# Contents

CHAPTER 1

# Introduction

Many budding ecologists have their imaginations captured by a seemingly simple question: why do we find different types and numbers of species in different places? The question is the same whether the setting is birds in the forest, plants along a mountainside, fish in lakes, invertebrates on a rocky shore, or microbes in the human body. Some parts of the answer to this question are glaringly obvious just from a short walk more or less anywhere on earth. Strolling through any city or town in eastern North America, we can see that the plant species growing in sidewalk cracks and dry roadsides are different from those growing in wet ditches, which are different still from those growing in wooded parks. Some birds reach very high abundance in dense urban areas, while others are found exclusively in wetlands or forests. So, we can observe everyday evidence that environmental variation selects for different species in different places (Fig. 1.1).

As we begin to look more closely, however, the story is not so simple. Some places that seem to present near-identical environmental conditions are nonetheless home to very different sets of species. Some pairs of species seem to live in very similar types of environments but almost never in the same physical place. Two places experiencing a very similar disturbance event (e.g., a drought or fire) subsequently follow very different successional trajectories. A hectare of one type of forest might contain 100-fold more species than a hectare of another type of forest. A major scientific challenge is thus to devise theories that can explain and predict such phenomena. Over the past 150 years ecologists have risen to this challenge, devising hundreds of conceptual or theoretical models that do just this. However, because almost every such model

**Figure 1.1.** The east-facing slope of Mont Saint-Joseph in Parc national du Mont Mégantic, Québec, illustrating spatial relationships between environmental conditions and community composition. The cold, upper part of the slope (~850–1100 m above sea level) is boreal forest (dark coloration) dominated by balsam fir (*Abies balsamea*). The lower slope is deciduous forest (light coloration) dominated by sugar maple (*Acer saccharum*). The photo was taken in springtime (8 May 2013), prior to the flushing of deciduous leaves. The foreground is relatively flat terrain (~400 m a.s.l.) composed mostly of a patchwork of young forest stands on private land, with a variety of different tree species. From left to right, the image spans roughly 4 km.

is relevant to at least one type of community somewhere on earth, the list of explanations for community patterns gets only ever longer, never shorter.

We are thus faced with a serious pedagogical challenge: how to conceptually organize theoretical ideas in community ecology as simply as possible to facilitate ecological understanding. We have for a long time organized ecological knowledge (in textbooks or other synthetic treatments) according to subareas into which researchers have self-organized rather than fundamental ecological processes that cut across these subareas. For example, a treatment of plant community ecology might have sections on herbivory, competition, disturbance, stress tolerance, dispersal, life-history tradeoffs, and so on (Crawley 1997, Gurevitch et al. 2006). Similarly, a conceptual treatment of community ecology might present many competing theories: island biogeography, priority effects, colonization-competition models, local resource–competition theory, neutral theory, metacommunity theory, and so on (Holyoak et al. 2005, Verhoef and Morin 2010, Morin 2011, Scheiner and Willig 2011, Mittelbach 2012). As a result, if each student in an undergraduate or graduate class is asked to write down a list of processes that can influence community composition and diversity (I have done this several times), the result will be a long list from each student, and collectively no fewer than 20–30 items.

The central argument to be developed in this book is as follows. Underlying all models of community dynamics are just four fundamental, or "high-level," processes: selection (among individuals of different species), ecological drift, dispersal, and speciation (Vellend 2010). These processes parallel the "big four" in evolutionary biology—selection, drift, migration, and mutation—and they allow us to organize knowledge in community ecology in a simpler way than by using the conventional approach. What seems like a jumble of independent theoretical perspectives can be understood as different mixtures of a few basic ingredients. By articulating a series of hypotheses and predictions based on the action of these four processes, we can thus build a general theory of ecological communities. As explained further in Chapter 2, the theory does not apply equally to all topics under the broad umbrella of community ecology. For example, models of species on the same trophic level interacting via competition and/or facilitation (sometimes called "horizontal" communities) fall cleanly within the theory, whereas models involving trophic interactions fit within the theory largely to the extent that they make predictions concerning properties of horizontal components of the larger food web (which they often do). Nonetheless, following the tradition set by MacArthur and Wilson (1967, *The Theory of Island Biogeography*) and Hubbell (2001, *The Unified Neutral Theory of Biodiversity and Biogeography*), I call my theory and therefore my book *The Theory of Ecological Communities*.

## 1.1. WHAT THIS BOOK IS

My overarching objective in this book is to present a synthetic perspective on community ecology that can help researchers and students better understand the linkages among the many theoretical ideas in the field. The initial sketch of these ideas was presented in Vellend (2010), and this book is a fully fleshed-out version of the theory, reiterating the key points of the earlier paper but going well beyond it in many ways:

- First, I more thoroughly place the theory of ecological communities in historical context (Chap. 3), and I present a novel perspective (gleaned from philosopher Elliott Sober) on why high-level processes (in this case selection, drift, dispersal, and speciation) represent an especially appropriate place to seek generality in community ecology (Chap. 4).
- I describe in detail how a vast number of different hypotheses and models in community ecology fit as constituents of the more general theory (Chap. 5).
- I provide simple computer code in the R language that (i) generates predictions for empirical testing, (ii) illustrates how changing a few basic "rules" of community dynamics reproduces a wide range of well-known models, and (iii) allows readers to explore such dynamics on their own (Chap. 6).
- After outlining some key motivations and challenges involved in empirical studies in ecology (Chap. 7), I then put the theory of ecological communities to work by systematically articulating hypotheses and predictions based on the action of selection (Chap. 8), drift and dispersal (Chap. 9), and speciation (Chap. 10), in each case evaluating empirical evidence supporting (or not) the predictions. In essence, Chapters 8–10 serve to reframe the corpus of empirical studies in community ecology according to a general theory that is considerably simpler than typically found in a textbook treatment of the discipline.

- Chapters 11 and 12 present some overarching conclusions and a look to the future.

### 1.1.1. Reading This Book as a Beginner, an Expert, or Something in Between

This book is aimed at senior undergraduate students, graduate students, and established researchers in ecology and evolutionary biology. It is the book I would have liked to read during grad school. I believe it presents the core conceptual material of community ecology in a new and unique way that makes it easier to grasp the nature of the key processes underlying community dynamics and how different approaches fit together. This has been my experience in using it as a teaching tool. I also hope to stimulate established researchers to think about what they do from a different perspective, and perhaps to influence how they teach community ecology themselves. Thus, I approached the writing of the book with the dual goals of pedagogy (beginning-student audience) and advancing a new way of thinking about theory in community ecology (expert audience). I suspect that readers who are somewhere on the pathway from beginner to expert—that is, grad students—have the most to gain from reading this book.

A pervasive challenge in scientific communication (including teaching) is to keep the most knowledgeable members of an audience engaged without "losing" those with the least preexisting knowledge of the topic. Readers can get the most out of this book if they are already somewhat familiar with the kinds of community-level patterns of species diversity and composition that ecologists aim to explain, as well as some of the factors commonly invoked to explain such patterns—environmental conditions, competition, disturbance, and so on. I begin explanations at a fairly basic level and provide what I consider the essential background (Chaps. 2–3), but even so, a full understanding of various historical advances in ecology (Chap. 3) and some of the more sophisticated empirical studies (Chaps. 8–11) requires delving into the primary literature. At the other end of the spectrum, expert readers will no doubt encounter sections they can skim, but I hope that all chapters of the book contain enough novel perspectives, approaches, or modes of traversing well-trodden ground to engage even the most expert reader. If you are an expert and pressed for time, you may choose to skip to the end of Chapter 3 (Sec. 3.4), where I begin the transition from background material to the details of my own distinct perspective and theory. Feedback on earlier versions of the book suggested that experts will find the most "new stuff" in the latter part of the book (Chaps. 8–12).

### 1.1.2. Unavoidable Trade-Offs

This book covers a very broad range of topics (models, questions, methods, etc.), which necessarily involves a trade-off with detail in several respects. First, the depth to which I explore each individual topic is limited. So, while readers will learn, for example, about the strengths and weaknesses of different approaches to testing for signatures of ecological drift or spatially variable selection, they will not learn all the detailed ins and outs of how to implement particular empirical methods. I am not myself an expert on all such details, and even for topics I do know quite well, I have deliberately limited the detail so as not to distract from the big-picture conceptual issues on which I want to focus. Plenty of references are provided for readers interested in digging deeper. Second, I present very few formal statistics, despite their ubiquity in ecological publications. I report a great many empirical results from the literature, but almost entirely in graphical form, allowing readers to see for themselves the patterns in the data. Interested readers can consult the original publications for $p$-values, slopes, $r^2$, AIC, and the like. Finally, I cannot claim to have cited the original paper(s) on all topics. My emphasis has been on communicating the ideas rather than tracing each of their histories to the origin, although I do dedicate a whole chapter to the history of ideas, and I hope I have managed to give credit to most of those papers considered “classics” by community ecologists.

### 1.1.3. Sources of Inspiration

By way of ensuring that I have appropriately credited the ideas that form the basic premise of this book, I end this introductory chapter by acknowledging those publications that inspired me by calling attention to the striking conceptual parallels between population genetics and community ecology (Antonovics 1976, Amarasekare 2000, Antonovics 2003, Holt 2005, Hu et al. 2006, Roughgarden 2009). Many additional researchers have taken notice of these parallels, especially following the importation into ecology of neutral theory from population genetics (Hubbell 2001). That said, I can say from experience that most community ecologists have *not* thought of things in this way, and there has been no systematic effort to find out whether it’s possible to reframe the bewildering number of theories, models, and ideas in community ecology as constituents of a more general theory involving only four high-level processes. This book is my attempt to do so.

# Index