useful as a guide to the problems and methods treated, with a view to practical applications. The book may be considered as a sequel to the author's *Linear Programming in Industry* (1960, 4th revised and enlarged edition 1974), but it can be used independently by readers familiar with the elements of linear programming models and techniques. The two volumes constitute an introduction to the methods of mathematical programming and their application to industrial optimization problems. The author feels that the vast and ever-increasing literature on mathematical programming has not rendered an introductory exposition superfluous. The general student often tends to feel somewhat lost if he goes straight to the special literature; he will be better equipped for tackling real problems and using computer systems if he has acquired some previous training in constructing small-scale programming models and applying standard algorithms for solving them by hand. The book is intended to provide this kind of training, keeping the mathematics at the necessary minimum. The text contains numerous exercises. The reader should work out these problems for himself and check with the answers given at the end of the book. The text is based on lectures given at the University of Copenhagen.

This book provides a practical introduction to computationally solving discrete optimization problems using dynamic programming. From the examples presented, readers should more easily be able to formulate dynamic programming solutions to their own problems of interest. We also provide and describe the design, implementation, and use of a software tool that has been used to numerically solve all of

the problems presented earlier in the book.

Mathematical Economics

Dynamic Programming Models Applied to Health, Risk, Employment, and Financial Stability

Integer Programming

Analysis and Methods

Applied Dynamic Programming

Iterative Dynamic Programming

Introduction to mathematical theory of multistage decision processes takes a "functional equation" approach. Topics include existence and uniqueness theorems, optimal inventory equation, bottleneck problems, multistage games, Markovian decision processes, and more. 1957 edition.

The trustworthiness of models and optimization is limited because the associated systems might be changing and data about them can be limited, i.e., there is "parametric" uncertainty. This dissertation provides applications and theory related to mitigating the effects of changing systems and data limitations in optimal decision-making. The primary application considered relates to reducing the maintenance costs associated with cyber security. By selecting optimal policies addressing data limitations, losses from stolen information and maintenance costs can be balanced. The approximated expected savings from implementing the suggested policies at a large Midwestern organization is over \$14M with a discount factor of 0.95 monthly. The dissertation also integrates data and dynamic programming models for project management decision-making that accounts for coordination and planning costs. This facilitates more accurate schedules with significant cost savings. Insights are provided into the choice between traditional planning methods and agile project management methods that reduce planning complexity. In many situations, we find that the so-called optimal approaches are suboptimal because they fail to address sizable coordination and planning costs. Two types of parametric uncertainty are explored here, each of which results in fundamentally different formulations and solution schemes. The first type of uncertainty considered relates to system parameters fluctuating over time randomly. The related models differ from ordinary inhomogeneous approaches because the specific parameters are not known and are assumed to fluctuate with known distributions. Associated decision problems are referred to as "Markov decision processes with random inhomogeneity" and proposed optimal solutions methods. Proof is given that the solution produced by backward induction is optimal for the finite horizon problems, and that the value-iteration-based algorithm gives solutions converging to the infinite horizon solutions, together with results regarding monotonicity property and rate of the convergence. The second type of parametric uncertainty is caused by insufficient data for parameter estimation, i.e., "data-driven" uncertainty. Previous researchers studying data-driven Markov decision processes declare the problem is intractable. Therefore, they propose approximation methods. We prove

that their methods can approximate suboptimal solutions by a numerical example. We also provide a dynamic programming algorithm to generate data-driven optimal policies with learning. We do this by demonstrating that the problem is equivalent to partially observable Markov decision processes. Further, by exploiting the structure of the problem and bounds assuming perfect information, we develop a bounding heuristic method for the infinite horizon problems.

Dynamic Programming Based Operation of Reservoirs

Nonlinear Programming

Nonlinear Optimization

Quantitative Methods and Applications

Abstract Dynamic Programming