Solving Stochastic Dynamic Programming Problems A Mixed

This book offers a systematic introduction to the optimal stochastic control theory via the dynamic programming principle, which is a powerful tool to analyze control problems. First we consider completely observable control problems with finite horizons. Using a time discretization we construct a nonlinear semigroup related to the dynamic programming principle (DPP), whose generator provides the Hamilton–Jacobi–Bellman (HJB) equation, and we characterize the value function via the nonlinear semigroup, besides the viscosity solution theory. When we control not only the dynamics of a system but also the terminal time of its evolution, control-stopping problems arise. This problem is treated in the same frameworks, via the nonlinear semigroup. Its results are applicable to the American option price problem. Zero-sum two-player time-homogeneous stochastic differential games and viscosity solutions of the Isaacs equations arising from such games are studied via a nonlinear semigroup related to DPP (the min-max principle, to be precise). Using semi-discretization arguments, we construct the nonlinear semigroups whose generators provide lower and upper Isaacs equations. Concerning partially observable control problems, we refer to stochastic parabolic equations driven by colored Wiener noises, in particular, the Zakai equation. The existence and uniqueness of solutions and regularities as well as Itô's formula are stated. A control problem for the Zakai equations has a nonlinear semigroup whose generator provides the HJB equation on a Banach space. The value function turns out to be a unique viscosity solution for the HJB equation under mild conditions. This edition provides a more generalized treatment of the topic than does the earlier book Lectures on Stochastic Control Theory (ISI Lecture Notes 9), where time-homogeneous cases are dealt with. Here, for finite time-horizon control problems, DPP was formulated as a one-parameter nonlinear semigroup, whose generator provides the HJB equation, by using a time-discretization method. The semigroup corresponds to the value function and is characterized as the envelope of Markovian transition semigroups of responses for constant control processes. Besides finite time-horizon controls, the book discusses control-stopping problems in the same frameworks.

As is well known, Pontryagin's maximum principle and Bellman's dynamic programming are the two principal and most commonly used approaches in solving stochastic optimal control problems. * An interesting phenomenon one can observe from the literature is that these two approaches have been developed separately and independently. Since both methods are used to investigate the same problems, a natural question one will ask is the fol lowing: (Q) What is the relationship betwccn the maximum principlc and dy namic programming in stochastic optimal controls? There did exist some researches (prior to the 1980s) on the relationship between these two. Nevertheless, the results usually werestated in heuristic terms and proved under rather restrictive assumptions, which were not satisfied in most cases. In the statement of a Pontryagin-type maximum principle there is an adjoint equation, which is an ordinary differential equation (ODE) in the (finite-dimensional) deterministic case and a stochastic differential equation (SDE) in the stochastic case. The system consisting of the adjoint equa tion, the original state equation, and the maximum condition is referred to as an (extended) Hamiltonian system. On the other hand, in Bellman's dynamic programming, there is a partial differential equation (PDE), of first order in the (finite-dimensional) deterministic case and of second or der in the stochastic case. This is known as a Hamilton-Jacobi-Bellman (HJB) equation.

The Wiley-Interscience Paperback Series consists of selected booksthat have been made more accessible to consumers in an effort toincrease global appeal and general circulation. With these newunabridged softcover volumes, Wiley hopes to extend the lives ofthese works by making them available to future generations ofstatisticians, mathematicians, and scientists. "This text is unique in bringing together so many resultshitherto found only in part in other texts and papers. . . . Thetext is fairly self-contained, inclusive of some basic mathematicalresults needed, and provides a rich diet of examples, applications,and exercises. The bibliographical material at the end of eachchapter is excellent, not only from a historical perspective, butbecause it is valuable for researchers in acquiring a goodperspective of the MDP research potential." —Zentralblatt fur Mathematik ". . . it is of great value to advanced-level students,researchers, and professional practitioners of this field to havenow a complete volume (with more than 600 pages) devoted to thistopic. . . . Markov Decision Processes: Discrete Stochastic DynamicProgramming represents an up-to-date, unified, and rigoroustreatment of theoretical and computational aspects of discrete-timeMarkov decision processes." —Journal of the American Statistical Association

Discussion Paper 816. Solving Stochastic Dynamic Programming Problems Using Rules of ThumbSolving Stochastic Dynamic Programming Problems Using Rules of ThumbKingston, Ont. : Institute for Economic Research, Queen's UniversityAdvances in Stochastic Dynamic Programming for Operations ManagementLogos Verlag Berlin GmbH

I develop a numerical method that combines functional approximations and dynamic programming to solve high-dimensional discrete-time stochastic control problems under general constraints. The method relies on three building blocks: first, a quasi-random grid and the radial basis function method are used to discretize and interpolate the high-dimensional state space; second, to incorporate constraints, the method of Lagrange multipliers is applied to obtain the first order optimality conditions; third, the conditional expectation of the value function is approximated by a second order polynomial basis, estimated using ordinary least squares regressions. To reduce the approximation error, I introduce the test region iterative contraction (TRIC) method to shrink the approximation region around the optimal solution. I apply the method to two Finance applications: a) dynamic portfolio choice with constraints, a continuous control problem; b) dynamic portfolio choice with capital gain taxation, a high-dimensional singular control problem.

This two-volume set of texts explores the central facts and ideas of stochastic processes, illustrating their use in models based on applied and theoretical investigations. They demonstrate the interdependence of three areas of study that usually receive separate treatments: stochastic processes, operating characteristics of stochastic systems, and stochastic optimization. Comprehensive in its scope, they emphasize the practical importance, intellectual stimulation, and mathematical elegance of stochastic models and are intended primarily as graduate-level texts.

Consisting of two parts, this book presents papers describing publicly available stochastic programming systems that are operational. It presents a diverse collection of application papers in areas such as production, supply chain and scheduling, gaming, environmental and pollution control, financial modeling, telecommunications, and electricity.

This book provides a straightforward overview for every researcher interested in stochastic dynamic vehicle routing problems (SDVRPs). The book is written for both the applied researcher looking for suitable solution approaches for particular problems as well as for the theoretical researcher looking for effective and efficient methods of stochastic dynamic optimization and approximate dynamic programming (ADP). To this end, the book contains two parts. In the first part, the general methodology required for modeling and approaching SDVRPs is presented. It presents adapted and new, general anticipatory methods of ADP tailored to the needs of dynamic vehicle routing. Since stochastic dynamic optimization is often complex and may not always be intuitive on first glance, the author accompanies the ADP-methodology with illustrative examples from the field of SDVRPs. The second part of this book then depicts the application of the theory to a specific SDVRP. The process starts from the real-world application. The author describes a SDVRP with stochastic customer requests often addressed in the literature, and then shows in detail how this problem can be modeled as a Markov decision process and presents several anticipatory solution approaches based on ADP. In an extensive computational study, he shows the advantages of the presented approaches compared to conventional heuristics. To allow deep insights in the functionality of ADP, he presents a comprehensive analysis of the ADP approaches.

The main purpose of the book is to show how a viscosity approach can be used to tackle control problems in insurance. The problems covered are the maximization of survival probability as well as the maximization of dividends in the classical collective risk model. The authors consider the possibility of controlling the risk process by reinsurance as well as by investments. They show that optimal value functions are characterized as either the unique or the smallest viscosity solution of the associated Hamilton-Jacobi-Bellman equation; they also study the structure of the optimal strategies and show how to find them. The viscosity approach was widely used in control problems related to mathematical finance but until quite recently it was not used to solve control problems related to actuarial mathematical science. This book is designed to familiarize the reader on how to use this approach. The intended audience is graduate students as well as researchers in this area.

Decision Theory An Introduction to Dynamic Programming and Sequential Decisions John Bather University of Sussex, UK Mathematical induction, and its use in solving optimization problems, is a topic of great interest with many applications. It enables us to study multistage decision problems by proceeding backwards in time, using a method called dynamic programming. All the techniques needed to solve the various problems are explained, and the author's fluent style will leave the reader with an avid interest in the subject. * Tailored to the needs of students of optimization and decision theory * Written in a lucid style with numerous examples and applications * Coverage of deterministic models: maximizing utilities, directed networks, shortest paths, critical path analysis, scheduling and convexity * Coverage of stochastic models: stochastic dynamic programming, optimal stopping problems and other special topics * Coverage of advanced topics: Markov decision processes, minimizing expected costs, policy improvements and problems with unknown statistical parameters * Contains exercises at the end of each chapter, with hints in an appendix Aimed primarily at students of mathematics and statistics, the lucid text will also appeal to engineering and science students and those working in the areas of optimization and operations research.

Current research results in stochastic and deterministic global optimization including single and multiple objectives are explored and presented in this book by leading specialists from various fields. Contributions include applications to multidimensional data visualization, regression, survey calibration, inventory management, timetabling, chemical engineering, energy systems, and competitive facility location. Graduate students, researchers, and scientists in computer science, numerical analysis, optimization, and applied mathematics will be fascinated by the theoretical, computational, and application-oriented aspects of stochastic and deterministic global optimization explored in this book. This volume is dedicated to the 70th birthday of Antanas Žilinskas who is a leading world expert in global optimization. Professor Žilinskas's research has concentrated on studying models for the objective function, the development and implementation of efficient algorithms for global optimization with single and multiple objectives, and application of algorithms for solving real-world practical problems.

Josef Anton Strini analyzes a special stochastic optimal control problem. The problem under study arose from a dynamic cash management model in finance, where decisions about the dividend and financing policies of a firm have to be made. Additionally, using the dynamic programming approach, he extends the present discourse by the formal derivation of the Hamilton-Jacobi-Bellman equation and by examining the verification step carefully. Finally, the treatment is completed by solving the problem numerically.

Many tasks in operations management require the solution of complex optimization problems. Problems in which decisions are taken sequentially over time can be modeled and solved by dynamic programming. Real-world dynamic programming problems, however, exhibit complexity that cannot be handled by conventional solution techniques. This complexity may stem from large state and solution spaces, huge sets of possible actions, non-convexities in the objective function, and uncertainty. In this book, three highly complex real-world problems from the domain of operations management are modeled and solved by newly developed solution techniques based on stochastic dynamic programming. First, the problem of optimally scheduling participating demand units in an energy transmission network is considered. These units are scheduled such that total cost of supplying demand for electric energy is minimized under uncertainty in demand and generation. Second, the integrated problem of investment in and optimal operations of a network of battery swap stations under uncertain demand and energy prices is modeled and solved. Third, the inventory control problem of a multi-channel retailer selling through independent sales channels is modeled and optimality conditions for replenishment policies of simple structure are proven. This book introduces efficient approximation techniques based on approximate dynamic programming (ADP) and extends existing proximal point algorithms to the stochastic case. The methods are applicable to a wide variety of dynamic programming problems of high dimension.

The availability of today's online information systems rapidly increases the relevance of dynamic decision making within a large number of operational contexts. Whenever a sequence of interdependent decisions occurs, making a single decision raises the need for anticipation of its future impact on the entire decision process. Anticipatory support is needed for a broad variety of dynamic and stochastic decision problems from different operational contexts such as finance, energy management, manufacturing and transportation. Example problems include asset allocation, feed-in of electricity produced by wind power as well as scheduling and routing. All these problems entail a sequence of decisions contributing to an overall goal and taking place in the course of a certain period of time. Each of the decisions is derived by solution of an optimization problem. As a consequence a stochastic and dynamic decision problem resolves into a series of optimization problems to be formulated and solved by anticipation of the remaining decision process. However, actually solving a dynamic decision problem by means of approximate dynamic programming still is a major scientific challenge. Most of the work done so far is devoted to problems allowing for formulation of the underlying optimization problems as linear programs. Problem domains like scheduling and routing, where linear programming typically does not produce a significant benefit for problem solving, have not been considered so far. Therefore, the industry demand for dynamic scheduling and routing is still predominantly satisfied by purely heuristic approaches to anticipatory decision making. Although this may work well for certain dynamic decision problems, these approaches lack transferability of findings to other, related problems. This book has serves two major purposes: ? It provides a comprehensive and unique view of anticipatory optimization for dynamic decision making. It fully integrates Markov decision processes, dynamic programming, data mining and optimization and introduces a new perspective on approximate dynamic programming. Moreover, the book identifies different degrees of anticipation, enabling an assessment of specific approaches to dynamic decision making. ? It shows for the first time how to successfully solve a dynamic vehicle routing problem by approximate dynamic programming. It elaborates on every building block required for this kind of approach to dynamic vehicle routing. Thereby the book has a pioneering character and is intended to provide a footing for the dynamic vehicle routing community.

This thesis describes a stochastic ship routing problem with inventory management. The problem involves finding a set of least costs routes for a fleet of ships transporting a single commodity when the demand for the commodity is uncertain. Storage at consumption and supply ports is limited and inventory levels are monitored in the model. Consumer demands are at a constant rate within each time period in the deterministic problem, and in the stochastic problem, the demand rate for a period is not known until the beginning of that period. The demand situation in each time period can be described by a scenario tree with corresponding probabilities. Several possible solution approaches for solving the problem are studied in the thesis. This problem can be formulated as a mixed integer programming (MIP) model. However solving the problem this way is very time consuming even for a deterministic problem with small problem size. In order to solve the stochastic problem, we develop a decomposition formulation and solve it using a

Branch and Price framework. A master problem (set partitioning with extra inventory constraints) is built, and the subproblems, one for each ship, involve solving stochastic dynamic programming problems to generate columns for the master problem. Each column corresponds to one possible tree of schedules for one ship giving the schedule for the ship for all demand scenarios. In each branch-and-bound node, the node problem is solved by iterating between the master problem and the subproblems. Dual variables can be obtained solving the master problem and are used in the subproblems to generate the most promising columns for the master problem. Computational results are given showing that medium sized problems can be solved successfully. Several extensions to the original model are developed, including a variable speed model, a diverting model, and a model which allows ships to do extra tasks in return for a bonus. Possible solution approaches for solving the variable speed and the diverting model are presented and computational results are given.

This book presents the latest findings on stochastic dynamic programming models and on solving optimal control problems in networks. It includes the authors' new findings on determining the optimal solution of discrete optimal control problems in networks and on solving game variants of Markov decision problems in the context of computational networks. First, the book studies the finite state space of Markov processes and reviews the existing methods and algorithms for determining the main characteristics in Markov chains, before proposing new approaches based on dynamic programming and combinatorial methods. Chapter two is dedicated to infinite horizon stochastic discrete optimal control models and Markov decision problems with average and expected total discounted optimization criteria, while Chapter three develops a special game-theoretical approach to Markov decision processes and stochastic discrete optimal control problems. In closing, the book's final chapter is devoted to finite horizon stochastic control problems and Markov decision processes. The algorithms developed represent a valuable contribution to the important field of computational network theory.

The authors consider the solution of stochastic dynamic programs using sample path estimates. Applying the theory of large deviations, they derive probability error bounds associated with the convergence of the estimated optimal policy to the true optimal policy, for finite horizon problems. These bounds decay at an exponential rate, in contrast with the usual canonical (inverse) square root rate associated with estimation of the value (cost-to-go) function itself. These results have practical implications for Monte Carlo simulation-based solution approaches to stochastic dynamic programming problems where it is impractical to extract the explicit transition probabilities of the underlying system model.

The concept of a system as an entity in its own right has emerged with increasing force in the past few decades in, for example, the areas of electrical and control engineering, economics, ecology, urban structures, automaton theory, operational research and industry. The more definite concept of a large-scale system is implicit in these applications, but is particularly evident in fields such as the study of communication networks, computer networks and neural networks. The Wiley-Interscience Series in Systems and Optimization has been established to serve the needs of researchers in these rapidly developing fields. It is intended for works concerned with developments in quantitative systems theory, applications of such theory in areas of interest, or associated methodology. Of related interest Stochastic Programming Peter Kall, University of Zurich, Switzerland and Stein W. Wallace, University of Trondheim, Norway Stochastic Programming is the first textbook to provide a thorough and self-contained introduction to the subject. Carefully written to cover all necessary background material from both linear and non-linear programming, as well as probability theory, the book draws together the methods and techniques previously described in disparate sources. After introducing the terms and modelling issues when randomness is introduced in a deterministic mathematical programming model, the authors cover decision trees and dynamic programming, recourse problems, probabilistic constraints, preprocessing and network problems. Exercises are provided at the end of each chapter. Throughout, the emphasis is on the appropriate use of the techniques, rather than on the underlying mathematical proofs and theories, making the book ideal for researchers and students in mathematical programming and operations research who wish to develop their skills in stochastic programming.

This volume in the SpringerBriefs in Energy series offers a systematic review of unit commitment (UC) problems in electrical power generation. It updates texts written in the late 1990s and early 2000s by including the fundamentals of both UC and state-of-the-art modeling as well as solution algorithms and highlighting stochastic models and mixed-integer programming techniques. The UC problems are mostly formulated as mixed-integer linear programs, although there are many variants. A number of algorithms have been developed for, or applied to, UC problems, including dynamic programming, Lagrangian relaxation, general mixed-integer programming algorithms, and Benders decomposition. In addition the book discusses the recent trends in solving UC problems, especially stochastic programming models, and advanced techniques to handle large numbers of integer- decision variables due to scenario propagation

This book presents the first part of a planned two-volume series devoted to a systematic exposition of some recent developments in the theory of discrete-time Markov control processes (MCPs). Interest is mainly confined to MCPs with Borel state and control (or action) spaces, and possibly unbounded costs and noncompact control constraint sets. MCPs are a class of stochastic control problems, also known as Markov decision processes, controlled Markov processes, or stochastic dynamic pro grams; sometimes, particularly when the state space is a countable set, they are also called Markov decision (or controlled Markov) chains. Regardless of the name used, MCPs appear in many fields, for example, engineering, economics, operations research, statistics, renewable and nonrenewable re source management, (control of) epidemics, etc. However, most of the lit erature (say, at least 90%) is concentrated on MCPs for which (a) the state space is a countable set, and/or (b) the costs-per-stage are bounded, and/or (c) the control constraint sets are compact. But curiously enough, the most widely used control model in engineering and economics--namely the LQ (Linear system/Quadratic cost) model-satisfies none of these conditions. Moreover, when dealing with "partially observable" systems) a standard approach is to transform them into equivalent "completely observable" sys tems in a larger state space (in fact, a space of probability measures), which is uncountable even if the original state process is finite-valued.

This is a book on stochastic dynamic macroeconomics from a Keynesian perspective. It shows that including Keynesian features in intertemporal models considerably contributes to resolve major puzzles arising in the context of the Dynamic General Equilibrium (DGE) model. It also demonstrates that including microeconomic intertemporal behavior of economic agents in macroeconomics is not inconsistent with Keynesian economics.

This book consists of, apart from the introduction, the chapters - DA Stochastic Dynamic Programming with Random Disturbances, - The Problem of Stochastic Dynamic Distance Optimal Partitioning (SDDP problem), - Partitions-Requirements-Matrices (PRMs). DA („decision after“) stochastic dynamic programming with random disturbances“ is characterized by the fact that these random disturbances are observed before the decision is made at each stage. In the past, only very moderate attention was given to problems with this characteristic. In Chapter 2 specific properties of DA stochastic dynamic programming problems are worked out for theoretical characterization and for more efficient solution strategies of such problems. The (DA) Stochastic Dynamic Distance Optimal Partitioning problem (SDDP problem) is an extremely complex Operations Research problem. It shows several connections with other problems of operations research and informatics such as stochastic dynamic transportation and facility location problems or metric task systems and more specific k-server problems. Partitions of integers as states of SDDP problems require an enormous amount of storage space for the corresponding computer programs. Investigations of inherent characteristic structures of SDDP problems are also important as a basis for heuristics. Partitions-requirements-matrices (PRMs) (Chapter 4) are matrices of transition probabilities of SDDP problems which are formulated as Markov decision processes. PRMs „in the strict meaning“ include optimal decisions of certain reduced SDDP problems, as is shown (in many cases) toward the end of the book. PRMs (in the strict meaning) themselves represent interesting (almost selfevident) combinatorial structures, which are not otherwise found in literature. In order to understand the investigations of this book, previous knowledge about stochastic dynamic Programming and Markov decision processes is useful, however not absolutely necessary since the concerned models are developed from scratch.

This book has the purpose of providing the "state of the arts" concerning bio-economic modelling dealing with agricultural systems. In most cases, the contributions use a methodology combining the use of biophysical and economic models, in all cases, an engineering production function approach is totally or partially applied. This practice is being developed in the last years as a response to concrete policy matters: agricultural policies are increasingly combined with environmental and natural resources policies, and this reality involves the need of an integrated assessment, that current economic models are not able to provide.

The model developed solves the multiperiod multireservoir water resources problem with stochastic inflows. Of unique importance is the development of a generalized network model which solves nonlinear nonseparable quadratic problems. Quadratic functions are used to measure the future value of water to the system. The nonseparable form stems from the realization that interaction exists between the benefits to be gained from a multireservoir system. Historically this interactive nature has been ignored due to the computational difficulty of measuring and solving such relationships. Also developed is a stochastic dynamic programming approach which utilizes the results to the network optimization as data for a least squares regression analysis. A quadratic function is fit to this data and is used to represent the future value of water to the system for the next period in the dynamic programming approach. This functional representation of the future value of water replaces the standard discrete matrix representation of dynamic programming and greatly reduces the dimensionality problems associated with the dynamic programming approach. In the end, this work represents a rare combination of generalized-nonlinear network flow programming, stochastic dynamic programming and regression analysis.

Uncertainties and changes are pervasive characteristics of modern systems involving interactions between humans, economics, nature and technology. These systems are often too complex to allow for precise evaluations and, as a result, the lack of proper management (control) may create significant risks. In order to develop robust strategies we need approaches which explic itly deal with uncertainties, risks and changing conditions. One rather general approach is to characterize (explicitly or implicitly) uncertainties by objec tive or subjective probabilities (measures of confidence or belief). This leads us to stochastic optimization problems which can rarely be solved by using the standard deterministic optimization and optimal control methods. In the stochastic optimization the accent is on problems with a large number of deci sion and random variables, and consequently the focus ofattention is directed to efficient solution procedures rather than to (analytical) closed-form solu tions. Objective and constraint functions of dynamic stochastic optimization problems have the form of multidimensional integrals of rather involved in that may have a nonsmooth and even discontinuous character - the tegrands typical situation for "hit-or-miss" type of decision making problems involving irreversibility ofdecisions or/and abrupt changes ofthe system. In general, the exact evaluation of such functions (as is assumed in the standard optimization and control theory) is practically impossible. Also, the problem does not often possess the separability properties that allow to derive the standard in control theory recursive (Bellman) equations.

Portrays dynamic programming as a methodology, identifying its constituent components, and explaining how it approaches problems and tackles them. Does not consider it as a practical tool, nor how it might address any actual situations in the real world. Assumes calculus, set theory, and some optimi

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.

Providing an introduction to stochastic optimal control in infinite dimension, this book gives a complete account of the theory of second-order HJB equations in infinite-dimensional Hilbert spaces, focusing on its applicability to associated stochastic optimal control problems. It features a general introduction to optimal stochastic control, including basic results (e.g. the dynamic programming principle) with proofs,

and provides examples of applications. A complete and up-to-date exposition of the existing theory of viscosity solutions and regular solutions of second-order HJB equations in Hilbert spaces is given, together with an extensive survey of other methods, with a full bibliography. In particular, Chapter 6, written by M. Fuhrman and G. Tessitore, surveys the theory of regular solutions of HJB equations arising in infinite-dimensional stochastic control, via BSDEs. The book is of interest to both pure and applied researchers working in the control theory of stochastic PDEs, and in PDEs in infinite dimension. Readers from other fields who want to learn the basic theory will also find it useful. The prerequisites are: standard functional analysis, the theory of semigroups of operators and its use in the study of PDEs, some knowledge of the dynamic programming approach to stochastic optimal control problems in finite dimension, and the basics of stochastic analysis and stochastic equations in infinite-dimensional spaces.
Praise for the First Edition "Finally, a book devoted to dynamic programming and written using the language of operations research (OR)! This beautiful book fills a gap in the libraries of OR specialists and practitioners." —Computing Reviews This new edition showcases a focus on modeling and computation for complex classes of approximate dynamic programming problems Understanding approximate dynamic programming (ADP) is vital in order to develop practical and high-quality solutions to complex industrial problems, particularly when those problems involve making decisions in the presence of uncertainty. Approximate Dynamic Programming, Second Edition uniquely integrates four distinct disciplines—Markov decision processes, mathematical programming, simulation, and statistics—to demonstrate how to successfully approach, model, and solve a wide range of real-life problems using ADP. The book continues to bridge the gap between computer science, simulation, and operations research and now adopts the notation and vocabulary of reinforcement learning as well as stochastic search and simulation optimization. The author outlines the essential algorithms that serve as a starting point in the design of practical solutions for real problems. The three curses of dimensionality that impact complex problems are introduced and detailed coverage of implementation challenges is provided. The Second Edition also features: A new chapter describing four fundamental classes of policies for working with diverse stochastic optimization problems: myopic policies, look-ahead policies, policy function approximations, and policies based on value function approximations A new chapter on policy search that brings together stochastic search and simulation optimization concepts and introduces a new class of optimal learning strategies Updated coverage of the exploration exploitation problem in ADP, now including a recently developed method for doing active learning in the presence of a physical state, using the concept of the knowledge gradient A new sequence of chapters describing statistical methods for approximating value functions, estimating the value of a fixed policy, and value function approximation while searching for optimal policies The presented coverage of ADP emphasizes models and algorithms, focusing on related applications and computation while also discussing the theoretical side of the topic that explores proofs of convergence and rate of convergence. A related website features an ongoing discussion of the evolving fields of approximation dynamic programming and reinforcement learning, along with additional readings, software, and datasets. Requiring only a basic understanding of statistics and probability, Approximate Dynamic Programming, Second Edition is an excellent book for industrial engineering and operations research courses at the upper-undergraduate and graduate levels. It also serves as a valuable reference for researchers and professionals who utilize dynamic programming, stochastic programming, and control theory to solve problems in their everyday work.
The aim of stochastic programming is to find optimal decisions in problems which involve uncertain data. This field is currently developing rapidly with contributions from many disciplines including operations research, mathematics, and probability. At the same time, it is now being applied in a wide variety of subjects ranging from agriculture to financial planning and from industrial engineering to computer networks. This textbook provides a first course in stochastic programming suitable for students with a basic knowledge of linear programming, elementary analysis, and probability. The authors aim to present a broad overview of the main themes and methods of the subject. Its prime goal is to help students develop an intuition on how to model uncertainty into mathematical problems, what uncertainty changes bring to the decision process, and what techniques help to manage uncertainty in solving the problems. In this extensively updated new edition there is more material on methods and examples including several new approaches for discrete variables, new results on risk measures in modeling and Monte Carlo sampling methods, a new chapter on relationships to other methods including approximate dynamic programming, robust optimization and online methods. The book is highly illustrated with chapter summaries and many examples and exercises. Students, researchers and practitioners in operations research and the optimization area will find it particularly of interest. Review of First Edition: "The discussion on modeling issues, the large number of examples used to illustrate the material, and the breadth of the coverage make 'Introduction to Stochastic Programming' an ideal textbook for the area." (Interfaces, 1998)
Dynamic Programming and Stochastic Control
Applied Stochastic Models and Control for Finance and Insurance presents at an introductory level some essential stochastic models applied in economics, finance and insurance. Markov chains, random walks, stochastic differential equations and other stochastic processes are used throughout the book and systematically applied to economic and financial applications. In addition, a dynamic programming framework is used to deal with some basic optimization problems. The book begins by introducing problems of economics, finance and insurance which involve time, uncertainty and risk. A number of cases are treated in detail, spanning risk management, volatility, memory, the time structure of preferences, interest rates and yields, etc. The second and third chapters provide an introduction to stochastic models and their application. Stochastic differential equations and stochastic calculus are presented in an intuitive manner, and numerous applications and exercises are used to facilitate their understanding and their use in Chapter 3. A number of other processes which are increasingly used in finance and insurance are introduced in Chapter 4. In the fifth chapter, ARCH and GARCH models are presented and their application to modeling volatility is emphasized. An outline of decision-making procedures is presented in Chapter 6. Furthermore, we also introduce the essentials of stochastic dynamic programming and control, and provide first steps for the student who seeks to apply these techniques. Finally, in Chapter 7, numerical techniques and approximations to stochastic processes are examined. This book can be used in business, economics, financial engineering and decision sciences schools for second year Master's students, as well as in a number of courses widely given in departments of statistics, systems and decision sciences.
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.÷ In÷Lectures on Stochastic Programming: Modeling and Theory, Second Edition, the authors introduce new material to reflect recent developments in stochastic programming, including: an analytical description of the tangent and normal cones of chance constrained sets; analysis of optimality conditions applied to nonconvex problems; a discussion of the stochastic dual dynamic programming method; an extended discussion of law invariant coherent risk measures and their Kusuoka representations; and in-depth analysis of dynamic risk measures and concepts of time consistency, including several new results.÷
Linear programs have been formulated for many practical situations that require decisions made periodically through time. These dynamic linear programs often involve uncertainties. Deterministic solutions of these problems may lead to costly incorrect decisions, and, when a stochastic solution is attempted the problem may become too large. In this report, we present methods for reducing the computational cost of these stochastic programs, and we show conditions under which the stochastic program need not be solved. Our methods are based on the large-scale programming techniques of decomposition,

partitioning, and basic factorization. (Author).