Contents

CHAPTER 1

Introduction

A harrier glides over a field in search of voles, mice, and insects. Many of those rodents and insects themselves search for seeds while on vigil for the harrier. Other rodents, leafhoppers, grasshoppers, and caterpillars munch on the leaves and stems of the plants that produced those seeds. The plants search for mineral resources, water, and sunlight as they stand in place and suffer the consequences of losing tissue to these tormenters. However, the plants also offer up nectar and pollen for bees, butterflies, beetles, and birds as an enticement to induce pollination. Mycorrhizal fungi and the roots of the plants exchange water, nutrients, and carbohydrates. Spiders build webs in the plants' branches to ensnare various insects. Viruses and bacteria infect all these species to perpetuate their own existence. All of these interacting species form a biological community.

Current estimates suggest that 1.5–8.7 million species may currently inhabit the Earth's biosphere (May 1992, Gaston and Blackburn 2000, Mora et al. 2011). Any location where individuals of any one of those millions of species is found could be described similarly. That includes old-growth or second-growth forests, small creeks or large rivers, vernal ponds or large lakes, the abyssal ocean or intertidal coastline, or any city street or suburban neighborhood.

At every location, many species must be able to thrive with all the other species that are present there. Most have little direct effect on one another: a Cooper's hawk perched on the branch of a cherry tree probably has no direct effect on the demographic performance of that tree. The Cooper's hawk does have strong direct effects on the smaller birds and rodents on which it feeds. This reduction in the abundances of bird and rodent species reduces the food available for the other accipiters and buteos, snakes, foxes, weasels, and all the other carnivores that also feed on those species. These direct effects of the Cooper's hawk on bird and rodent abundances also indirectly benefit the cherry trees by reducing the abundances of the herbivorous rodents that feed on the cherry trees' seedling offspring, and indirectly harm the trees by reducing the abundances of birds that disperse their seeds. The bees directly benefit those cherry trees by pollinating their flowers, and in so doing indirectly benefit the caterpillars that feed on the leaves of those cherry trees. Those pollinators also indirectly harm the other trees and shrubs that compete with those cherry trees for sunlight and mineral resources and the understory plants that must deal with the allelochemicals that leach from the fallen cherry leaves and branches.

Each species at some location must deal either directly or indirectly with all the other species that are there. However, not all of those species are thriving at that site. Some species are present only because individuals continually immigrate from other areas. Some species are present now but will not be present sometime in the future. The individuals of some are vagrants. Only the species that can support populations in the face of dealing with all those other species can be considered to be thriving. These thriving species and how each deals with all the others to be able to sustain its local population are the foci of this book.

WHAT IS "COEXISTENCE" IN GENERAL TERMS?

For a community ecologist, saying that a species is "coexisting" in a community implies something much more than the colloquial meaning of the word. In general usage, coexistence simply means living together, which would suggest that all the species found in the same place would be coexisting. However, to a community ecologist, coexistence implies a very precise statement.

The word *coexistence* also typically connotes a comparative statement about a small group of species embedded in a larger community. For example, we typically might speak about two resource competitors as coexisting with one another. However, this narrow focus on only these two species in the community is too limiting a perspective. All the other species in the community must also be considered to make a definitive statement about each one's ability to coexist. Most directly, whether each of these two resource competitors will coexist depends on the dynamics of the resources over which they are competing and the dynamics of any predators, pathogens, or mutualists they may have. The dynamics of these other species will in turn depend on how they interact with other species in the community; for example, mutualists that the resources may have and additional prey that the predators may have. What determines the success or failure of each of these species are the direct impacts that the resource, predator, mutualist, and pathogen abundances have on its demographic performance. The other resource competitor only indirectly affects its competitor's success via its own impacts on these other species. Certainly, the presence or absence of a resource competitor may strongly influence the success of a particular species of interest, but this influence is mediated through the network of species interactions, and this resource competitor is just one node in that causal network. Moreover, these two resource competitors will not be able to coexist everywhere, because the conditions affecting all these other species will differ among locations.

Thus, the issue of coexistence is not a comparison of a small collection of species taken in isolation from all the other species in the community. Coexistence is a property of each species in the community. That property is its demographic performance in all the interactions with other species in the context of the abiotic environment in which it must engage. For me, the correct comparison is not limited to some subset of species, but rather pertains to whether this species can coexist with all the other species in the community.

In community ecology, the thriving species are coexisting. The simplest and most general definition of coexistence in the community ecology meaning of the word is the following:

> *A species is coexisting in a community if it can maintain a population in the local ecological conditions it experiences.*

This sentence is simple enough to seem like a platitude, but the devil is in the details, and those details are the subject of this book.

This definition of coexistence has three important issues. The first issue is that coexistence is a statement about a single species embedded in a community and ecosystem. That community defines the web of species interactions that take place at that location, and the ecosystem context defines the regime of abiotic conditions in which those species interactions take place. However, this issue must be evaluated separately for each species in the community. Therefore, coexistence must be considered simultaneously in both granular and holistic contexts from the perspective of each species individually and all species at once.

The second issue is the meaning of the clause "maintain a population." The colloquial meaning implies merely that the population exists. However, individuals of a species may be present at a site for many different reasons. If individuals are simply vagrants that are quickly passing through, one must question whether those individuals are persistent members of the community. As geese migrate south for the winter, they stop for a few hours to days at ponds, lakes, and rivers along the way. These stopovers may have important influences on the dynamics of the local community (e.g., defecation causing significant nutrient inputs: Manny et al. 1975, Olson et al. 2005). However, these geese do not "maintain a population" at these ponds. Questions about which species are coexisting are not the same as questions about which species influence the structure of a community: as we will see, species need not be coexisting to influence community structure.

A distinction is also drawn between species that maintain a population because of the balance of local per capita birth and death rates and those that would be locally extirpated were it not for continual immigration from some other area. Again, species maintained by immigration may have important effects on local community structure: individuals of such a species still consume resources, are themselves consumed by predators, interact with mutualistic partners, and support diseases. However, their populations are present not because they are relatively successful in the local ecological conditions, but rather because their immigration rate is higher than their local rate of population decline (Shmida and Ellner 1984, Holt 1985, Pulliam 1988).

Coexistence is limited to those species that can maintain a population because of the positive balance between their per capita birth and death rates determined by local ecological conditions. This focus traces through the ideas about a species' niche (Grinnell 1917, Elton 1927, Leibold 1995), in particular G. E. Hutchinson's set-theoretic definition of the niche. Hutchinson's initial statement (1958, p. 416) is this:

> Consider two independent environmental variables x_1 and x_2 which can be measured along ordinary rectangular coordinates. Let the limiting values permitting a species S_1 to survive and reproduce be respectively x'_1, x''_1 for x_1 and x'_2, x''_2 for x_2. An area is thus defined, each point of which corresponds to a possible environmental state permitting the species to exist indefinitely.

Thus, Hutchinson's statement of the niche is a statement about where a species' per capita birth and death rates permit it to "exist indefinitely." He defines niche axes as features of the environment and considers only competition with other species as being important—which has led to much confusion about whether other species interactions, such as predation, disease, and mutualisms, are part of a species' niche (e.g., see review by Leibold 1995). My point here is not to debate the definition of a niche (that will come later), but rather to highlight the distinction that the phrase "maintain a population" is limited to the balance of per capita birth and death processes caused by local ecological conditions.

The third issue is how to define and identify "the local ecological conditions it experiences." Given that the critical issue of coexistence is whether the per capita birth and death rates of the species foster maintaining the local population indefinitely, these local ecological conditions must be evaluated with respect to those demographic rates. Specifically, what local ecological conditions influence the values of these demographic rates and how do these rates change with the changing abundances of all the species in the community, including the species of interest? These demographic rates are functions of both the local abiotic environmental regime, as Hutchinson's (1958) niche definition identified, and interactions with all the species that exist there. These include all the abiotic resources, prey, predators, competitors, mutualists, and pathogens in the community; in other words, the entire ecological milieu that influences the demography of each species. To give primacy to one particular type of species is to ignore much of this causal structure.

The central focus of this book is to explore how these demographic rates for each species in the community change because of the abiotic environmental regime and species interactions with other community members. The set of important ecological conditions may be different for every species in the community. In fact, the differences in these sets among species are typically the defining mechanistic features that permit the coexistence of each (Hutchinson 1958, MacArthur and Levins 1967, Levin 1970, Chesson 2000b). This makes our current one-size-fits-all approach to understanding coexistence inappropriate.

INVASIBILITY

How do we assess whether a species can maintain a population in a local area indefinitely? As discussed above, the mere presence of the species does not justify such a conclusion. As Hutchinson's niche definition implies, this assessment must be based on the per capita birth and death rates that shape the population dynamics of the species at that location. For clarity, I outline this issue here purely from a population ecology perspective, ignoring how interactions with other species define these demographic rates. The rest of the book incorporates those species interactions into this framework.

Because no species has increased to infinite abundance, we can infer that the abundances of all species are regulated to some degree, or at least are bounded by their ecological surroundings. The essence of population regulation is the change in the constituent demographic rates that comprise per capita population growth rate with a change in the abundances of one's own species and those of other species. A species may have only indirect impacts on its own demographic rates through how the abundances of its prey (resource limitation) or predators (predator limitation) or mutualists (mutualist augmentation) or pathogens (pathogen limitation) change as its own abundance changes, but for now I ignore the mechanisms causing these demographic relationships to exist. When viewed from this purely population ecology perspective (i.e., ignoring interactions among species), the signature of population regulation is the change in these per capita demographic rates with a species' own abundance. This implies that either per capita birth rate decreases, per capita death rate increases, or both happen simultaneously as its own abundance increases.

These ideas can be formalized mathematically. A generic model describing the change in total population growth rate caused by these demographic processes for a generic species N is

$$\frac{dN}{dt} = N\left[b\left(N,\ldots\right) - d\left(N,\ldots\right)\right], \tag{1.1}$$

where $b(N, \ldots)$ and $d(N, \ldots)$ are the per capita birth and death rates, respectively, which are each functions of its own abundance (N), the abundances of all the other species in the community (the effects due to other species are represented by the ellipses because they are being ignored for now), and the abiotic environmental conditions in which these species interactions take place (typically defined by the structure and the parameter values of the function). Throughout this book, I use this continuous-time formulation of population change, but all the results presented here can also be derived in discrete-time formulations (Hassell 1978, Murdoch et al. 2003, Turchin 2003).

What are the logical implications of a population being able to "exist indefinitely," in Hutchinson's terms, given such population regulation? From a population dynamics perspective, one of three dynamical situations is implied. The first is that the community is being regulated to a stable point equilibrium in abundance, and so the population of species N is being regulated to a stable point equilibrium (Case 2000). An equilibrium abundance for N is the value where the population growth rate does not change (i.e., $dN/dt = 0$). Based on equation (1.1), that occurs when $N = 0$ or at the abundance, giving $b(N, \ldots) - d(N, \ldots) = 0$; in other words, that abundance at which per capita birth and death rates are equal (ignoring per capita immigration and emigration rates). Throughout the book, I identify such equilibrium abundances with a superscript star: N^*.

For this latter equilibrium with $N^* > 0$ to be stable, population abundance must move toward the equilibrium when it starts near the equilibrium. Thus, if the species' abundance is below this equilibrial value, the per capita birth rate of the species must be greater than its per capita death rate, and so its abundance will increase. In contrast, if its abundance is above this equilibrial value, its per capita birth rate is less than its per capita death rate, and so its abundance will decrease (fig. 1.1).

The second possible dynamical situation is that the community displays limit cycles or chaotic dynamics. These are dynamical features in which the community orbits a stationary equilibrium (i.e., an orbital attractor): the equilibrium point itself is unstable, but it is surrounded by a stable limit cycle or chaotic shell (Hirsch et al. 2012, Strogatz 2015). The cycling of hare and lynx abundances across Canada is one of the most famous examples in ecology (Elton 1942, Stenseth et al. 1997), but many other communities show analogous types of dynamics (Elton 1942, Turchin 2003, Korpimäki et al. 2004, Myers 2018). In this case, the environmental conditions remain constant, but the internal dynamics of the community cause the cycling (Krebs et al. 2013, 2018). In other words, the functions $b(N, \ldots)$ and $d(N, \ldots)$ do not change, and so the location of N^* does not change either. However, the dynamics caused by interacting with the ignored species create cycles in the abundance of N around N^*. Despite the continually changing demographic relationships, the same population regulation features hold if the

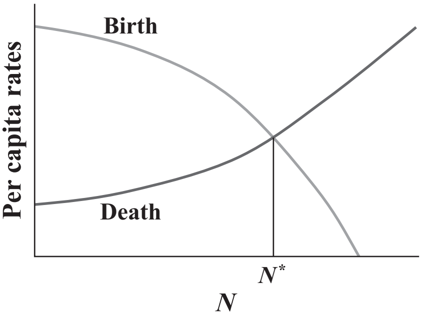

FIGURE 1.1. A stylized representation of the per capita birth and death rates for a population of some species. The birth and death rate curves are density dependent, meaning that their values change as the abundance of the population changes. The equilibrium abundance occurs at the abundance where birth and death rates are equal, and this equilibrium abundance is signified by N^*. The equilibrium pictured here is stable because per capita birth rate is greater than death rate at abundances lower than N^*, and death rate is greater than birth rate at abundances above N^*.

species is to be present: when the species reaches low abundance in the cycle, its per capita birth rate becomes greater than its per capita death rate so that it then increases in abundance; and when the species reaches high abundance in the cycle, it begins to decline because its per capita death rate becomes greater than its per capita birth rate.

The final possible dynamical situation is that temporal variability in the environmental conditions drive variability in these demographic rates for community members, which introduces variability to the dynamics of their interactions. Stochastic year-to-year weather variability can be a potent source of such demographic variability in many species (e.g., Leirs et al. 1997, Kausrud et al. 2008, Schmidt et al. 2018). One can think of this as continual change in the shapes of the per capita demographic rate relationships $b(N, \ldots)$ and $d(N, \ldots)$ caused by ecological variability through time (Tuljapurkar and Orzack 1980, Tuljapurkar 1989, Lande 1993, 2007, Lande et al. 2003). This generates temporal variation in the position of the equilibrium N^*, and this shifting position of the equilibrium drives population change. For a species to persist in this community, its long-term average per capita birth rate must be greater than its long-term average per capita death rate when at low abundance:

$$\overline{b\left(N \approx 0,\ldots\right)} > \overline{d\left(N \approx 0,\ldots\right)},$$

where the overbars signify averages (Turelli 1977, 1978b, Levins 1979, Chesson and Ellner 1989, Chesson 2000b, Schreiber et al. 2011).

All of these imply that a species will be present if its abundance on average increases when its abundance is below the equilibrium value. Taken to its extreme, this implies that the species' abundance will increase when it is extremely rare: that is, when its abundance is surely below $N^* > 0$. In fact, this is the criterion identifying coexistence for a species, which is commonly termed *invasibility* (MacArthur 1972, Holt 1977, Turelli 1978b, 1981, Chesson 2000b, Siepielski and McPeek 2010). This concept has also been referred to as *permanence* (Hutson and Law 1985, Hutson and Schmitt 1992, Law and Morton 1993, 1996, Morton and Law 1997), but I use invasibility throughout. Invasibility implies the success of the species in a number of contexts. For example, if some perturbation caused the species to be knocked to very low abundance (e.g., a disease outbreak left only a few surviving individuals of this species) or it cycles to low abundance, the species would be able to recover and increase in abundance. Alternatively, if it were initially absent, the species could successfully invade the community—hence *invasibility*. Also, note that invasibility is not evaluated with respect to any other particular species but with respect to the entire community. Thus, if invasibility is the criterion to evaluate coexistence, this argues that the concept of coexistence is a property of each species, given the community and ecosystem in which it finds itself.

A technical definition of invasibility is as follows:

> *A species satisfies its invasibility criterion if its per capita birth rate exceeds its per capita death rate when it is extremely rare and all other species in the community are at their demographic steady states in its absence.*

Expressed mathematically, this definition becomes

$$b\left(N \approx 0,\ldots\right) > d\left(N \approx 0,\ldots\right) \tag{1.2}$$

(see references in the previous paragraph). Each species that can satisfy this invasibility criterion is coexisting in the community: this is the community ecologist's meaning of coexistence. Obviously, a population with few individuals will be subject to the whims of demographic stochasticity that may cause its extinction even though its expected per capita population growth rate is positive (Adler and Drake 2008). Thus, not every invasion of a potentially coexisting species or return from very low abundance will be successful, and cycling species are at greater risk of extinction when near the nadir of their abundance cycle. However, satisfying this invasibility criterion is the hallmark of every coexisting species, at least theoretically.

Biological communities do not contain only coexisting species. As mentioned above, some species that are present at a location are not sustaining themselves because of local ecological conditions, and these species can be as important to local community structure as the coexisters. By identifying which species are coexisting in a community, the species that are present but not coexisting are also being identified, namely, *co-occurring species* (Leibold and McPeek 2006). In addition to the vagrants, three other types of co-occurring but not coexisting species are present in communities.

The first type of co-occurring but not coexisting species are what Daniel Janzen called *walking dead species*. They have also been called the extinction debt (Tilman et al. 1994) and the living dead (Hanski 1998). These are species that are headed to extinction but not there yet. Species headed to extinction do not suddenly just disappear. Extinction also has a temporal dynamic based on the differences in local and average regional birth and death rates (Lewontin and Cohen 1969, Turelli 1977, Raup 1992, McPeek 2007), and these extinctions may take a very long time. For example, two million years were required to drive 12 bryozoan species in the southwestern Caribbean basin extinct; their extinction was caused by the reduction in productivity precipitated by the closing of the Isthmus of Panama and the introduction of new and presumably more competitive species via speciation (O'Dea et al. 2007, O'Dea and Jackson 2009).

The second type of co-occurring but not coexisting species are *sink species*. Sink species have their local per capita death rate exceed their local per capita birth rate, but their local population is maintained by continual immigration from other local communities in the region where they are sustaining (i.e., source populations) (Shmida and Ellner 1984, Holt 1985, Pulliam 1988). Because sink species can only emerge as a result of movement between communities from source populations to sink populations, their existence and effects on community structure are considered in chapter 9 on spatial variability.

The final type of co-occurring but not coexisting species are *neutral species* (Hubbell 1979, 2001, Hubbell and Foster 1986). When present, neutral species are a guild of ecologically identical species. For example, in eastern North America, the 8–12 *Enallagma* species found in each lake with fish display all the hallmarks of neutral species (Siepielski et al. 2010). Because they are ecologically identical, interactions with other species in the community regulate the total abundance of all guild members instead of each guild member separately (Hubbell 2001, Siepielski et al. 2010, McPeek and Siepielski 2019). As a result, the abundance of each species may change through time following a random walk within this constraint of total abundance. Neutral species are technically not coexisting species, because if one guild member is already present in the community, any additional guild member would have its per capita birth rate equal to its per capita death rate when it invades. In effect, what is coexisting in the community is the entire guild of neutral species simultaneously, even though no single member of the guild is coexisting. Neutral species are considered in chapter 10.

THE BOUNDS OF A COMMUNITY

The boundaries of a community are difficult if not impossible to define. I have worked on the communities of organisms in ponds and lakes for most of my career. Community ecologists who work in terrestrial ecosystems generally think that we aquatic types have it made: "the community" in a lake clearly ends at the water's edge, and "the community" is completely self-contained in the water of that lake.

However, even in something as discrete as a lake, multiple communities may exist, and the boundaries between them are not clear. The assemblage of crustaceans, insects, annelids, mollusks, salamanders, and fish that live clinging to and growing around the macrophytes in the littoral zone of a lake is very different from the assemblage of crustaceans, insects, and fish that live suspended in the pelagic open water zone, and the assemblage of annelids, crustaceans, insects, and mollusks that live in the mud of the benthic lake bottom zone is different still. The populations of some species span across the boundaries of these zones, but many taxa are restricted to only one. Because interactions among individuals of various species define the dynamics of a community, each of these zones might be considered separate communities of interacting organisms. Even the assemblage of organisms in each of these "zones" may not represent a single community. For example, if the macrophyte beds of a lake are disjointed, each bed may represent a separate dynamical unit with little movement of individuals between them. Individuals of many species are restricted to only one of these communities within a lake for much or all of their life cycle.

However, each of these lake zones is not an integrated and wholly separate community. Individuals of some species move between these various communities within the lake. Zooplankton migrate between the epilimnetic and hypolimnetic zones of the pelagic over the course of a day; some fish species forage in these different areas on a daily basis; some fish undergo ontogenetic habitat shifts so that they spend one life stage in one zone and another life stage primarily in another. Moreover, the water's edge is no discrete boundary. Individuals of many species, such as ducks, geese, herons, gulls, and moose, routinely forage in multiple lakes and so link the dynamics of multiple lakes. Likewise, some fish species use the stream connections to move between lakes. Does the fact that the populations of some species span these zones unite them into a single community, or are these separate communities that are linked by dispersal? Where is the line between two populations with limited dispersal versus a single continuous population? Is trying to draw a community boundary even useful if the populations of different species operate on such disparate spatial scales?

The boundaries of communities in terrestrial ecosystems are similarly obscure. The assemblage of species found in a light gap in the forest is different from that found in a patch of the continuous forest surrounding it. Some species are common to both and some are unique to each. Some species forage across multiple light gaps, while others may spend their entire lifetimes in only one patch. In the surrounding forest, the individuals of the various species that are interacting primarily with one another may be generally found over the expanse of a few hectares to square kilometers, but the forest may continuously extend over many thousands of square kilometers. Where is the boundary of this community to be drawn? Exactly the same issues emerge, but on a much smaller spatial scale, when you consider the microbiomes found in various parts of the gastrointestinal tract, the respiratory tract, and skin on your own body.

Defining the boundaries of a community or communities that develop along abiotic gradients also poses conceptual, empirical, and philosophical challenges. The plant assemblages change in a regular patterning along the altitudinal gradient on the mountains of the southern Appalachians. In the deep ravines cut by streams at the bottom of the mountain, you will find cove forests dominated by basswood, poplar, magnolias, and birches, while the dry mountaintops are dominated by oaks, hickories, and beeches (Braun 1935, 1940, Whittaker 1956). Some species can be found over much or all of this gradient, while others are restricted to particular ranges. Even salamanders that may move only a few meters in their lifetimes show distributional patterning along these altitudinal gradients (Hairston 1951, 1980, Jaeger 1971). The same issues arise over centimeters in the patterning of species distributions and community structure on the rocks of the marine intertidal (Connell 1961, Paine 1966, 1969, Menge 1976, Sousa 1979). Should the boundary of a community be placed at the distributional limits of each segregating species along the gradient or at the limits of only key species, or do the limits of the most widely distributed species set the community's boundaries?

If a biological community is not a defined and integrated unit that has no discernible boundaries, is the coexistence of species in a community something that can or should be studied? In fact, these are exactly the issues that led Ricklefs (2008, pp. 741–742) to argue that "community" is an "epiphenomenon that has relatively little explanatory power," and "coexistence can be understood only in terms of the distributions of species within entire regions" (see also Gleason 1926). If all the species present at a location are simply a sampling of all the species in the surrounding region, with each species having little to no demographic impact on any other, changes in species composition that occur as you hike up a mountain or swim from the littoral to pelagic zones of a lake would simply reflect the physiological tolerances of those various species to the local abiotic conditions.

While I completely agree with Ricklefs's arguments about the importance of regional processes and macroevolutionary dynamics in shaping local species assemblages (Ricklefs 1987, 1989, 2008, 2010) and his arguments have inspired my own work on how macroevolutionary processes shape current species diversity and the phylogenetic patterns of component taxa (McPeek and Brown 2000, 2007, McPeek 2007, 2008b), I completely disagree with the deduction that little if anything can be learned by studying coexistence in the framework of community ecology. The literature is replete with examples where adding or removing one species causes others to go locally extinct and thus permits others to locally thrive that were not previously there. Remove *Balanus balanoides* barnacles from rocks in the intertidal of Scotland, and *Chthamalus stellatus* barnacles will colonize and thrive (Connell 1961). Remove *Pisaster ochraceus* seastars from rocks in the intertidal of Washington, and as many as 25 species of barnacles, sponges, anemones, snails, chitins, urchins, and algae are replaced by a single species, *Mytilus californianus* mussels (Paine 1974). Introduce largemouth bass, *Micropterus salmoides*, to a lake in Wisconsin, and the species composition of insects, cladocerans, copepods, rotifers, and algae throughout the pelagic food web shifts dramatically (Carpenter and Kitchell 1993). This large body of results proves that interactions among species on a local scale shape not only community membership and the local abundances of those species but also larger patterns of species composition and diversity.

For me, community ecology is the study of the network of species interactions across local and regional scales. A community is an assemblage of species connected in a network of species interactions (Paine 1980, Martinez 1991, 1992, Dunne et al. 2002, Bascompte and Melián 2005). If one species directly impacts the demographic rates of another species, those two species are directly linked in this interaction network. Indirect species interactions include the effects of one species on another that passes through one or more intermediate species in the network because of changes in abundances (i.e., abundance-mediated indirect effect: Abrams 1995), or the effect that one species has on the strength of the interaction between two other species (i.e., trait-mediated indirect effect: Abrams 1995, Werner and Peacor 2003). What is included in any particular study depends on the taxa and the scope of the interaction network that are the primary foci of the researcher's interest, the linkages of those taxa within the local interaction network, and the linkages of networks among locations.

Invasibility is then a test of whether the species of interest is persistent in the local network because of the network's structure and its own local demographic success. If the species passes the test, it is a persistent member of the community. If it fails, it may be a redundant community member (i.e., neutral species), it will disappear from the network (i.e., walking dead species), or it exists primarily because of linkages to networks in other locations (i.e., sink species).

Community ecologists can ask a multitude of different questions of various scopes about these interaction networks, and different questions will require the researcher to consider the interaction network at different spatial and temporal scales. Most community ecologists are initially interested in the role of particular taxa or assemblages or trophic levels in the community, and so all questions begin with these species as the focus. Initially, only the components of the interaction network that directly impact their demographic rates and abundances will be important. Thus, for a community ecologist interested in understanding some aspects of the community ecology of *Daphnia* species found in the pelagic zone of a lake, the components of this interaction network that directly affect *Daphnia* species' demographic rates and abundances are most relevant; that is, the algal species on which they feed and the predators that feed on them. Depending on the questions being addressed, it may also be important to understand more distant parts of the interaction network that influence the abundances of these species. And so, typically, one will start with some specific component of the web and work out through various connections, because components of the interaction web that do not directly or indirectly influence the components of interest are irrelevant to study.

The extent of the network of interactions that one wants to or is forced to study is primarily determined by the questions being addressed by the researcher. This is because species composition in one part of a local food web intimately depends on the species that occupy other parts of the food web. In fact, a food web represents much (but not all) of the causal network of species interactions that constitutes a biological community. For some taxa and some types of questions, only adjacent portions of the local interaction network are necessary. For others, linkages between the networks in multiple locations are critical (Holyoak et al. 2005, Leibold and Chase 2017). I have organized this book to systematically work through the exploration of these interaction networks and how various features of the networks foster or constrain various features of communities (i.e., opportunities for new types of species to establish, consequences of species perturbations or deletions, patterns of species distributions, and abundances among communities scattered along environmental gradients).

THE WHY QUESTION

Understanding this network of species interactions is also key to understanding *why* each species is successful or unsuccessful in a community, namely, the mechanism for why it can coexist. Testing whether a species can invade a community is in principle a fairly simple exercise (but exceedingly difficult in practice for most communities), and such tests can be accomplished without any reference to why the species has a positive per capita population growth rate when it is rare. However, each coexisting species in a community is present because of its abilities to balance the conflicting demographic demands of the various species interactions it faces at its position in the food web. Understanding comes from not simply asking whether this or that species is coexisting; rather, understanding comes from asking why this species is able to coexist in this community.

One of the main problems I see today with community ecology is that we have largely gotten away from asking the why questions at a mechanistic level. This is not a new problem: Tilman (1987) actually made exactly this same point specifically about the study of resource competition; Schoener (1986) made a similar argument without focusing on a broader range of interactions, and he sketched the philosophical designs of a mechanistic research program. This book is an expanded argument for their perspectives. However, since the 1950s, coexistence has been explored primarily through the prism of a general theory of competition that is devoid of any mechanism (Gause 1934, Gause and Witt 1935, Slobodkin 1961, May 1973, Vandermeer 1975, Chesson 2000b). In fact, Chesson (2000b, p. 345) states explicitly that “these models are phenomenological: They are not defined by a mechanism of competition.” The central question in that approach is, What is the effect one species has on another relative to the effect a species has on itself (Gause and Witt 1935, Slobodkin 1961, Chesson 2000b)? However, these effects are characterized devoid of any causes for those effects. As we explore in chapter 2, the justification for this analysis comes from the ability to transform simple models of two types of species interactions into a common form that subsumes the causal chain of interactions across multiple species into what appears to be a single direct interaction between two species (MacArthur 1969, 1970, Schoener 1974b, Abrams 1975, Chesson 1990, Chesson and Kuang 2008). These two types of interactions are (1) two or more consumers competing for resources and (2) two or more prey that are fed upon by the same predators. In fact, we use the language of “competition” to describe both of these interactions, namely, resource competition and apparent competition, respectively.

Although important conceptual insights have been gained and I am sure will continue to emerge from this theory, I think that relying solely on general competition theory to understand species coexistence is problematic for a number of reasons, which I explore in greater depth in chapters 2 and 11. My main reason for concern has both philosophical and practical bases. If you compare and contrast resource competition with apparent competition, you realize that the abilities and phenotypic properties that make a species successful at utilizing some resource are completely different from the abilities and phenotypic properties that make a species successful at reducing mortality from some predator—in fact, they are often in antagonistic opposition! However, when couched in the language of competition, most minds immediately gravitate to the interpretation of resource competition, which in turn focuses the mind back to this specific mechanism.

Answer these questions honestly to yourself. If I say that two species differ in competitive ability, what do you think of? Do their relative abilities to avoid a common set of predators enter your conception?

> *The words, language, and models we use to conceptualize nature define how we see nature and, consequently, how we measure and study nature.*

As an example, consider Hutchinson's (1961) famous *paradox of the plankton*. Ponds and lakes of moderate size support dozens to hundreds of phytoplankton species (Smith et al. 2005). For example, 84 phytoplankton species were recorded in Trout Lake, Wisconsin, USA, from samples taken in 2005–2006 (data downloaded 22 August 2019 from https://portal.edirepository.org/nis/mapbrowse?scope=knb-lter-ntl&identifier=238). Hutchinson's paradox explains how so many phytoplankton species can all coexist on just a handful of limiting resources (he listed light, CO_2, phosphorus, nitrogen, and sulfur compounds, plus 14 elements), given that "natural waters, at least in the summer, present an environment of striking nutrient deficiency, so that competition is likely to be extremely severe" (Hutchinson 1961, p. 137). He offered a number of possible hypotheses for how so many species could be found together and possibly coexist. The two hypotheses that have resonated over the last 60 years or so are that either temporal variability never permits the system to reach an equilibrium and so competitive exclusion is never completed (here the species are not coexisting but simply co-occurring because they are on long transients caused by this variation: Chesson and Huntly 1997), or species really do coexist because the spatial and temporal ecological variation causes continual reversals and shuffling of the competitive dominance of species. Theoretical explorations attempting to explain the paradox of the plankton primarily approach the problem by considering such spatiotemporal variation in nutrient limitation and competitive dominance reversals (depending on the model; Armstrong and McGehee 1980, Chesson 1994, Huisman and Weissing 1999, 2001b, Huisman et al. 2001, Abrams and Holt 2002, Klausmeier 2010, Li and Chesson 2016).

However, competition for limiting nutrients is not the only ecological process that shapes the demographic rates of phytoplankton species. In fact, Hutchinson (1961) discusses two other types of species interactions, although the fact that he did so has been largely forgotten. (Note to graduate students: You actually need to carefully read classic papers yourself. Do not take other people's word for what papers say.) He devotes a paragraph to symbiosis and commensalism and offers some conjectures about how these may foster the coexistence of some species: some may benefit from acquiring vitamins that are released by others into the water column. Hutchinson (1961, p. 141) then spends exactly two sentences to suggest that predation may also foster the coexistence of some species:

> It can be shown theoretically, as Dr. MacArthur and I have developed in conversation, that if one of two competing species is limited by a predator, while the other is either not so limited or is fed on by a different predator, co-existence of the two prey species may in some cases be possible. This should permit some diversification of both prey and predator in a homogeneous habit.

To my knowledge, MacArthur and Hutchinson never published these insights. We explore these conjectures explicitly in chapter 4 (see also Grover 1994).

Predation by herbivorous and omnivorous zooplankton species are actually a very likely additional set of species interactions that may foster the coexistence of many different phytoplankton species (e.g., Leibold 1989, Grover 1994, Leibold et al. 2017). For example, Trout Lake has 68 species of herbivorous or omnivorous copepods, cladocerans, and rotifers that feed on those 84 phytoplankton species (data from https://portal.edirepository.org/nis/mapbrowse?scope=knb-lter-ntl&identifier=37). As we explore in the coming chapters, enemies (i.e., herbivores, predators, omnivores, pathogens) also serve as limiting factors that foster coexistence just as limiting resources do. In fact, these 68 phytoplankton enemies plus the 19 abiotic nutrients that Hutchinson listed give a total of 87 potentially limiting factors, which are more than enough to account for 84 phytoplankton species (see Levin 1970). These numbers do not include the many herbivorous and omnivorous protozoan species that also feed on phytoplankton and the many viruses that also attack phytoplankton, which have not been enumerated in surveys of Trout Lake to my knowledge.

Thus, with great deference and all due respect to Hutchinson (whom I revere), I see no "paradox" at all, given the number and diversity of nutrients and the number and diversity of enemy species that potentially influence the per capita birth and death rates of each species in the phytoplankton assemblage of a lake. That is not to discount the contributions of spatial and temporal variation in these factors at all; we explore their contributions to fostering coexistence in chapters 8 and 9. Rather, the point of this book is to begin to synthesize a robust framework presented by dozens of theoretical ecologists over the past six decades, which simultaneously incorporates this diversity of factors that can contribute to the coexistence of the various species found in a community. The paradox only exists if we limit ourselves to considering competition for limiting resources exclusively. From what I see, community ecology's focus on a general theory of "competition" to explain and understand coexistence of every species biases us to think only in terms of limiting resources and so blinkers us from exploring the full gamut of demographic processes and limiting factors that can easily account for a diversity of species to be successful members of a community.

This focus on a general competition theory also orphans other types of species interactions outside the mainstream framework for how we understand coexistence. In its generic definition, competition is a (−,−) interaction between two species, meaning that each species has a negative impact on the per capita growth rate and abundance of the other. Both resource competition and apparent competition are (−,−) indirect interactions that are mediated through other species. Because predation is a general (+,−) direct interaction (the two species directly interact, and the interaction has a positive effect on the predator but a negative effect on the prey), a general theory of "competition" to explain coexistence is completely agnostic in explaining the coexistence of a predator and its prey: predator and prey do not coexist because each limits its own abundance more than it limits the other species' abundance. If you want a real paradox, this is it: the direct interactions in the chains of causation that generate the indirect effects of both resource and apparent competition are all predator-prey interactions, and so this competition theory cannot explain the interactions that underlie what generates the competitive effects. The effects of pathogens, parasites, and parasitoids are similarly ignored, given that they influence their hosts in similar ways to predators influencing their prey.

Likewise, mutualisms have been largely orphaned outside the general competition theory. The general definition of a mutualism is a (+,+) interaction between two species, either direct or indirect. General competition models can be converted into models of mutualisms by simply making competition coefficients positive instead of negative (e.g., Gause and Witt 1935, Vandermeer and Boucher 1978, Goh 1979, Dean 1983, Boucher 1985). However, they are subject to the "orgy of mutual benefaction" where both species increase without bounds (May 1973), and they are still only caricatures that lack any true mechanism of interaction among mutualists, just like the competition forms are.

Moreover, we need a framework for studying coexistence that explores how mutualisms alter the results of other types of interactions in which species engage. For example, imagine a field filled with plant species. From the general competition theory perspective, the natural question to ask is, How do the interspecific competitive effects between species relative to their intraspecific effects on themselves promote their coexistence? All of these plant species also have suites of enemies that depress their abundances (e.g., Reader 1992, Carson and Root 1999, 2000). Many of them also have pollinators that increase their seed yield and mycorrhizal and rhizobial symbionts that provide them with extra nutrients to grow larger (reviewed in Bronstein 2015a). These symbionts inflate the individual sizes and population abundances of their plant partners, and so differential benefits from their various mutualistic partners may decidedly shift the balance of competition among the plants. Such mutualistic effects are outside the bounds of the general competition theory framework. However, models that better capture the mechanisms of interactions among multiple species in a food web simultaneously can explore such scenarios. We do so in chapter 6.

Answering the why question is the essence of the inquiry into the mechanism of causation (Pearl and Mackenzie 2018). Testing the invasibility of a species is a necessary step to definitively evaluate whether that species is coexisting in the community. In practice, invasibility tests are almost never done because they are simply logistically impossible to perform correctly in real communities (Siepielski and McPeek 2010). However, models of community structure that allow us to probe the reasons why a species would be successful at invading can identify definitive predictions that should then be empirically tested to validate the operation of those mechanisms in promoting or retarding that species' coexistence. This is one reason why working with models that capture the essence of mechanisms is so crucial. Moreover, even if an invasibility test can be performed in a real community, the test only answers whether inequality (1.2) is true. Why it is or is not true must still be left to these same studies addressing the existence and operation of various causal mechanisms.

ASSEMBLING A MULTITROPHIC-LEVEL COMMUNITY

Analyzing a model of a causal network of interactions among multiple species at multiple trophic levels can be a very daunting task. Because of this, I have organized this book to follow a systematic approach. In this multitrophic-level food web, some of the species had to invade before others, which sets up a natural order to the assembly of a community. The algae, bacteria, and protozoa that form the foundation of the green and brown food webs of the water column in a lake must invade and become established before any of the species that base their existence on consuming them can invade. A rotifer species such as *Keratella cochlearis* cannot invade and coexist in the lake unless a resource is available to support its population. In turn, species like the predatory copepod *Mesocyclops edax* cannot invade and coexist until species such as *Keratella* establish. Thus, a logical ordering of species invasions exists. Moreover, as we will see for many of these invasions, the same kinds of criteria for invasibility emerge over and over for very different types of species at very different food web positions.

Consequently, I have structured the book around the idea of starting with an empty ecosystem that is devoid of all biological species and adding one species at a time until the multitrophic-level community is complete. In effect, this approach is organized around the idea of building community modules (Holt 1997). A *community module* is a network configuration of interacting species that denote "multispecies extensions of pair-wise interactions" (Holt 1997, p. 333). A consumer feeding on a resource is the simple pairwise interaction that forms the backbone of the community module approach. Once the criteria for each of them to coexist in the ecosystem of interest are established, we can then ask what is required for a second consumer to invade to establish a "resource competition" module, or for a second resource to invade to establish an "apparent competition" module, or for a predator of the consumer to invade and thus establish a three-trophic-level "food chain" module.

For each invading species, the criteria necessary for it to be able to invade when rare will be the first concern. Once established, each species will potentially affect the other species already present, either by altering their abundances or actually being driven extinct because of its invasion. These changes in abundance of the species already present will in turn alter the criteria that subsequent invaders must satisfy, and so how the invasion of one species affects the abilities of other species to invade is also a central issue of inquiry here.

This approach also focuses attention explicitly on the theoretical predictions that form the basis for empirical testing. What are the species interactions that directly impact the demographic rates of the invading species? How are the abundances of those species determined by the chain of indirect interactions with other species in the existing food web? After each species invasion, I will briefly discuss the critical empirical tests that are needed to both evaluate invasibility and identify signature features of the mechanisms that foster or retard that invasion.

The mathematical backbone of this analysis is a variation on the classic Rosenzweig-MacArthur model of consumer-resource interactions (Rosenzweig and MacArthur 1963, Rosenzweig 1969, 1971). This is a versatile model that can capture many of the mechanistic features of pairwise species interactions for the full gamut of mechanisms needed here. The mathematical exposition of the model is postponed until chapter 3, but I describe the general features of the model here. The basic model assumes that the resource species grows according to logistic population growth in the absence of the consumer, and the consumer has a linear or saturating functional response for feeding on the resource (Rosenzweig and MacArthur 1963, Rosenzweig 1969, 1971). These features cause this basic model to display either a stable point equilibrium or a limit cycle, depending on parameter values. Also, if a predator is added to produce a three-trophic-level food web, the community may display chaotic dynamics in addition to a point equilibrium or a limit cycle (Hastings and Powell 1991, McCann and Yodzis 1994). Thus, the model can display the full range of dynamical features we need for this analysis.

Additional mechanistic features can be easily added to the model to explore how they modify outcomes between consumer and resource (Murdoch and Oaten 1975). For present purposes, I limit this to exploring two different ways to incorporate direct self-limitation for consumers and predators into the model. Self-limitation can arise when species display such features as being cannibalistic or territorial, or interfering with the feeding of conspecifics (Beddington 1975, DeAngelis et al. 1975, Gatto 1991, Fryxell et al. 1999, Amarasekare 2002, McPeek 2012, 2014).

Finally, the Rosenzweig-MacArthur consumer-resource model is also the basis for more mechanistic models of interactions among mutualistic partners (Holland and DeAngelis 2009, 2010). The mechanism generating the benefits to partners in many mutualisms is based on trade, where each species allows its partner to consume some material that the other has in return for some benefit (Wyatt et al. 2014, Bronstein 2015b). Plants offer up nectar and pollen for pollinators to consume in return for pollination services. Other mutualistic interactions, such as between plants and mycorrhizal fungi, plants and rhizobial bacteria, or corals and zooxanthellae, involve each partner consuming some material from the other. Thus, formulation of mutualistic interactions using this consumer-resource model fundamentally characterizes their mechanism of interaction (Holland et al. 2005, Holland and DeAngelis 2010).

Of course, all models are mechanistic abstractions to some degree (Levins 1966). Models can either capture the mechanism directly (e.g., a saturating functional response) or mimic the patterns of the consequences of some mechanistic component (e.g., logistic population growth, various forms of direct consumer self-limitation) (Fryxell et al. 1999). Also, often the exact mathematical function used is not important to the outcome but rather only the functional shape: for example, various functions can be used to model a saturating functional response and all produce the same outcomes (Seo and Wolkowicz 2018). Thus, I believe that the model results presented here adequately capture major qualitative and quantitative features of coexistence mechanisms.

I am no mathematician, but over the years I have trained myself to understand the models presented here. This training came mainly from teaching an undergraduate course in community ecology where these models are the central focus. Students are typically apprehensive about their math skills. On the first day, I ask them two questions: “How many of you passed seventh-grade algebra?” and “How many of you understand what a derivative is from the first semester of calculus?” I tell them that if they can answer yes to both of these questions, they are prepared for the course. The same is true for this book.

In fact, most of what is required is to be able to draw and interpret a special kind of graph. The workhorse of our analysis will be graphs of the isoclines that result from the dynamical equations. An *isocline* is a function that maps out when a species’ population growth rate is zero based on the abundances of all the species in the community. Thus, a graph of this function tells you the combinations of all species abundances where the abundance of the species in question is not changing. Moreover, if the combination of all abundances is not on the isocline, its position relative to the isocline tells you whether the species' abundance will increase or decrease. Thus, the overall dynamics of the community can be deduced from the graph of the isoclines of the interacting species (Hirsch et al. 2012, Strogatz 2015).

Biological mechanism is also embedded in the geometry of the isoclines. Each mechanistic feature of a species interaction (e.g., logistic growth, linear or saturating functional response, feeding interference, direct self-limitation) imparts a specific shape to the isoclines of the species engaged in those interactions. Also, the dynamical responses of species to one another are manifest in the geometrical relationships among the isoclines of those various species. For example, whether a consumer can invade and sustain a population on some resource species can be determined directly from the relative positions of their isoclines. A major focus for the reader should be to identify common features associated with invasibility as more species at different trophic levels are added. Increasing the dimensionality of the system by adding more mechanistic features to the model of each species and simply adding more species quickly make the mathematics quite complex and analytically intractable, particularly for someone of my modest mathematical skills. However, teaching experience tells me that understanding can emerge from considering the geometry of the isoclines in these much more complicated multispecies problems.

If this sounds like a lot of ground to cover, it is. Therefore, I cannot cover a number of important topics related to species coexistence in communities in this book. The importance of age-, size-, and stage-structure within species to promoting their own coexistence and fostering or retarding the coexistence of other species is not addressed here (e.g., Mylius et al. 2001, Rudolf and Lafferty 2011, de Roos and Persson 2013, Wollrab et al. 2013). I also do not address how adaptive behavioral shifts of species in response to one another or the resulting trait-mediated indirect effects can affect coexistence (e.g., Abrams 1992, 2010, Křivan 1998, 2000, Werner and Peacor 2003, Křivan and Schmitz 2004, Valdovinos et al. 2013, Bachelot and Lee 2018). Finally, species cannot evolve in response to one another (e.g., Slatkin 1980, Lande 1982, Taper and Case 1985, Abrams et al. 1993, Abrams and Chen 2002, McPeek 2017b, 2019a). The models I use here to explore coexistence are all easily extended to explore these issues (see the references cited here), so a thorough exploration of coexistence here sets a strong foundation for exploring these added complexities.

THEORY AND EMPIRICAL TESTING

Models are fundamental to all aspects of the scientific endeavor. However, models are constructed for many different purposes, and we must be clear what our purpose is for using any particular type of model (Levins 1966).

A *model* is a formal statement about the relationships that are deemed important among some set of variables. Some scientific models are purely statistical associations that describe covariance structures among a set of relevant variables. For example, quantitative genetic models of phenotypic associations describe the statistical associations among phenotypes based on genetic principles, but these models do not describe the causal mechanisms that map genic, genomic and developmental interactions into the phenotype. Likewise, models using artificial intelligence with big data are currently just mining statistical relationships embedded in the data trove (Pearl and Mackenzie 2018). If the covariance structures remain stable, these types of models can be excellent at making very precise predictions within the bounds of the data that are used to construct them. Yet, no matter how good its predictive powers are, this type of model is opaque to the causal mechanisms that generate the covariance structure (Woodward 2010). This is the basis of the aphorism “correlation is not causation.”

Other models are derived specifically to capture relevant features of some causal network. These models are derived by first defining the causal network and stating a set of assumptions about the interactions in the network. The causal network and assumptions thereabout are meant to represent some feature of nature under study. A mathematical description of the network and the assumptions are then derived and analyzed to generate descriptors of the dynamics generated by those assumptions. Those model dynamics serve as the basis to generate predictions that can be tested in the real world to examine how well the causal model represents the feature of nature under study. If the predictions are supported by empirical testing, one can then have some level of confidence that the causal model captures important features of the real system under study. Discrepancies will require refinement of the model and then subsequent empirical testing—which is the nature of science. The models discussed in this book fall into this category.

Although the feedback loop defined by causal model development and empirical testing is the basis of the scientific method, theoreticians and empirical scientists often have a hard time communicating with one another. The gulf may not be as wide for community ecologists as it is in other scientific disciplines, but the gulf does exist. As an empirical community ecologist myself, I have tried to write this book with the empirical community ecologist in mind. However, this is a book of mathematical models, so there is no way around *doing the math*.

First, I have tried to present the models at a level that, if you exercise your seventh-grade algebra, basic calculus, and graphing skills, you will be able to rederive everything I present here. And I strongly encourage empiricists to do just that! Read this book with a pad of paper and a pencil, and do the algebra to get from model to result and isocline picture. All the insight is embedded here. You will learn that repeated issues arise in these models, and these analyses become easier with the practice of doing the analysis alongside the book as a guide.

I have also emphasized only the mathematical features that are critical for developing biological insights. Therefore, only a few exact formulas for things like equilibria are given. Instead, I simply refer to the equilibrium where it exists, namely, where the isoclines cross at a single point. To extract biological insights, that is all that's required.

Because I am encouraging understanding through graphs of isoclines, visualization of these dynamical graphs is crucial. To aid the reader, I have also constructed computer simulators for many of the community modules to be explored. Simulation modules considered in this book can be found at https://mechanismsofcoexistence.com/. These simulators allow the user to examine the geometry of the isoclines for multiple parameter combinations and show the dynamics of species abundances given those geometries. Links to the corresponding simulators are found at appropriate places throughout the book.

Finally, to guide the reader in developing empirical tests of these models, each major section also contains brief expositions of what I see as some of the critical predictions and the empirical data, both observational and experimental, that would test those predictions. These are not exhaustive lists of the predictions or of the studies that could or should be done, but rather are obvious and major features that must be tested to evaluate whether the model adequately captures a community dynamic and the mechanism permitting species coexistence. Empirical research is an exercise in creativity, and so no exposition I could provide should be taken as a definitive statement about how to approach any particular model prediction. I offer these only as guides and starting points for in-depth analyses. I hope they are suggestive and inclusive of the best tests. In addition, I hope this will exercise the skills of empirical ecologists to extract the critical predictions from models. At least my students have told me this is one benefit from the approach I have taken here.

Index