Preface

This book grew out of my undergraduate course Methods in Mathematical Biology, which I have taught at the National University of Singapore since 2014 and which in turn draws substantially from Simon Levin's long-running Methods of Mathematical Biology course at Princeton University. When I was designing the course, I wanted to implement all the numerical exercises in R, which has been a popular programming language among ecologists for over two decades. But when looking for a suitable textbook that covered the range of topics I wanted to cover, I could not find any that also used R. I instead prepared my own materials and then, several years later, I began to turn these materials into course notes, which eventually became this book.

This book is aimed at the upper undergraduate and graduate levels. It is assumed the reader will have a background in high school calculus, algebra, and probability, a strong interest in or familiarity with ecology, and a motivation to learn how to build and analyse mathematical models of ecological dynamics. Readers who tackle the R coding parts of the book (presented in boxes and exercises) should have a basic background in R, such as one might obtain from a biostatistics course, or have experience programming in another language. The programming philosophy presented here is to write code from first principles as far as possible, without relying on existing packages. The only package needed to run the example code in the boxes is `deSolve`, which implements differential equation solvers. The code for all the examples is also provided online for readers' convenience.

The primary aim of the book is to teach the reader to think like a theoretical ecologist. What is a theory? What is a model? How do we build models? What are the uses of models? In addition, the book seeks to impart mathematical skills relevant to theoretical ecology, mostly related to the analysis of dynamical systems. The R programming examples are deployed in the service of the book's primary goal of teaching theoretical ecology, but also teach basic coding principles relevant to ecology.

The book places particular emphasis on how to go from a verbal description of a biological problem to a mathematical one, and vice versa. In my experience, many students are comfortable with the biology and they can also successfully learn to analyse a given mathematical model, but they struggle with how to translate ideas between the two domains. This can be a limitation when they move beyond the classroom and into original research, because then they must usually build the mathematical models themselves. Likewise, when they read the ecological literature, the mathematical models presented are rarely the same as the models taught in class, and they must learn how to interpret the components of these models in terms of the underlying biology.

The book also focuses on the key question of what models are useful for. Again, in my experience, students are able to learn the details of particular models, but when asked about the strengths and weaknesses of a model, they often lack the confidence to draw their own

conclusions. This skill is essential for engaging critically with the scientific literature in students' future careers.

In terms of content, this book shares many topics with Alan Hastings's 1997 classic, *Population Biology*, although that book has no programming content and I omit the topics of population genetics and infectious disease dynamics in favour of giving a broader and more comprehensive treatment of macroecological topics. (For an introduction to population genetics modelling with applications in R, I recommend the recent book *Population Genetics with R: An Introduction for Life Scientists*, by Áki Láruson and Floyd Reed; for disease modelling, I recommend Ottar Bjørnstad's excellent 2018 book *Epidemics: Models and Data Using R*.) The topics covered here also overlap with the 2019 book *Quantitative Ecology: A New Unified Approach* by Lehman et al., which also uses R, but here I present a more comprehensive and rigorous mathematical treatment suited to a quantitatively inclined audience. Similarly, the book *A Primer of Ecology with R* by Henry Stevens uses R and gives a broad treatment of topics in general ecology, statistics, and theoretical ecology, but because of its breadth by necessity gives a less thorough treatment of the latter than is given here. This book also complements several other quantitative ecology books, such as Sarah Otto and Troy Day's *A Biologist's Guide to Mathematical Modeling in Ecology and Evolution*, that do not use R. This book is not a statistics book. It is about how to build mathematical models of ecological systems from mechanistic principles, rather than how to fit these models to data. Although I present several examples comparing model predictions to empirical data, I do not delve into statistical methods for rigorous model fitting. I leave this extremely broad topic to other R-focused books, such as *Mixed Models and Extensions in Ecology with R* by Alain Zuur et al.

This book contains 20 chapters on a range of ecological topics (Fig. 0.1). A typical chapter deals with one topic and centres around two or three models that are used to gain insights into the ecology. Mathematical content that is essential for understanding and analysing the models is presented throughout, with detailed mathematical derivations given in boxes. A few of the boxes give crash courses on mathematical topics that some readers may not be familiar with (matrices, complex numbers, partial differentiation). Boxes are also used to present R coding examples, which explore the models numerically. Exercises at the end of each chapter probe students on the conceptual, mathematical, and numerical (R coding) aspects of the focal topic. Exercises marked with asterisks are harder and should only be attempted by students who have successfully tackled the preceding exercises. Occasional exercises marked with double asterisks are particularly challenging. For the reader looking for a not-too-technical introduction to theoretical ecology, the book can be read without referring to boxes or doing the exercises. Readers looking for more depth may read each chapter first and then refer to boxes depending on needs and interests.

The most essential mathematical skills are imparted in Chapters 1–5, which deal with single-species models, and 9–10, which deal with multispecies models (these chapters are shaded in light grey in Fig. 0.1). Chapters 6 and 8, on age-structured models and spatial models, can be considered optional as later chapters do not depend on them. Chapter 7, on metapopulation models, is essential only for Chapter 15, on metacommunity theory. After Chapter 10, readers are mathematically equipped to handle a range of topics, from community ecology (Chapters 11–13 and 16–17; dark grey shading in Fig. 0.1), island biogeography

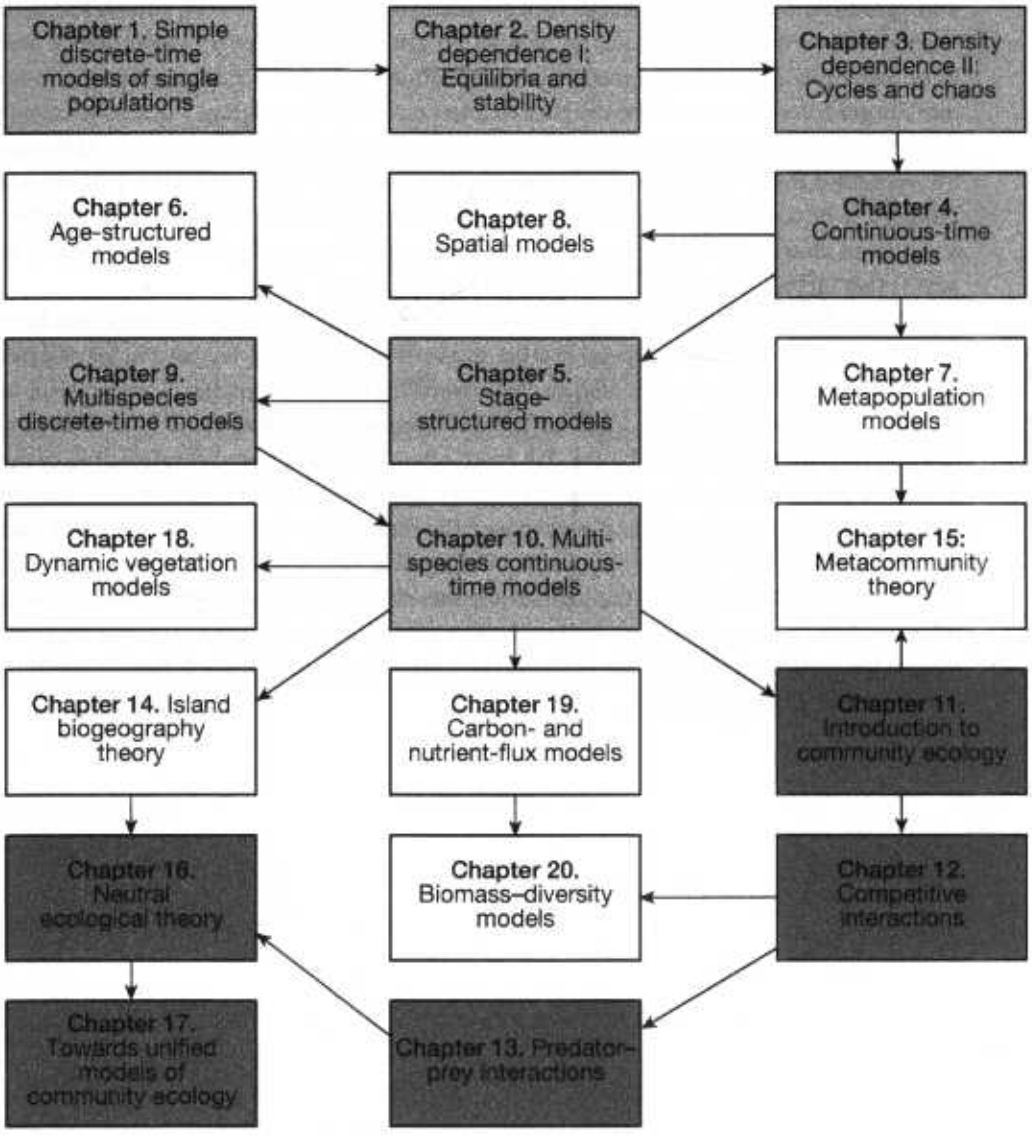


Fig. 0.1. A map of chapter dependencies in this book. Light grey boxes indicate core chapters that teach the central mathematical topics of the book. Dark grey boxes indicate a group of chapters relevant to community ecology.

theory (Chapter 14), dynamic vegetation models (Chapter 18), carbon- and nutrient-flux models (Chapter 19), and biomass–diversity models (Chapter 20). There is more content in the book than one could feasibly teach in a single-semester course. A recommended course structure would be one based on Chapters 1–5 and 9–10 (light-coloured boxes in Fig. 0.1), incorporating other chapters and potentially other resources depending on the interests of the lecturer and the needs of the students. For example, in a typical semester in my

undergraduate class, I teach material corresponding to Chapters 1–5, 8–14, and 16, and then augment this with one week each on population genetics, evolutionary game theory, and infectious disease modelling. I incorporate a selection of the remaining chapters each year into my graduate class, where for the later chapters especially I encourage students to engage more with the relevant primary literature. Each chapter of this book has a short reference list, separated into "References cited" and "Further reading." These reference lists are not intended to be comprehensive but rather to give readers different entry points into the literature: from classic studies to more modern works; and from purely theoretical papers to papers that include experimental work or data analysis.

I would like to acknowledge the many undergraduate and graduate students who have read the course notes that comprised earlier drafts of the book chapters and in many cases picked up typographical errors and identified sections that were unclear. I am grateful to Duncan Menge, who provided ideas and critical comments for Chapter 19. I would in particular like to thank Lynette Loke and Tak Fung, who each read a near-complete draft of the book and provided extensive feedback. Lynette's comments helped improve the clarity of the text for the general ecological reader, and she created the cup-and-ball diagrams that illustrate concepts of stability. Tak, with his technical expertise and attention to detail, picked up a number of subtle mathematical issues that would have been missed by most other readers. Three anonymous reviewers also provided constructive criticism that improved the final book. I would also like to thank my wife, Maria Elizabeth Rodríguez Ronderos, for her unfailing support.