Preface

The goal of this book is to prepare readers to apply the optimal control theory to nonlinear processes beyond the standard applications. The examples investigated in depth are drawn from drug policy, corruption, and counter-terror, so this book will be of particular interest to economists, epidemiologists, and other scholars interested in the economics of crime, health, and violence. However, the book is designed to simultaneously appeal to an even larger audience of students in mathematics, economics, biosciences, and operations research for whom it fills a gap in the literature between "cookbooks," that give recipes but not deeper understanding, and more formal mathematical textbooks that do not explain how to apply the theory to practical problems.

This innovative balance is inspired by what we and our students have discovered to be essential to productively harnessing optimal control to address novel application areas. In our experience, new models are rarely just minor variants of others whose solution already appears in the literature. So searching through a typical applications-oriented textbook for the right recipe is fruitless; the new model usually has some twist or idiosyncrasy that requires modifying the cookbook recipes, not just applying them. Such modifications require understanding grounded in a more mathematical presentation than cookbooks usually provide. On the other hand, students who read only the more mathematical presentations often have little clue how to go about solving applied problems, either in terms of writing down the original model or in terms of carrying out the numerical computations. For the curious reader, we provide references to formal proofs of key theorems, but do not always reproduce them here, preferring to focus on intuition and applications.

Indeed, even standard applications-oriented textbooks do not always prepare students for studying novel domains. Some authors focus on a specific application; their books are wonderful for students who know they will devote their careers to problems within that one application area, but students' interests and future career trajectories are rarely so well defined. Other authors sprinkle their texts with examples drawn from across the spectrum of application areas. This opens students' minds to the power and flexibility of these methods, so we do some of that ourselves in suggestions for further reading and end of chapter exercises drawn from optimal growth, marketing, production/inventory management, and environmental planning, among other applications. The risk of a broad but shallow approach, however, is that none of the examples is developed in depth, and careers cannot be made by working textbook-like toy problems.

So a second principal innovation of this text is developing in detail examples from three exciting and topical areas: drug, corruption, and counter-terror control. This treatment in depth allows students to learn by example some of the art and craft of good modeling, including the idea of developing layered models starting with highly stylized depictions and then building in greater and greater realism. Compared to most texts, we spend more time discussing why certain functional forms might or might not capture well the underlying system dynamics and how to think about the limitations of modeling approximations, rather than writing down the complete model all at once and simply asserting that it is an accurate reflection of the problem at hand.

We do not expect the majority of students to work professionally in these areas, although the fields of health economics and the economics of crime are growing rapidly. From classroom experience, however, we know students enjoy seeing how mathematics can be used to address these serious and intrinsically fascinating problems. Thinking about how to fight terror, corruption, or illegal drugs is simply more fun for most students than thinking about problems in firms (e.g., inventory management) or mechanics (e.g., balancing an inverted pendulum). Furthermore, the case studies are not just cartoons; the model development in this text mirrors the trajectory of the actual research. This allows the reader to experience some of the same excitement we felt when discovering new results or briefing interpretations that resonated with actual policy-makers with whom we have worked on these problems.

A third innovation is extensive treatment of models with multiple equilibria separated by DNSS points, which are a sort of tipping point. We have found that nonlinear models often have such features, at least for certain parameter constellations, and misleading results can be obtained if the analyst is not sensitized to this possibility.

From a pedagogical perspective, the fourth and perhaps most important innovation is that we have developed a toolbox (`OCMat`)[1] using the MATLAB language to analyze optimal control problems specifically for discounted infinite time horizon models. This toolbox allows (in general) an immediate analysis of the models after the problem formulation, since the necessary files for computation are generated (semi)-automatically. The core framework is the formulation of a boundary value problem, where the infinite time horizon is truncated to a finite horizon and the usual transversality conditions are replaced by so-called asymptotic transversality conditions. This approach is very flexible. It allows consideration of mixed and pure state constraints

[1] `OCMat` is available via `http://www.eos.tuwien.ac.at/OR/OCMat`.

and multi-stage problems, as well as the computation and continuation of DNSS points, junction points between boundary and interior arcs, and heteroclinic connections. Multiple examples are offered for both one and higher-dimensional models.

We are aware that not all readers come to optimal control with the same mathematical background. Hence, after a brief introduction (Chap. 1), Part I of this book continues by summarizing prerequisite material on continuous dynamical systems (Chap. 2). We stress the term "summarizes." A student who has never taken a course in differential equations cannot learn differential equations just by reading Chap. 2. However, for a student who studied differential equations a few years ago or who merely passed the course without attaining a deep understanding of all the material, Chap. 2 offers a valuable review focused specifically on the topics needed to understand optimal control. Appendix A does the same for material from linear algebra, topology, and calculus. That appendix also defines notational conventions, most of which are standard and should be familiar, but are included nonetheless for completeness.

Part II begins with material specifically on nonlinear optimal control. Chapter 3 is in some sense the heart of the book, introducing, explaining, and applying Pontryagin's Maximum Principle. Truly understanding how to apply the theory requires computing numerical solutions, not just proving propositions, so the text refers the student forward to Chap. 7 (Numerical Methods) at appropriate places for tutorials on how to perform the computations. Chapter 4 takes an innovative approach to deepening students' intuition concerning Pontryagin's principle, "deriving" it from standard discrete optimization by viewing continuous optimization as the limit of better and better discrete approximations to the continuous system, where we put "deriving" in quotes because the focus is on building intuition, not mathematical rigor. For a more mathematical class, particularly one for operations research students with a strong background in discrete optimization, Chap. 4 provides some basic ideas for the theory presented in Chap. 3. For others, Chap. 4 may be better appreciated later in the course, after the students have seen and worked on more applied problems.

In either case, the discussion in Chap. 5 of multiple equilibria should be read before proceeding to Part III and, as with Chap. 3, it includes pointers forward to Chap. 7 (Numerical Methods).

Real applications often involve more than just one-dimensional models, so we view the discussion in Chap. 6 of higher-dimensional as another core chapter. What makes higher-dimensional models so much harder is not primarily more advanced theory, but rather practical problems encountered when trying to analyze higher-dimensional systems. Hence, Chap. 6 comprises three detailed case studies, one each from drugs, corruption, and counter-terror. Each case study offers not just a solution, but also a description of the thought processes – including even some of the dead ends we encountered – that went into solving the problem, again with forward references to the parts of Chap. 7 needed for the numerical computations. We believe a complete course on optimal control ought to cover at least one or two of these case studies, but they are not cumulative, so the instructor can pick any subset of them in any order. Chapter 8 discusses advanced topics which is not required material for a basic course in optimal control. Rather, we include just brief discussions of three areas (multi-stage problems, differential games, and distributed parameter models) to give students a sense of the possibilities that exist beyond the standard ordinary differential equation or "compartmental" models.