From trait patterns to species lifetimes: Effects of niche differentiation on coexistence and community structure

by

Rafael D'Andrea Rocha

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Ecology and Evolutionary Biology)
in The University of Michigan
2016

Doctoral Committee:

Associate Professor Annette M. Ostling, Chair
Professor Bradley Cardinale
Professor Deborah E. Goldberg
Professor Mercedes Pascual

Acknowledgements

Annette Ostling, who tirelessly coached and supported me as a PhD student. Annette made a point to be always available for feedback on my work, even when she had tons of other things to do, which was all the time. She pushed me on when I faltered but was always friendly and never talked down to me. We collaborated closely—and intensely—on the papers we co-wrote. Her academic and financial support made possible my participation at numerous events that advanced my PhD career, including nine conferences, two courses and workshops, and two visits during her sabbatical year. She hired me twice as a GSI for her General Ecology class, which was an important learning experience, and also great fun. She introduced me to other ecologists, including James O'Dwyer, whose lab I will be joining after completing my PhD. Besides being there as my mentor and advisor, Annette doubled down as a friend and confidante. I am immensely grateful for the considerable amount of time and effort Annette dedicated to helping me during my five or six years at EEB. I would not be here without her.

Mercedes Pascual, Deborah Goldberg, and Brad Cardinale, who generously agreed to join my committee and provided valuable guidance to improve the dissertation.

Gyuri Barabás, my colleague, role model, collaborator, and good friend. Gyuri is an inspiration: devoted to science, incredibly smart, humble, and a heck of a great guy! I was lucky to join Annette's lab when Gyuri was there.

Rosalyn Rael, who was a postdoc in the lab when I joined. We collaborated on papers, enjoyed Indian food together, and had lots of laughs as ESA travel companions.

John Vandermeer, for believing in my application to EEB despite my scant background in biology.

Douglas Futuyma, whose Evolutionary Ecology course at Stony Brook started it all, and who suggested U of M.

Jane Sullivan and Cynthia Carl, EEB Grad Coordinators who made my PhD life so much easier.

NSF, EEB, and Rackham, for the funding.

Table of Contents

List of Figures

List of Appendices

Abstract

From trait patterns to species lifetimes:
Effects of niche differentiation on coexistence and community structure
by
Rafael D'Andrea Rocha

Chair: Annette M. Ostling

One of the most enduring questions in ecology is what accounts for coexistence among trophically similar organisms. Niche differentiation is one answer, but so is neutrality: species can coexist either because of ecological differences or ecological similarities. Despite being diametrically opposite, these two theories can be difficult to separate in nature. Because neutral theory posits that species traits are irrelevant for ecological performance, trait patterns are commonly used in niche inference, but results are mixed. This dissertation argues that widely accepted ideas about trait pattern driven by niche differentiation must be updated in light of recent findings, and takes the first steps in that direction. We contend that the current theory of trait patterning is incomplete, and progress requires exploring patterns across a variety of niche models. It emerges from this exploration that stochastic niche dynamics may result in the spontaneous formation of species clusters, under the qualification that pattern will not form if idiosyncratic regulation mechanisms allow arbitrarily similar species to maintain distinct ecological strategies. We provide a new metric for identifying and quantifying species clusters, which outperforms existing metrics in rejecting neutrality in our pseudodata. Another major theme is that the null hypothesis is critical in inference tests. Process-based null models are superior to statistical null hypotheses based on randomization of data because the latter destroys pattern caused by forces unrelated to niche differentiation. For example, we show that when species can randomly mutate into similar species, clusters may occur even in the absence of niche differentiation. A final key theme is that the effect of niche differentiation on species

dynamics and pattern may be more complex than is currently appreciated when stochastic dynamics under immigration is considered. Species richness, lifetimes, and extinction rates in niche-differentiated assemblages may be lower or higher than neutral assemblages, depending on the ratio between regional diversity and the number of niches available. The findings of this dissertation contribute to our theoretical understanding of niche differentiation as an important coexistence mechanism, and to realizing the potential of trait patterns in assessing its prevalence in nature.

Chapter 1

Introduction

1.1 Background

Ecology is not a new science, and yet it is still challenged by some of the same fundamental questions of its infancy. Among them is explaining the great levels of biodiversity observed in certain ecosystems such as tropical forests. The easy answer is that local biodiversity at a certain place and time results from a balance between species introductions via speciation and immigration, and extirpation via drift, overexploitation, and competitive exclusion. The hard part is explaining why the balance should occur, as it is entirely within the realm of statistical probability that it may not. Obviously gains cannot exceed losses forever, as the world is finite, but what stops an imbalance in the other direction?

Niche theory postulates that species avoid competitive exclusion by differing in ecological strategies, such as specializing on different resources or lifetime schedules. A niche-differentiated community is stable because species whose abundances are depressed tend to bounce back as the niche they occupy is suddenly left open (**?**). Without stabilization, local competitive dynamics inevitably leads to monoculture, unless species stocks are constantly replenished by sufficient immigration (**?**). Niche differentiation is therefore a key ingredient of coexistence theory and potentially a major part of the answer to the biodiversity question.

On the other hand, consider the Darwinian Demon—the super organism that maximizes all aspects of fitness. It cannot exist, for if it did we would not have biodiversity. The Darwinian Demon thought experiment proves that species are ecologically identical *on average*. Is it possible that they are also ecologically identical *all the time*—so-called ecological neutrality?

The idea that the obvious differences in species morphology, phenology, and behavior make no difference to their ecological performance defies belief, and yet **?** shocked ecologists by claiming that his neutral theory of biodiversity reproduces patterns of species abundances and common in nature. Debates ensued over whether neutral theory actually fared better than process-free statistical theories (**???**), whether the patterns described by neutral theory are informative of community assembly process (**??**), and whether species coexist because of ecological differences, ecological similarities, or both (**??**). Some authors proposed that at the very least, neutrality could replace statistical null hypotheses in ecology, because it is process-based rather than just a randomization of data, which may destroy pattern that could be caused by factors unrelated to niche differentiation (**??**).

These debates have highlighted the need for better tools to distinguish a neutral community from a niche differentiated one. Experiments are the golden standard for proving that competition relates to species similarity (**???**) and that coexistence is stable (**???**). However, one often cannot perform experiments at the scale of an entire community, and even at the individual scale they are typically limited to short-lived species. Among alternative methods, a particularly practical and hence commonly used one is inference of process from ecological pattern. Unfortunately, linking pattern to process is complicated when one process can lead to multiple patterns or when one pattern can be caused by multiple processes. The latter is often the case in nature (**??**); in fact, some common ecological patterns such as log-series species abundance distributions and power laws in species-area relationships may be common precisely because they arise from many different processes and not because they reflect some law of species interactions. This is the basis for statistical theories such as maximum entropy (**??**).

Cue in traits-based ecology. Trait data can be used to parametrize process-free statistical theories (**??**), but have unique potential for process inference: traits can be more revealing of niche differentiation because neutral theory proposes that traits are irrelevant. Indeed, trait patterns have been extensively used as an inference tool (e.g. **????**). Trait-based approaches have several advantages over strictly taxonomic approaches in that they are quantitative, easily generalizable, and have explicit ties to ecological strategy and performance (**??**). Furthermore, trait data are relatively cheap and easy to collect compared with genetic sequencing, and inference based on trait patterns does not require censuses spanning extended periods of time, which is particularly handy in the study of long-lived organisms such as trees. However, as we will see in Chapter 1, the theory connecting niche differentiation with trait patterns, although decades old, is still a work in progress.

This dissertation is about the connections between niche differentiation, coexistence, and patterns of community structure, especially as pertaining to species traits. We review early and recent literature, examine niche models for their ability to explain coexistence and generate pattern, analyze current inference-based methods, and offer new ones. The main goal is to advance theory connecting niche differentiation to coexistence and pattern, and contribute to improving niche inference methods.

1.2   Dissertation structure

In Chapter 1, we review the literature on using traits to infer niche differentiation. We find that expectations of trait pattern under niche assembly are largely based on the classical MacArthur-Levins model (**?**), which ignores the influence of stochasticity and immigration, and makes restrictive assumptions about the way competitive interactions relate to species traits. Furthermore, some well-established metrics of trait pattern are founded primarily on intuitive hypotheses linking niche differentiation and functional overdispersion (**??**), which however have not been validated with niche models. Classical predictions such as limiting similarity have been questioned in the past (**?**) and some questions remain unanswered. We argue that progress in the field requires theory development, which should entail investigating patterns across conceptual and system-specific niche-axis models. This research program is performed in Chapter 5, but first we reveal in Chapter 2 how a common but unrealistic assumption in niche models leads to inflated estimates of coexistence and lack of trait pattern. Chapter 3 applies the lessons of Chapter 2 to revise two prominent niche models, which are then studied for pattern in Chapter 5.

Many niche models assume, either implicitly or explicitly, that factors regulating species dynamics can be idiosyncratic to species regardless of their traits. For example, a species-specific parasite that affects a single species host and no other, no matter how similar the hosts are (**??**). As pointed out in Chapter 2, this assumption amounts to allowing clearly distinct ecological strategies between arbitrarily similar species. This seemingly innocuous "discontinuous strategy" assumption is at once biologically unrealistic and extremely consequential. For example, these models may predict stable coexistence between arbitrarily similar species, so-called continuous coexistence. In apparent contradiction with previous theoretical results (**?**), continuous coexistence in these models is not destroyed by small fluctuations in environmental conditions[1]. This has grave consequences for the

---

[1] There is no real contradiction, as such results do not apply to models with non-differentiable kernels.

prospects of using pattern to infer niches, because if a niche mechanism allows coexistence regardless of traits, then it cannot be expected to create trait pattern. In Chapter 2 we mathematically prove the connection between the discontinuous strategy assumption and robust continuous coexistence, and discuss why such assumption is fundamentally unrealistic and cannot be expected to be met in any real system.

In Chapter 3, we explore two influential niche models which implicitly make the discontinuous strategy assumption, namely Muller-Landau's tolerance-fecundity tradeoff (**?**), where there is an inverse relation between species stress tolerance and number of propagules produced, and Tilman's competition-colonization tradeoff (**?**), where species trade off fecundity with the ability to displace other species. We show that these models contain strategy discontinuity, which has a very strong impact on predictions of limits to similarity, species diversity, and trait pattern. In particular, both of these models predict robust continuous coexistence. We then offer revised formulations where strategy discontinuity is removed, and show that continuous coexistence disappears. The removal of strategy discontinuity from these models allows for better informed predictions of pattern, which will be examined in Chapter 5.

Chapter 4 examines how niche assembly differs from neutral assembly in terms of species lifetimes, that is, the amount of time between a species introduction to a local community by immigration, and its extirpation due to drift or competitive exclusion. We focus on species assemblages where the number of niches available is less than the number of species immigrating into the community – in other words, when not every species can occupy their own niche, and some species must therefore compete for the same niche. We show that, counterintuitively, average lifetimes in the niche-differentiated assemblage are typically shorter than in neutral assemblages. This occurs because, although some species are never extirpated as they are optimally adapted to their niche, most are outcompeted and thus quickly removed. We conclude that niches may lessen extinction rates and raise richness and mean lifetimes, or they may have the opposite effect, depending on the number of niches available to species. In particular, in a species assemblage similar to the tree community on Barro Colorado Island, extinction rates may actually be higher than neutral expectations if the number of niches is not sufficiently high. The results of this chapter qualify the sense in which niche differentiation enhances coexistence relative to neutrality.

Answering the call from Chapter 1, Chapter 5 studies trait pattern in a suite of niche

See Chapter 2.

models spanning several different niche mechanisms. We explore the recent findings that, contrary to widespread ideas of limiting similarity and trait spacing between coexisting species, niche assembly may lead to the spontaneous formation of clusters of similar species (**?**). Recognizing that much of existing pattern theory is based on models that ignore the influence of stochasticity and immigration, we explore trait pattern in stochastic niche assembly happening in local communities under propagule pressure from a regional pool. We ask whether clusters are a general phenomenon across niche mechanisms and can therefore be used to infer niche differentiation in nature. This chapter features the first use of an abundance-weighted metric for identifying and quantifying species clusters on a trait axis. We validate our metric by applying it on the MacArhtur-Levins model, where clusters are known to arise (**??**). We conclude that clusters are a typical outcome of models where competition is primarily determined by similarity in species traits, but are more likely to appear if immigration is low, the ratio between the number of species and niches is high, and competitive sorting is fast.

Chapter 6 shows that in a context where species may mutate into similar species, neutrality will also generate clusters, which needs to be accounted for in inference-based tests. One recent paper in microbiology developed a new test for inferring selection from abundance structure in microbial communities (**?**). The test is based on the premise that selection will cause nonrandom abundance structure among operational taxonomic units (OTUs, a common proxy for species in microbiology). Specifically, the authors hypothesize that if some OTUs are favored by selection, there should be a correlation between the abundance of other OTUs and their genetic distance to those favored OTUs, leading to clusters in OTU space. This test was used to reject neutrality in vertebrate gut microbiomes. Chapter 6 shows that neutral mutations acting alone or with immigration can also cause such clustering, a fact that was not considered by the authors in their null model. We argue that clustered patterns in OTU space may indeed be revealing of selection, but only if clustering caused by selection can be distinguished from clustering caused by a neutral regime where mutation and immigration are considered.

1.3 Publication status

As of August 3 2016, four of this dissertation's research chapters have been published in peer-review journals, while the remaining two are at the preparation and review stages for publication.

Chapter 1 was published in D'Andrea, R., & Ostling, A. (2016). Challenges in linking trait patterns to niche differentiation. *Oikos*.

Chapter 2 was published in Barabás, G., D'Andrea, R., & Ostling, A. M. (2013). Species packing in nonsmooth competition models. *Theoretical Ecology*, 6(1), 1-19.

Chapter 3 was published in D'Andrea, R., Barabás, G., & Ostling, A. (2013). Revising the tolerance-fecundity trade-off; or, on the consequences of discontinuous resource use for limiting similarity, species diversity, and trait dispersion. *The American Naturalist*, 181(4), E91-E101.

Chapter 4 is under review by Ecology.

Chapter 5 is in preparation for submission to Ecology Letters.

Chapter 6 was published in D'Andrea, R., & Ostling, A. (2016). Can Clustering in Genotype Space Reveal "Niches"?. *The American Naturalist*, 187(1), 130-135.

The research described in Chapters 1, 4, 5, and 6 was done in collaboration with Annette Ostling, and Chapters 2 and 3 were done in collaboration with Gyuri Barabás and Annette Ostling. In both of the latter two chapters, Gyuri and I took equal part in the research and development. Gyuri wrote Chapter 2 and I wrote Chapter 3, as well as all other chapters.

Chapter 2

Challenges in linking trait patterns to niche differentiation

Abstract

Among approaches to establish the importance of niche differentiation for species coexistence, the use of functional traits is attractive for its potential to suggest specific coexistence mechanisms. Recent studies have looked for trait patterns reflective of niche differentiation, building on a line of research with a deep but somewhat neglected history. We review the field from its foundation in limiting similarity theory in the 1960s to its resurgence in 2000s, and find the theory of trait patterning still in a stage of development. Elements still to be accounted for include environmental fluctuations, multidimensional niche space, transient dynamics, immigration, intraspecific variation, evolution, and spatial scales. Recent empirical methods are better than early approaches, but still focus on patterning arising in simplistic models, and should rigorously link niche space with trait space, use informative null models, and adopt new metrics of pattern as theory develops. Because tests based on overly simplistic expectations of trait pattern are of little value, we argue that progress in the field requires theory development, which should entail exploring patterns across a set of conceptual and system-specific models of competition along trait axes.

2.1 Introduction

The idea that competition relates to species similarity is central to theoretical community ecology, and has been verified in experimental studies (**???**). Its corollary, that coexistence is predicated on species differences, forms the basis of niche theory. Recent debates over the importance of niche differentiation (see glossary) for coexistence and community structure (**??**) have reignited empirical efforts to parse how strongly niche differentiation

drives community assembly and maintains biodiversity (**?**). Key approaches include direct verification of frequency-dependence (**?**), experimental removal of stabilization (**?**), searching for phylogenetic signatures of niche differentiation, e.g. (**?**), and examining the goodness of fit of coexistence models to observed species abundance distributions (**?**). As fruitful as these approaches may prove to be, they leave open questions about the specific niche mechanisms at work, and their role in shaping the presence and relative abundance of species.

In parallel, there has been a recent drive among ecologists to understand communities in terms of traits impacting individual fitness via their effects on growth, reproduction and survival, so-called "functional traits" (**?**). According to **?**, "statements about traits give generality and predictability, whereas [research based on species identities] tends towards highly contingent rules and special cases", because traits are more easily tied to physical, chemical, and biological processes. Of particular interest has been the use of trait-based patterns as evidence of niche differentiation: sorting species by their trait value, one looks for a pattern in species occurrence and/or abundance along the trait axis, (e.g. **?**). This approach to establishing niche differentiation could point towards specific niche mechanisms and a generalizable understanding of their influence on species abundances.

Linking trait patterns to niche differentiation is not new; rather, it represents the return of a classical line of inquiry dating back to the 1960s. Early research in this area was the subject of critical reviews highlighting methodological issues and the many simplifications of models used to derive expectations of pattern (**??**). In the 2000s, the subject acquired renewed momentum. A number of studies have reported patterns of trait dispersion that are interpreted as reflecting niche differentiation, (e.g. **????**, etc).

Despite the rich literature and existing reviews on related topics (**???**), there hasn't been a synthesis of the theoretical basis for pattern-based trait analysis specifically. Here, we revisit this science from early work to latest developments, assess how current research deals with past shortcomings, identify remaining and new issues, and map out next steps to meet those challenges. The central message of our review is that further theoretical development is critical for moving forward in this line of research, as current theory of trait patterning arising from niche differentiation is too rudimentary to inform us of the pattern we should truly expect in nature.

Section 2 reviews the origins of modern trait-pattern work in the classical limiting similarity theory of the 1960s and 70s, and describes the methodological and theoretical criticisms raised at the time. Section 3 considers new insights on the concept of limiting

similarity, but suggests that this concept is just one element of a more comprehensive theory of trait patterning arising under niche differentiation. Section 4 argues that new empirical approaches have met some but not all of the early criticisms. Section 5 describes new challenges that have come to light as interest in trait pattern research resurged. Section 6 proposes that examining trait pattern across a set of conceptual and system-oriented niche models can help refine predictions of trait pattern under niche differentiation.

2.2 Early trait pattern research: rise and fall of limiting similarity

The general expectation about species coexisting in the teeth of competition is that there should be limits to their similarity. This is an old idea in ecology, found in one form or another in the works of Gause, Hutchinson, and others. In his famous Santa Rosalia piece, **?** proposed heuristically that competition leads to a pattern of regularity in the body size ratio of trophically similar species, which he backed up with empirical evidence from waterbugs. A few years later, a similar assembly rule was derived by **?**. Using a simple Lotka-Volterra competition model with species arranged along an axis of resource preference, and assuming that the strength of competition between species is proportional to the overlap in resource preference ("niche overlap"), they obtained a quantitative limit to species similarity consistent with coexistence. In particular, they showed that when carrying capacities are equal across species, a new species can only invade and coexist between two residents if it is separated from each by at least one standard deviation of the resource use curve (Fig. 2.1A). This niche overlap rule marked the beginnings of the theory of limiting similarity.

? refined the theory by adding considerations of robustness. They showed that limits to similarity depend on species' carrying capacities, and that arbitrarily similar species can stably coexist if those values are just right, but the carrying capacity range consistent with coexistence decreases drastically with species similarity (Fig. 2.1B). Limiting similarity occurs because coexistence of very similar competitors depends upon fine-tuned parameters, whereas in nature those parameters are expected to fluctuate. By that logic, one could expect limits to similarity to be contingent upon the degree of environmental variation. However, **?** argued that the one-standard-deviation niche overlap rule proposed by MacArthur and Levins was universal, on two grounds. First, he highlighted that when distance between resource use curves d is less than one standard deviation σ, only a small range of carrying capacity parameters allows for coexistence, but once $d > \sigma$,

that range rises sharply (Fig. 2.1B). Second, upon adding white environmental noise to carrying capacities in the Lotka-Volterra competition model, he found that the rule holds for species-rich communities regardless of the magnitude of environmental variation, so long as it is not too large (Fig. 2.1C).

These papers, along with Hutchinson's original conjectures on coexistence constraints, triggered extensive empirical exploration focused on specific rules of trait separation, such as the aforementioned niche overlap rule and Hutchinson's constant body size ratios. By the mid-1970s, the literature featured numerous reports of confirmation of such rules (reviewed in (**?**)). Later, however, these studies were questioned regarding the quality and consistency of the results and methodology. **?** contended that overlap in commonly examined traits such as body size and bill dimensions may not be truly representative of niche overlap, and listed a number of studies where overlap was found to be in disagreement with the MacArthur-Levins rule. Furthermore, in a series of influential papers sparking heated debate among community ecologists, Simberloff and co-workers (**??**) explained away much of the claimed evidence for limiting similarity as either lacking statistical rigor or not distinguishable from null expectations. Their relentless campaign for a Popperian approach to community ecology led to the widespread incorporation of null models in subsequent studies.

The criticism didn't stop at the empirical evidence; the theoretical basis for the rules were questioned as well **?**. **?** showed that the minimum trait difference between resident species and a successful invader depends heavily on the shape of the resource use curve. **?** contended that the MacArthur-Levins rule and May and MacArthur's expansion of it were problematic because they ignored extinction of resources and hinged on the particular way overlap in resource use translates to Lotka-Volterra competition coefficients in the model. He emphasized that different assumptions lead to different formulas for the competition coefficients, which in turn lead to different limiting similarity rules. **?** criticized May's mathematical approach to environmental stochasticity, and through an alternative approach concluded that "practical limits to similarity can be obtained by simply ignoring the stochastic terms" (i.e. by using species' average carrying capacities). **?** used this to argue that limits to similarity will depend on the system and the species. Further, he suggested that observed spacing between species may often not conform to its theoretical minimum, as any such minimal separation would be enlarged by evolutionary character displacement. Later, he pointed out that theoretical studies that consider coevolution between competitors often find that the ecological and evolutionary effects of competition

differ, either increasing or decreasing similarity depending on the circumstances (**?**).

Given all these issues, one might conclude that limiting similarity is irrelevant to natural systems. That, however, is unjustified. Abrams did not dismiss limiting similarity as much as he expanded its meaning and highlighted its system-specific character. "What is needed instead is a broader definition of limiting similarity. The concept should be represented as a relationship between the difference in competitive ability and the maximum similarity that will permit coexistence. Such a relationship has the potential to be different for every different pair of species and will generally be different for a given pair of species in different environments" (**?**). Abrams saw a need for system-specific, field-parameterized competition models, and comparison of limits to similarity predicted by those models with similarity observed in nature.

Technically, the principle that limits to similarity always exist is unfalsifiable. One can always argue that the actual limits are smaller than those probed empirically. But this is mostly of academic interest, as most ecologists ask not whether limits to similarity exist, but whether they shape community structure. This can be answered by comparing observed trait structure with null hypotheses.

The term "limiting similarity" has lost purchase over the past two decades (**?**). Yet, most empirical trait-based studies of community assembly still expect biotic (niche) interactions to lead to trait overdispersion, an expectation that has its basis in limiting similarity theory.

2.3   New perspectives and prospects on limiting similarity

Abrams's redefinition of the concept of limiting similarity was recast by **?** in his increasingly influential framework for coexistence mechanisms. Chesson introduces fitness-equalizing forces, which decrease average differences in "fitness" between species, and stabilizing forces, which allow species to invade each other from low abundance. By "fitness" Chesson means the maximal population growth rate in the absence of resource limitation, scaled by the sensitivity of the species growth rate to resource availability. Chesson's "stabilizing" forces are those causing species' competitive influence on themselves to differ from that on other species and hence leading to species' population growth rates to be positive at low abundance. They require niche differentiation between species (**?**). Coexistence mechanisms involve one force or the other, or very often both. In fact, Chesson casts the population growth rate of an invading species as the sum of two terms: one quantifying the fitness differences between the invader and the residents, and the other

quantifying the effect of the stabilizing forces. Chesson revisits limiting similarity by pointing out that for coexistence to be possible, "the smaller the average fitness differences, the smaller the stabilizing niche differences can be."

Chesson's perspective has had a strong influence on community ecology (**?**), as it highlights how species differences can either favor or disfavor coexistence. However one must be careful in applying it. Chesson's ideas seem to naturally imply that a) if there are no fitness differences between species, there are no limits to similarity; b) conversely, if there are fitness differences, there must be limits to similarity. From a dynamic stability standpoint, both conclusions turn out to be incorrect for coexistence between more than two species. In the case of no fitness differences, the niche-axis Lotka-Volterra model proposed by **?** does not allow for stable coexistence between three or more arbitrarily similar species. It must be the case that Chesson's stabilization term of the growth rate of the potential third invader is actually negative unless the two resident species are far enough apart on the trait axis. On the other hand, some competition coefficient choices in the same model lead to coexistence of an indefinitely large number of arbitrarily similar species despite fitness differences caused by differences in carrying capacities (**?**).

Critically, these instances of stable coexistence of arbitrarily similar competitors do not warrant the conclusion that there may be systems in nature without limits to similarity, at least not yet. Limits to similarity pervade all competition models if one considers robustness to fluctuations in model parameters set by environmental conditions. For example, it has been shown for a large spectrum of models that small changes in the carrying capacity of a single species brings extinction to a number of its closest competitors (**???**), while changes in the carrying capacities of multiple species throughout the trait axis lead to spacing between all stably coexisting species (**?**).

The concept of robustness, or structural stability, has not gained as much traction as the concept of dynamic stability, even though it may be equally relevant for analysis of model predictions (**?**). In a dynamically stable state, species abundances return to their previous equilibrium values after a perturbation, whereas in a structurally stable state species equilibrium abundances change little when environmental conditions setting the carrying capacity are perturbed. Because both abundances and environmental conditions can be expected to fluctuate in nature, a predicted equilibrium state is arguably not biologically relevant unless it is both dynamically and structurally stable.

However, there is an unresolved tension between Turelli's work on environmental fluctuations, the robustness analysis of Barabás and coworkers, and work by Chesson on

environmental fluctuations (**?**). **?** concludes that small environmental fluctuations are of little consequence for limiting similarity theory, whereas Barabás and coworkers ascribe them primary importance. The latter approach differs from the former by considering potentially long-term changes in environmental conditions. Chesson suggests environmental fluctuations can enhance coexistence when species' responses to them are negatively correlated, but in this case environmental fluctuations open up new niche axes along which there will be limits to similarity[1]. Resolving this tension is needed to achieve a clearer picture of whether one does expect any special systems where limiting similarity does not hold, and a more general understanding of the likely role of environmental fluctuations relative to inherent population dynamics in setting the limits to similarity in nature. Doing so requires answering the empirical question of whether negatively correlated responses to the environment are common, and whether environmental fluctuations are fast and small enough for Turelli's analysis to apply, or longer term as assumed by the robustness analysis of Barabás et al.

Yet even once the nature of environmental fluctuations is better known, limiting similarity theory will be just one ingredient in a comprehensive theory of trait patterning. Imposing dynamic and structural stability constraints on simple niche models gives conceptual insight into what competition might do to communities, but it seems unlikely that it should be able to tell us what trait pattern to actually expect to find in nature, where so many additional complexities ignored in these models come into play. The next two sections will discuss the current status of trait pattern research in the context of old and new challenges to formulating such comprehensive theory and developing better empirical methods.

2.4   Recent empirical approaches in the context of past criticisms

A variety of recent studies look for patterns of spread and evenness in the distribution of species along a trait axis (**?????**—see (**?**) for a review). The observed distribution of species traits is compared with a null distribution, typically a random draw from the regional species pool, often weighted by regional abundance and constrained by the number of individuals or species in the sample. Analysis focuses on species differences, and in

---

[1]The reverse idea, that biodiversity stabilizes community-level responses to environmental fluctuations, has been borne out empirically (**?**). This has been theorized to occur if species-level responses to the environment are asynchronous (**?**), which presumably requires that species differ sufficiently in traits related to such responses.

particular how widely and evenly the observed species are distributed on one or more trait axes. Box 1 describes common metrics and test statistics used. If niche differences associated with that particular trait are an important driver of community assembly, or if evolution has driven niche differences in sympatric species, then species are expected to be more widely or evenly distributed than in the null models. Filtering of traits best suited for local abiotic conditions can also influence trait patterns in a local community and is expected to lead to the opposite patterns. Authors typically look for both processes, and in some cases factor filtering into the null model for niche differentiation by restricting the species pool (**??**), or employ metrics that should have little impact from filtering (see Box 1).

Such current approaches are an improvement over the first wave of empirical efforts of the 60-70s. Increased trait range and even spacing may arise regardless of the particular size of the niche spacing between species needed for robust and stable coexistence. Further, use of null models conferred the statistical rigor lacking in the purely observational studies that were previously the norm. Still, given the intense past debates in the literature and mixed results in detecting limiting similarity, careful assessment of recent approaches in the context of past methodological critiques is warranted.

Some recent empirical trait patterning studies take novel approaches, such as experiments to see how successful invading species impact community trait pattern (**?**), and looking at trait patterning along ontogenetic, successional, productivity, and environmental gradients (**????**) to gain additional evidence for community assembly processes (e.g. differences in trait pattern between saplings and adults may reveal processes influencing organism survival). We do not specifically address these approaches here, but our assessment of theoretical and methodological issues applies equally well to them.

2.4.1 Is evenness to be expected?

Despite improvements in the pattern expectations of empirical studies, a closer look on whether limiting similarity theory truly supports those revised expectations quickly turns up potential issues. In particular there are issues with the expectation of even trait spacing, which has been the focus of a number of studies due to the potential for this type of pattern to be indicative of niche differentiation. To our knowledge, **?** was the first to explicitly question even spacing as a general feature of competition, and he had previously argued that limits to similarity may vary between different species pairs in a single system (**?**).

The idea that niche assembly leads to even trait spacing follows intuitively if interspecific competition depends solely on species similarity. Indeed, evenness features strongly in

the MacArthur-Levins model (Fig. 2.1A). However, that model assumes a great deal of simplifying symmetry, including uniformity of carrying capacities and resource utilization function shapes and widths across species. It also assumes competition coefficients that depend exclusively on trait differences and not the trait value itself, a simplification that does not hold for traits such as plant size (**?**). It is also not seen in other models, for instance models with competitive hierarchies such as the competition-colonization tradeoff (**?**) and tolerance-fecundity tradeoff (**?**).

Simulations of the MacArthur-Levins model in a few example cases with heterogeneous carrying capacity exhibit even spacing (**??**), suggesting the effects of fitness differences on evenness may be small. In contrast, the dependence of competition on trait values can have large effects, as models with this dependence yield spacing rules that can involve irregular intervals between species (**??**). For example, examination of patterns in the Tilman-May-Nowak hierarchical competition model reveal no obvious signs of evenness in species placement along the trait axis, cf. figures in (**?**). Even in models that predict even spacing, evenness may only be expected when the niche space is full, i.e. when there are as many species present as can stably and robustly coexist. Lower diversity could allow for more variation in species spacing, unless fast-acting evolutionary forces drove species to more even spacing. We know of no multispecies evolutionary studies that investigate impact on even spacing (or more generally on the distribution of stably coexisting species, but see (**?**)), although those would be critical to establish whether evenness is to be expected at evolutionary time scales.

Some ecologists believe that models based on an indefinite number of consumer species feeding on a continuum of substitutable resources, such as the ones discussed in the previous paragraph, have little relevance for expected pattern in nature because most natural systems have a small number of resources and relatively limited species diversity. Of course any real system has a finite number of consumers and resources, so the real question is how large the number of species/resources needs to be for models based on continuums to be useful ("All models are wrong, but some are useful"). For the purposes of predicting features of trait pattern among local species sorting from a regional pool of species, continuum models seem appropriate when the number of species in the pool and number of potential local resources for them greatly exceeds the number of stably coexisting species allowed by the model. There are often more potential resources or "limiting factors" (**??**) than a first look might reveal. For example, terrestrial plants may be limited by the availability of patches of different ages since disturbance, as in successional

niche models (**??**).

Further exploration is in order before conclusions can be drawn, but if evenness is not a universal feature of niche differentiation, then empirical studies should shift focus to other aspects of pattern more robustly suggestive of niche differentiation.

2.4.2 Identifying niche axes and multidimensionality

Two important criticisms leveled at early trait studies continue to be overlooked or not fully addressed in much of the empirical literature. First, to rigorously connect trait patterning and niche differentiation, one must establish an explicit link between candidate niche axes in a given system – i.e. the aspects of organism performance in which differentiation confers stable coexistence – and the actual traits measured in field studies. Early trait studies were criticized because it was unclear that axes of variation being examined actually corresponded to niche differentiation.

Attempts at making such linkages have certainly been made, but there are stumbling blocks. Consider plant traits such as maximum height, seed size, and specific leaf area. These are known to influence species' distributions along abiotic gradients (**?**), but they are also thought to be proxies for species' positioning on life history axes (**?**) and to enable coexistence through a successional niche tradeoff (**??**). Many studies have connected these traits to demographic performance (e.g. seed output, or growth or mortality at various life stages); however, due to limitations in the prevalence of different light conditions in forests, what is typically considered is performance in full sun or in the understory, and not the demographic response to shading that is critical to successional niche differentiation (e.g. **??**, but see **?**). Beyond tree traits, theory is still developing regarding which types of trait variation and tradeoffs are related to stabilization and which are simply related to fitness differences and/or fitness equalization (**?**).

Another problem is that it is much easier to associate traits with individual performance than with population growth rates, but it is their effect on the latter that matters for niche differentiation. Very few studies have tried to verify that connection empirically (but see **?**). Considering competition models formulated in terms of individual performance, e.g. **?**'s (**?**) model of size-structured competition in forests, or even more detailed models like individual-based forest simulators (**?**), may aid this process. These hurdles have recently been circumvented in annual plant systems, where manipulative experiments can get at population dynamics and be used to parameterize competition models, to provide what seems at least compelling evidence of which trait differences contribute to stabilization

(**??**). But such manipulation is not always possible.

? provide a useful road map to the broad types of niche axes – spatial and temporal heterogeneity, natural enemies, and resource partitioning – that could be considered in forming linkages between traits and performance more informative of potential niche mechanisms. The next step is to develop mechanistic competition models geared to specific example systems, (e.g. **?**), in order to get a detailed sense of the candidate niche axes there, and then link traits to these niche axes through both empirical correlations and physiological theory, (e.g. **?**). **?** suggest that ecologists abandon pattern-based analysis altogether in favor of establishing trait-niche axis relationships. However, showing that trait variation drives variation along a viable niche axis does not in itself indicate that it is substantial enough, or of the right nature, to enable stable coexistence. Pattern-based trait analysis, when coupled with the establishment of trait-niche axis relationships, provides a means to gain suggestive evidence that the trait variation is actually enabling coexistence.

The second issue is that the niche space in most cases likely involves many dimensions. There is a dearth of models providing predictions for trait patterning caused by niche interactions mediated by multiple traits. Multidimensionality means more niche space, and hence more potential for empty niche space in low-diversity systems, and for trait patterns to be more complex or random-looking as a result. In addition, the number of dimensions of niche space can have effects on limits to similarity. For example, the minimum multidimensional spacing required for coexistence can be lower than the one-dimensional limit (**?**). Multidimensionality can also lead to greater sensitivity to environmental noise (May 1973, p. 164). Finally, perfectly even niche segregation occurring in multiple dimensions can elude detection on individual trait axes. For example, if the niche space is two-dimensional and niche axes are not parallel to measured trait axes, species that are evenly distributed in the full niche space might seem unevenly distributed when examined on a single trait axis. (Fig. 2.2).

Empirical trait studies typically deal with this by calculating distances in multidimensional trait space (**?????**) or performing a principal-component analysis to determine the axes where trait segregation is strongest, (e.g. **?**). Those are not without issues, though. Most of these multivariate metrics are direct extensions of their single-trait cousins (see Box 1), and therefore inherit the same limitations: namely, they are based on the idea that niche differentiation will lead to high trait dispersion and/or evenness, which may not always be justified. Furthermore, there are multiple ways to calculate distance in multidimensional space (**?**), and without some guiding theory one is left with finding

17

whatever leads to the strongest signal. Similarly, with PCA one heuristically finds which axes give the strongest pattern, and must interpret the result after the fact. The problem then is analogous to the Texas sharpshooter fallacy: by looking everywhere, one is bound to find pattern somewhere, but its meaning may be lost. Such heuristic measures may thus fail to identify combinations of traits from which the niche mechanisms involved can be surmised. An alternative approach is to develop and ground-truth system-specific mechanistic models to identify candidate niche axes, and then identify the multiple traits that may be driving them using physiological theory and phenomenology.

2.4.3   The null model

Trait studies may have universally accepted the importance of comparisons with a null model, but not enough attention has been given to what exactly the null model is supposed to represent. In recent trait studies, null assemblages normally consist of abundance-weighted random draws of a fixed number of species or individuals from the regional pool, or some related form of randomization of the observed data. Some authors use more sophisticated randomized statistical null models that, for instance, account for habitat filtering by selecting only species that have been observed to survive in the local environment (**??**). But are randomized null models our best option? Many ecologists say yes. **?**, citing **?**, argue that null expectations intrinsically involve randomizations of observed patterns. To them, null models should function as "a standard statistical null hypotheses, in contrast to a scientific hypothesis, which is a mechanism to explain the pattern".

Others disagree. Some authors, (e.g. **???**) argue that null models should, as far as possible, represent explicit ecological processes by which the community would have been assembled in the absence of the process of interest. In other words, process-based null models are preferable to statistical null models. For example, **?** incorporate historical processes of speciation and extinction ignored by statistical null models, and show that the latter may lead to false detection of limits to phylogenetic relatedness.

? pointed out that "fabricating random communities [. . .] does not bear the relationship to ecology that the neutrality hypothesis does to population genetics." However, little resulted from this insight because at the time there was no framework in community ecology comparable to neutral theory in population genetics. Yet the past fifteen years have seen the development of a neutral theory for community ecology (**??**), and we believe trait patterning studies should move towards using it. Indeed one study suggests the use

of a neutral model instead of a purely random null can lead to different results regarding patterns of species co-occurrence (**?**).

The idea of neutral theory as a null model in community ecology is not new, and has its opponents. Interestingly, while **?** reject it because neutral models are mechanistic rather than purely statistic, **?** rejects it because to him they are not mechanistic enough. Our perspective is that neutral theory presents some difficulties. First, it may require quite a bit of information about a system, and the formulation of complex neutral models accounting for that information. For example, details regarding the nature of dispersal limitation (**?**), size or age dependence of reproduction and death (**?**), and the spatial distribution of species in the regional pool, may all have important consequences for expected distributions in the absence of niche differentiation. However, less sensitivity to at least some of these details might be achieved by shifting focus from traditional species abundance distributions to distributions of species abundance in trait space, where more specific types of departures from neutral expectations might be expected. Power analyses and theoretical exploration of example tests could indicate what information is most vital, and in what situations the use of a neutral model rather than a pure random model is particularly important. Another issue is scale. It is very hard to get good information on the nature of neutral processes operating at the regional scale, such as speciation, and this makes it difficult to build a good neutral model at that scale. However, if data on current regional abundances are available and can be used to make adequate predictions of expected local abundances based on dispersal alone, one can test for non-neutral dynamics at the local scale without needing to make assumptions about neutrality at the regional scale.

It should be noted that alternative methods that do not use null models are also possible. For example, approximate Bayesian computation algorithms (**?**) have been used to estimate the relative importance of different community assembly processes in shaping trait patterns by comparing trait-based summary statistics produced under different model assumptions about process with summary statistics from the data (**?**).

2.5 Further elements to be considered in trait research

In addition to the classical methodological and conceptual issues discussed in the previous section, new challenges have come to light as interest in trait pattern research has resurged. As will be seen below, emergent clusters arising from classical niche models offer fresh perspectives on expectations of trait dispersion. Intraspecific variation

presents challenges in the degree of sampling needed to discern pattern, and could impact coexistence and hence trait patterning predictions. The spatial scale of both the mechanisms of niche differentiation and the sampling design impact the likelihood of detecting pattern and the nature of the pattern to be expected.

2.5.1 Emergent clusters

?'s (**?**) study and much subsequent work focus on system equilibria, as nonlinear models typically cannot be solved but their equilibria can often be found analytically. However, there is no guarantee that time to reach these equilibria is short enough that real-world communities can be expected to be at equilibrium. Advances in computer power have opened the door to simulations of model dynamics that were unavailable in the 1960s, and such simulations reveal that transient regimes may be in striking contrast to expectations from equilibria. Key among newly discovered phenomena is the spontaneous emergence of groups of similar species under competition. In simulations of the MacArthur-Levins model, starting with a large number of species randomly placed on a trait axis, **?** found a transient regime composed of clusters of similar species with high abundance, separated by sparsely populated regions on the axis (Fig. 2.3A).

The appearance of transient clusters in the competitive context can be informally understood as follows. Because competition increases with similarity, minimal interspecific competition is achieved by partitioning the trait axis into evenly spaced spots (determined by any small initial heterogeneity in species abundances). Species located at those spots will face the lowest competition and thus thrive, while all others are eventually excluded. This arrangement is the final equilibrium to which the system eventually converges. On the other hand, species close together on the trait axis are similar to one another in the competition they experience from other species. Transient clustering arises because it takes time for the species at the favored spots to exclude their slightly inferior neighbors. One can interpret the region surrounding each spot as a niche, and describe the dynamics as a two-stage process, the first characterized by species being sorted into niches, and the second characterized by within-niche competition leading to the exclusion of all but one species per niche.

We note that fitness differences (sensu Chesson 2000) play no necessary role in the formation of these clusters. In Scheffer and van Nes's (2006) model, all species have equal intrinsic growth rates and carrying capacities, and can be distinguished only by their position on the niche axis; yet, clusters arise spontaneously. Transient clusters could

also arise from fitness differences (**?**) and/or habitat filtering, but those would differ in trait space from Scheffer and van Nes's clusters: unless there were some substructure of resource availability driving a multimodal fitness curve, fitness differences would cause a single cluster around the fittest species. In contrast, the emergent clusters in discussion here consist of multiple regions on the trait axis—as many as the number of species that would stably coexist at equilibrium.

Scheffer and van Nes showed clustering patterns for a limited parameter range in one simple model, but clusters have since been seen in many studies. **?** verified the robustness of the clustering pattern in Lotka-Volterra niche overlap models with finite niche axes by showing that it arises as a transient under a variety of niche dimensions, competition functions, niche boundary conditions, species initial abundances, species intrinsic growth rates, and carrying capacities. Furthermore, Scheffer et al. showed that the addition of a stabilizing term to the population dynamical equations of the MacArthur-Levins model leads to permanent rather than just transient clusters (**??**, but see **?**). Further, clusters are permanent in niche models that include the effects of frequent environmental variation (**?**), mutation (**?**), and immigration (**?**)—see Fig. 2.3B. Lastly, Scheffer and van Nes found that implementing evolution by allowing species to move along the trait axis makes clusters more robust to variations in species carrying capacities (**?**). Indeed, clusters are no strangers to the evolution literature: it has been shown that phenotypic clustering can appear due to selection (**???**), or simply due to mutation in mutation-drift processes (**??**).

In nature, clustering in the positioning of species on trait axes has been seen in European aquatic beetles, Dutch freshwater algae, American prairie birds (**?**), marine and saline lagoon phytoplankton (**??**), and alpine woody plants (**?**). Some of the evidence is simply in the form of visually apparent modes in histograms of species counts along discrete portions of body size axes (**?**), but other studies used quantitative metrics such as looking for peaks in the curve of species density on the trait axis (**?**), or looking for a larger coefficient of variation in nearest neighbor distances than expected by chance (**?**). In addition to local peaks in richness on the trait axis, **?** and **?** also found evidence of local peaks in abundance, like those predicted in Lotka-Volterra models, using an abundance-weighted trait distance metric and a test looking for peaks in Shannon-Weaver diversity respectively.

In sum, clusters appear in nature and feature in Lotka-Volterra competition models, as well as evolutionary models. It seems reasonable to expect clusters to arise in a wide range of models of ecoevolutionary dynamics. More thorough examination using system-specific ground-truthed models is needed, in particular regarding whether clusters will arise

for realistic levels of environmental variation, mutation, and immigration. If clusters prove to be a general feature of niche models, this will have profound implications for trait-patterning studies, as a clustered community is the exact opposite of the type of trait dispersion generally associated with niche differentiation. Future metric development should take this into account (see Box 1).

2.5.2 Spatial scales

Spatial scales have been brought to bear extensively on examinations of community assembly patterns and ecosystem processes (see **?** and citing works), with focus on scale-dependence of species diversity, composition and spatial distribution, (e.g. **?**), phylogenetic relatedness, (e.g. **??**), and ecosystem function (**?**). Here, we highlight the relevance of scales to patterns of trait dispersion in the context of how niche mechanisms relate to scales of species dispersal and environmental heterogeneity.

Consider terrestrial plants. It may appear at first that competition occurs primarily among neighboring individuals, so niche patterns should occur at the neighborhood scale (**??**). But in reality, all individuals sending propagules to a location potentially compete with one another. Therefore, expectations of niche differentiation must encompass more than just local resource competition (**??**). In addition, niche mechanisms often require spatial variation in the environment to enable coexistence. For example, species may simply differ in habitat preference, so that variation in the conditions of local habitat are required for coexistence. At a small spatial scale encompassing a limited range of habitats, one would tend to find trait filtering, reflecting selection for tolerance to those habitats. At a larger scale, one would expect niche differentiation (Fig. 2.4; **??**). Appendix A describes how patterning across spatial scales might similarly emerge for niche mechanisms based on life history tradeoffs, and discusses some examples of traits and the respective scales at which niche patterning is expected.

Once one considers a variety of types of niche mechanisms, it becomes clear that the spatial scale at which trait-based niche patterns will be found depends on the particular niche mechanism. If trait variation stabilizes coexistence through spatial heterogeneity in the environment, then the grain of environmental heterogeneity will affect the expected scale of trait dispersion. This scenario presents a logistical hurdle for trait-based analyses that may be difficult to overcome: data may be needed from large swaths of land, and power issues may ensue (**?**).

Spatially explicit population dynamic models should be applied to more rigorously

derive and refine predictions of pattern. But, save for a few cases (**??**), this path remains essentially unexplored. Such studies must account for the potential interactions of organismal dispersal limitation and the scale of the niche mechanism. For instance, **?** showed that the degree to which Janzen-Connell effects maintain diversity in trees depends on the relative scales of dispersal limitation and enemy-mediated density dependence.

2.5.3 Intraspecific variation

Substantial trait diversity within populations is common in nature, and the subject has received extensive coverage in the past few years. Recent reviews have discussed potential implications for functional diversity, population productivity and stability, coexistence, and ecosystem processes (**????**). However, less is known about the potential implications for trait patterning.

Intraspecific variation raises obvious issues. For example, the common practice of using trait means to represent a species has been criticized for inaccurately representing resource use and strength of competition (**?**, Fig. 2.5). Indeed, empirical studies have found that considering intraspecific variation may reveal otherwise obscure pattern (**???**). For example, **?** found that detection of habitat filtering and niche differentiation is stronger when considering variation in trait expression within populations at the plot level and especially at the individual level, compared to using species means at the regional level.

Sampling effort is also a concern. Plant studies looking for niche differentiation rarely sample more than ten individuals per species, premised on the idea that trait variation is substantially larger across than within species. However, given observed levels of variation within populations, more recent studies suggest that sample sizes of 10 to 20 individuals are needed to reliably distinguish between species (**?**). Potential consequences of undersampling include false indication of the degree of trait spacing between species (**?**) and missing spatial structure in trait expression that is critical for niche differentiation. A sampling theory to guide empirical work on the appropriate sample size for estimating species means could be helpful, but one may ultimately need to move from using species means to indices of functional diversity designed with intraspecific variation in mind (**????**), or to grouping individuals by trait values rather than by species.

In addition to issues with sampling, likely impacts of intraspecific variation on trait patterning depend on how it affects coexistence (we discuss recent research on this question in Appendix B). For example, if intraspecific variation facilitates coexistence by slowing down competitive exclusion, then transient patterns such as emergent clusters

can be expected to be persistent and easily maintained by immigration. Alternatively, if intraspecific variation enables coexistence by allowing species to occupy more than one viable niche, then coexistence between species with any trait means is possible. Any trait mean in a particular range can be achieved by having some individuals at a viable niche on one end of that range and some at the other. Thus one would expect weak to no patterning in species mean traits. In contrast, abundance patterns of individuals grouped by trait values rather than species would be preserved.

Very few theoretical studies have tried to verify how intraspecific variation affects trait distribution between coexisting species (but see **?**). In simulations of the competition-colonization model, **?** found that the addition of intraspecific variation changes the region of trait space that allows for coexistence between two species. In particular, a species with intermediate fecundity and mean growth rate can coexist with another very similar species only if intraspecific variation is present. The authors conclude that intraspecific variation in that model leads to the disappearance of limits to similarity. (Importantly, robustness to parameter changes was not examined. Indeed, it seems that a slight variation in mean growth of a coexisting competitor could result in exclusion.) Moving forward, theoretical studies of the impact of intraspecific variation on coexistence and trait patterning should model intraspecific variation not just as a fixed parameter, as is often done, but as a dynamical quantity shaped by a combination of selection and the genetic details of trait inheritance.

2.6 Toward a theory of trait patterning under niche differentiation

The primary challenge in linking trait patterns with niche differentiation is gaps in the theoretical foundation from which empirical expectations are derived. Theoretical progress can be made by exploring a spectrum of niche-axis coexistence models for insight on the types of patterns to be found under different biological conditions. We suggest a two-pronged approach with analyses of both conceptual and ground-truthed system-specific models.

Simple conceptual models allow exploration of the impact of the most salient features of competitive interactions commonly found in nature. In the MacArthur-Levins model, competition strength depends strictly on trait differences. This is a common feature of nature (**???**) and the defining property of a niche mechanism. The pattern arising from the model can be taken to represent the effect of niche differentiation on trait distribution

when acting alone. The competition-colonization and tolerance-fecundity tradeoffs (**??**) are simple representations of competitive hierarchies that may be prevalent in nature (**?**) but do not fit into the MacArthur-Levins framework due to the dependence of competition on traits as well as trait differences. We suggest that the MacArthur-Levins model and simple hierarchical tradeoff models are a good starting point for developing trait pattern theory. This review discussed natural phenomena likely to contribute to pattern which are absent in these models, such as environmental fluctuations, multidimensional niche space, immigration, intraspecific variation, evolution, and spatial scale-dependent interactions. A systematic approach starting with these conceptual models and adding in simple representations of each of these additional factors in turn could provide insight into the separate and combined effects of different processes on trait dispersion.

Examining this spectrum of models will no doubt be revealing. Yet it does not exhaust the complexity of competitive interactions in nature. For example, competition between neighboring trees in forests is size-structured, and it is unclear how this local scale complexity translates into population-level competition. Hence to complement study of conceptual models, we suggest that a suite of system-specific models also be developed and studied for trait pattern predictions. Such models would be built on empirical measurements of how traits determine an individual's response to resources and other individuals in the system of interest. Here we envision models with a level of detail similar to the forest architecture hypothesis model (**?**). Additional phenomena like those mentioned above (environmental fluctuations, immigration, etc) can be incorporated to the degree to which measurements suggest they are prevalent in the system of interest.

These model-based studies will give insight into the range of trait patterns reflecting niche differentiation in species assemblages, and culminate in more sophisticated predictions for trait pattern and new metrics designed to better handle potentially complex pattern than the ones currently available (see Box 1). That said, it is entirely possible that no distinguishable pattern will be observed under some circumstances, such as high immigration and severe competitive asymmetries that depart from the distance-based MacArthur-Levins competition. In that case, trait-pattern analysis will probably be of little avail, and more data-heavy methods such as experiments to directly verify stable dynamics would be required. When trait pattern is distinctive, analysis of pattern across spatial scales and environmental gradients could lead to hypotheses about the predominant community assembly processes, to be then further tested with experiments or observational studies.

2.7 Conclusions

Empirical efforts should be informed by theory, and recent modeling results regarding the distribution of species along trait axes depart from classical expectations of overdispersion and even spacing in trait space. Recent field studies look for a less restrictive signature of niche differentiation than earlier ones, but continue to focus on patterns arising in overly simplistic models. Most still fall short of establishing firm links between trait axes and niche axes, and use null models that may not accurately represent expected background processes. The next generation of field studies should be founded on expectations that contemplate patterns coming out of more realistic implementations of niche models.

The theory of trait patterning is itself still a work in progress. Particularly needed are exploring more complex niche mechanisms than the distance-dependent interactions in the MacArthur-Levins model, elucidating the influence of environmental fluctuations on pattern, incorporating intraspecific variation and explicit spatial structure in models, and extending our purview to out-of-equilibrium phenomena such as transients and factors that interact with such dynamics, such as immigration and evolution. Further work on conceptual models may shed light on how different assembly processes shape communities, while system-specific models may clarify the range of patterns to be expected in nature. Ultimately, trait pattern analysis is but one tool in the toolbox of inference methods. There are of course many complementary empirical approaches to infer niche differentiation from a wide array of observational or experimental data, including phylogenetic dispersion (**?**), variations in species composition through time (**?**), direct measurement of frequency-dependent competition (**?**), etc. As trait patterning research matures, it can be used in tandem with other methods to paint a more complete picture of community assembly than currently available.

2.8 Glossary

Cluster: a set of species with similar trait values and relatively high abundance, separated from other such sets by sparsely populated regions on the trait axis.

Community-wide character displacement: the idea that species coexisting within a community will differ from each other more than expected by chance, which will be reflected as trait overdispersion.

Competition function, (competition kernel): function describing the competitive suppression of a species' per capita growth rate caused by one individual of the same or another species. In trait-based models, it is often a function of species traits. It is the extension of the community matrix in the limit of a continuum of species.

Expected scale of trait dispersion/Scale of the niche mechanism: The spatial scale of sampling required to observe trait pattern. Different niche mechanisms lead to trait patterning at different spatial scales. For example, Janzen-Connell effects are expected to lead to overdispersion at spatial scales within the range of interactions between individuals, whereas niche mechanisms that rely on habitat heterogeneity are expected to lead to overdispersion at scales commensurate with the grain of the environment.

Fitness (sensu Chesson, 2000): in the context of a consumer-resource model, it is the mean per capita growth rate of the species in the absence of resource limitation, scaled by the rate at which the per capita growth rate declines as resource abundance declines. For stable coexistence, fitness differences must be compensated by the stabilization term.

Limiting similarity: has taken on multiple meanings over time. In the 1960s it was understood as a rule for coexistence that imposed a specific minimum limit to the similarity of two coexisting species, namely the distance between two consumers' resource use curves could not be smaller than one standard deviation of those curves. Later, it was redefined more generally by **?** as the recognition that limits to the similarity of coexisting competitors always exist, but their value depends on the species/system involved. It is also sometimes equated with community-wide character displacement.

Niche axis: the axis of variation involved in a niche mechanism, if that variation can be sensibly arranged on an axis. For a successional niche tradeoff, the axis runs from gap-specialist to shade-tolerant species.

Niche differences/differentiation: species differences that lessen competition between them. As argued by **?**, these are necessary for stable coexistence. Niche differences will involve one or more niche mechanisms and may or may not fall on a one-dimensional

niche axis.

Niche mechanism: a process or set of circumstances whereby variation across species can lessen competition between them. In a niche mechanism, the nature of species interactions, plus sometimes the presence of tradeoffs, lead to opportunities for species to differ in their interaction with limiting factors (Levin 1970). Examples include enemy specialization, variation in resource use, and the successional niche tradeoff, where plant species differ by age of patches exploited. Resource use curve: specifies a consumer's preferences among a large array of substitutable resources. In the MacArthur-Levins 1967 paper, it describes the probability of use of a given resource in a unit time by an individual.

Robust coexistence: an equilibrium state that resists changes in parameters that regulate dynamics. For example, a species is robustly coexisting if it is still able to coexist upon a (small) reduction in its carrying capacity due to removal of nesting sites.

Stabilization: in Chesson's formalism, it is the contribution to the invasion growth rate from differences between species competitive influence on themselves and their competitive influence on other species. It contains the frequency-dependent part of a species' growth rate, and can enable stable coexistence by causing a negative correlation between a species' abundance and its growth rate.

Stable coexistence: an equilibrium state that resists (small) changes in species abundances. For example, if some individuals from one of the species are removed, the species will tend to bounce back towards higher abundance.

Trait: any measurable feature of an organism's physiology/ morphology/ behavior/ phenology, such as a canine's tooth size, or a plant's specific leaf area. Often "traits" is used implying specifically "functional traits", meaning those that impact individual fitness via their effects on growth, reproduction and/or survival.

Trait pattern: non-random distribution of species along a one-dimensional trait axis. This distribution may be based on presence-absence or weighed by species abundances.

Trait-niche axis relationships: niche differences are not necessarily available for easy measurement, so variation in more accessible traits are often used as proxy for niche differentiation. However, a relationship needs to be established between a species' placement on the niche axis and the trait axis before overdispersion in the latter can be taken as indication of niche differentiation.

2.9 Box 1: Trait-based metrics

A variety of trait metrics have been used to detect niche differentiation. Here we describe key metrics used in empirical studies, which can be presence-absence or abundance-weighted, one-dimensional or multivariate. Multivariate metrics tend to be direct analogues of one-dimensional metrics. In common, almost have been applied with the expectation that niche differentiation will lead to high trait dispersion and/or high evenness in the distribution of species across trait space. As discussed in the main text, this is problematic because more complex niche differentiation mechanisms (e.g. hierarchical tradeoffs) do not necessarily predict even spacing nor overdispersion (see Section 4), and emergent clusters may occur (see Section 5). We also describe some metrics that have been proposed to look for patterns of clustering, and offer suggestions for further metric development.

Presence-absence metrics

Trait range. If strength of competition declines as species differ, species traits should spread over a wider range than could be expected by chance (**?**). However, trait range is impacted in the opposite direction by habitat filtering (**?**), making it difficult to parse the influences of filtering and limiting similarity that may act on the same trait simultaneously. It is also sensitive to outliers (**?**), and does not incorporate abundance.
Minimum distance between nearest neighbors. This offers a direct measure of the potential limits to similarity, which would be reflected as larger-than-random minimum separation between species (**?**).
Mean distance between nearest neighbors. A larger value than expected by chance could be indicative of competition dependent on trait differences (**?**). Mean distance is not independent of trait range and hence subject to the same confounding influence of habitat filtering.
Trait evenness. Measures variation in trait differences between adjacent species on the trait axis. Like the mean, the standard deviation of nearest neighbor distances scales with trait range and hence is subject to the confounding effects of habitat filtering, see e.g. **?**. Normalized measures such as the coefficient of variation and the ratio between the standard deviation and the range are used to account for that effect (**?**). Limiting similarity is expected to be expressed as excess evenness and thus low coefficient of variation. For related metrics, see e.g. **?**.

Abundance-weighted metrics

Functional regularity index (FRO). An abundance-weighted measure of trait evenness, based on trait differences between adjacent species on the trait axis scaled by the total abundance of the pair (**?**).

Functional divergence (FDvar). Measures the abundance-weighted mean of the deviation of species traits from the mean species trait. A high functional divergence signals high niche differentiation (**?**).

Community-wide abundance-weighted distance. $\sum_{\text{all pairs}} d_{ij} n_i n_j$, where i and j represent two species, n is their abundance, and d is the absolute value of the trait difference (**?**). This index should be higher in communities driven by niche differentiation than in communities where trait value are unrelated to abundance. A similar metric, Rao's quadratic diversity, is defined as $\sum_{\text{all pairs}} d_{ij} p_i p_j$, where p is the relative species abundance, $p_i = n_i / \sum_k n_k$. It is the expected dissimilarity between two individuals randomly sampled from the community (**?**).

Multidimensional metrics

Functional richness (FRic), or convex hull volume. The hypervolume spanned by the trait means of observed species in multidimensional trait space (**??**). It is the multidimensional equivalent of the trait range, and hence suffers from the same problems as a metric for niche differentiation.

Minimum spanning tree. A measure of the functional diversity spanned by the species in a community. Essentially collapses multidimensional trait space into a one-dimensional space. It is constructed by joining all species together in a tree, with branches weighed by Euclidean distances in multidimensional trait space. Large trees and even spacing within trees are indicative of niche differentiation (**?**).

Functional evenness (FEve). Measures the regularity with which species abundances are distributed across the minimum spanning tree (**?**). It is a multivariate extension of the functional regularity index (see above).

Multivariate functional divergence (FDiv). Multivariate analogue of FDvar. Calculates the abundance-weighted deviations between species traits and the trait center of gravity in multidimensional trait space (**?**).

Functional dispersion (FDis). The abundance-weighted mean distance between species and their abundance-weighted centroid (**?**). It is related to Rao's quadratic diversity, which

expresses the mean Euclidean distance between two randomly selected individuals.

Metrics proposed to detect clustering on the trait axis

Trait clustering. Variation in trait differences between adjacent species on the trait axis can reflect clustering if it is higher than expectations from a null (**?**). It should be noted that many of the metrics listed above would work equally well to detect clustering, as that would be reflected as the opposite direction from overdispersion in the metric.

Distribution of species richness along the trait axis. Based on the idea that similarity in fitness between similar species, coupled with partitioning of the trait axis into niches, may lead to an uneven distribution of species along the axis (**?**).

Distribution of distance to nearest abundant species. Proposed by **?** for genetic sequences[2] of gastrointestinal organisms as an alternative to rank-abundance plots. A proportion of the sequences is arbitrarily selected as modal (highly abundant), and the distribution of distances of all other sequences to their nearest modal sequence (the most similar modal sequence) is obtained and compared with the expected distribution from a neutral community. However, appropriate null models must be used so that clustering from mutation alone can be discarded, see **?** (Chapter 6).

Potential sources for further metric development

Time series analysis. Species abundances on a trait axis can be taken as a time series, with the trait axis being time and the abundances being the signal. Fourier and wavelet analysis can be run to look for periodic waves of high and low abundance on the axis, corresponding to the emergent groups (niches) found by **?**.

Spatial descriptive statistics. Spatial descriptive statistics that are normally used in geographic information systems (GIS), such as Ripley's K, could be used to measure the degree of clustering in the distribution of species trait values in trait space. Clustering in trait space, as determined by such metrics, would indicate filtering, while overdispersion would indicate niche differentiation. This would be especially useful in multivariate analyses.

Trends in abundance as a function of distance to niche centers. Based on the idea that niche differentiation will lead to clustering, metrics focusing on patterns of species abundances as a function of the position of a species in relation to others within a cluster

---

[2]Specifically, operational taxonomic units, a commonly used proxy for species in microbiology.

could prove useful. To our knowledge, this has not been investigated yet, and remains a promising line of research.

2.10 Figures

Figure 2.1: The niche overlap rule. **A**. MacArthur and Levins (1967) concluded that if the separation d between nearest-neighbors were much smaller than the standard deviation w in resource use curves, coexistence would not be possible. **B**. The range of parameters allowing 3-species equilibrium ("coexistence bandwidth"), as a function of niche overlap, d/w. For simplicity, resident species 1 and 3 are chosen to have the same carrying capacity, $K_1 = K_3$. The bandwidth (shaded area) is narrow for similar species ($0 < d/w \ll 1$), and disappears as d/w approaches zero (adapted from May 1973). **C**. "The closest niche overlap d/w consistent with community stability in a randomly varying environment with fluctuations characterized by variance σ^2/K. Over a wide range, the variance has little influence on the species packing distance for a number of surviving species $n > 2$" (from **?**).
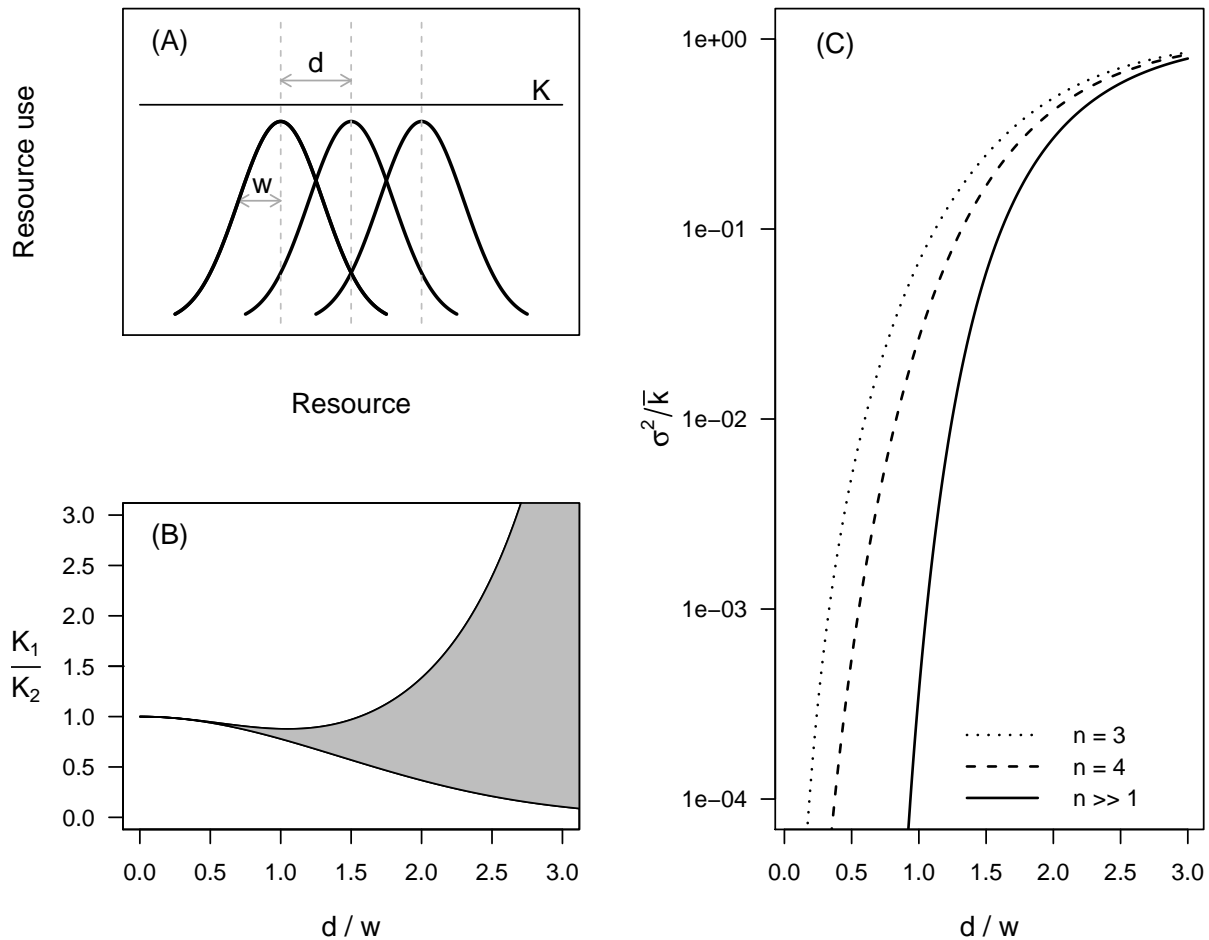
Figure 2.2: When separation between species is defined in multidimensional space, limiting similarity may be hidden from the viewpoint of a single axis. Here, even though species are evenly distributed when the full 2-dimensional trait space is considered, such structure is lost when we take a projection onto a single axis. For example, the coefficient of variation (CV) of the adjacent-neighbor distances between the projections on the axis shown above is 0.9. In contrast, the CV in 2-D is 0. When compared with 10,000 null distributions with the same number of species and same first and last positions but random positions in between, almost half of the null CVs are found to be lower than 0.9. Thus the assemblage above, as seen from that axis, would be indistinguishable from random. Note that different angles between the orientation of the lattice and the axis will lead to different CVs and p-values.
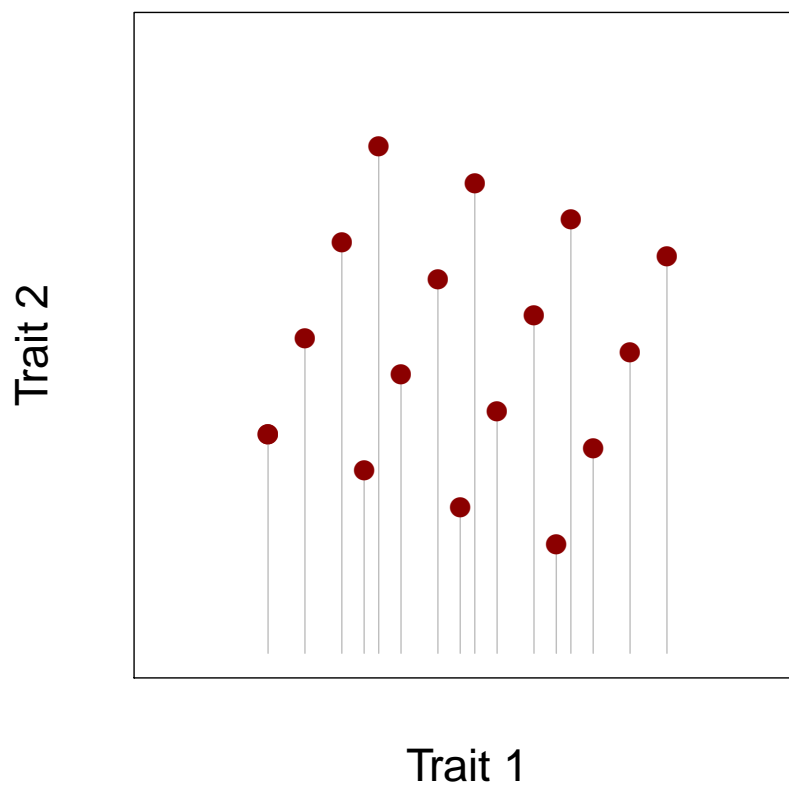
Figure 2.3: Emergent clusters. **A**. Transient regime in the MacArthur-Levins model. The community is initially packed with species lined along the trait axis. Before the final state of species at their limits to similarity is reached, clusters are formed around the favored sites on the axis. **B**. Outcome when immigration is added to the MacArthur-Levins model. The clusters differ from A in that they are maintained by immigration, and are thus a permanent feature of the community. This pattern was highlighted in (**?**) (Chapter 2).
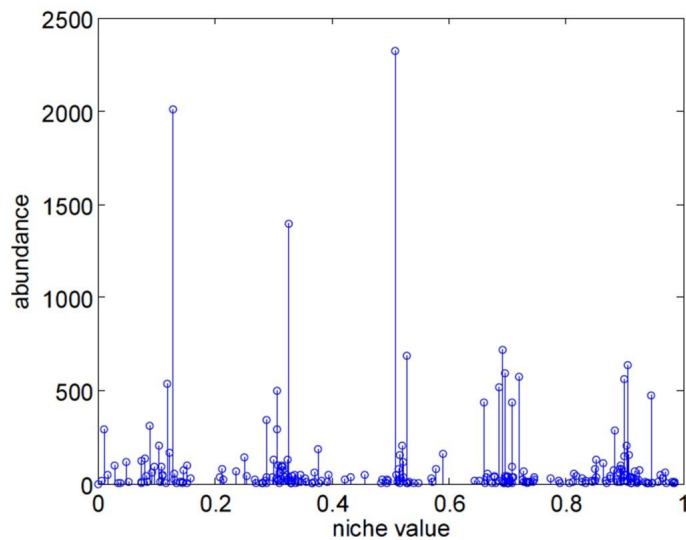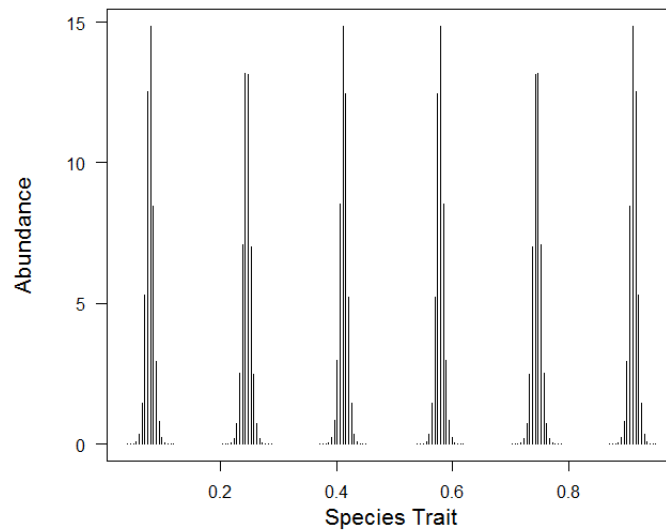
Figure 2.4: Niche mechanisms and spatial scales. Top panels: representation of niche assembly operating at different spatial scales. Different symbols represent individuals from different species distributed spatially. **A**: Case where coexistence relies on a heterogeneous environment. Species traits associate with environmental condition (represented by shading color). In this scenario, evidence of niche assembly is stronger at intermediate or large scales rather than the neighborhood scale. **B**: Case where the coexistence mechanism acts at the smallest scale, such as Janzen-Connell effects. In that case, evidence for trait dispersion is strongest at the neighborhood scale. Three sampling scales are shown in each scenario (dashed squares). **C**: Compared to a null assemblage, dispersion in local samples can be expected to decrease from small to intermediate or large scales if assembly mechanisms operate at the immediate vicinity of individuals (grey curves), but increase if assembly mechanisms rely on environmental heterogeneity (black curves).
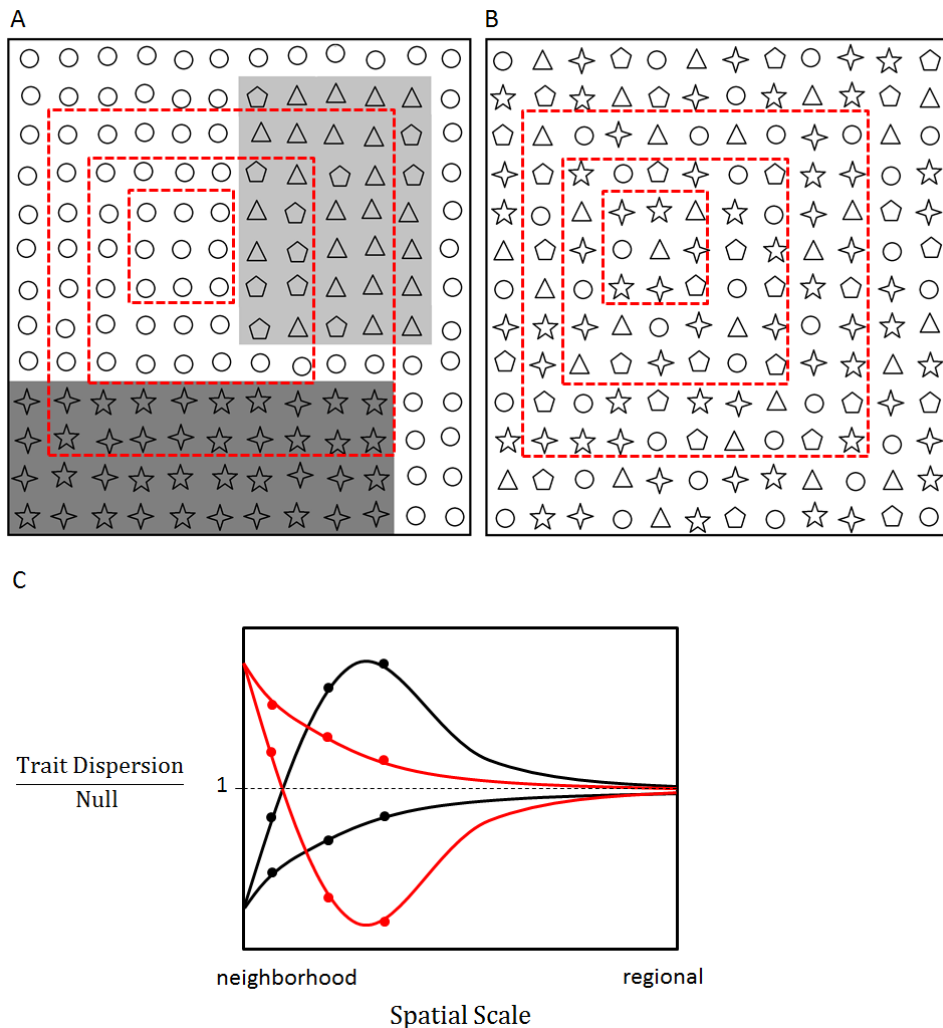
Figure 2.5: Effect of intraspecific variation on competition. We consider competition between individuals with traits x_1 and x_2 following $\alpha(x_1, x_2) = \exp[-(x_1 - x_2)^2]]$, as in **?**, and compare the strength of competition estimated from species mean trait values, $\alpha(\bar{x}_1, \bar{x}_2)$ to the average strength of competition between pairs of non-conspecific individuals, $\overline{\alpha(x_1, x_2)}$. **A** and **B**: Consider two species whose trait distribution and means are shown in the black and grey curves and lines, respectively. If intraspecific variation is neglected, competition is overestimated when the two species have the same means (A), but underestimated when the distance between the species means is large compared to the trait spread (B). **C**: The ratio $\overline{\alpha(x_1, x_2)}/\alpha(\bar{x}_1, \bar{x}_2)$ increases from lower to higher than 1 with distance between species means. For comparison, if competition depends linearly on species positions (black curve), intraspecific variation brings no difference, on average, to the competition between homogeneous species.
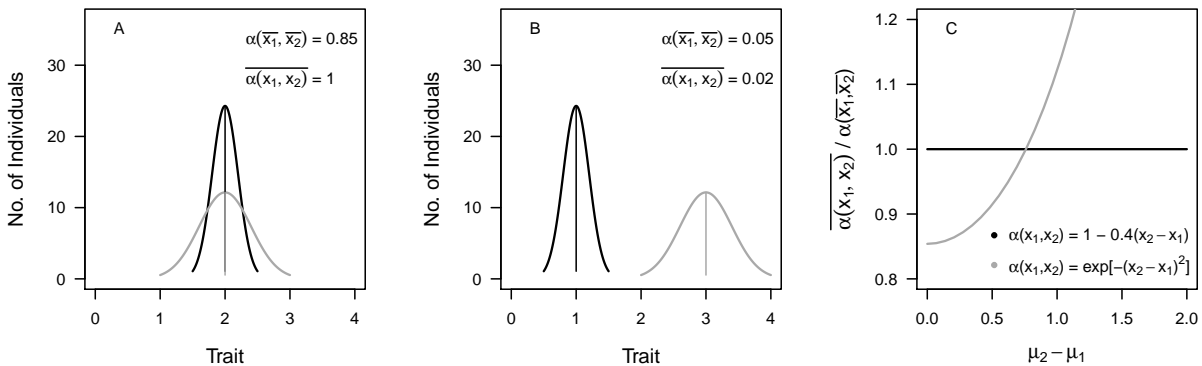
Chapter 3

Species packing in nonsmooth competition models

Abstract

Despite the theoretical possibility for competitive dynamics to allow coexistence of species with arbitrarily similar traits, prior work has shown that this should not occur in nature. A key reason is that known instances of continuous coexistence are fragile, requiring environmental conditions to be just right: a small fluctuation leads back to the classical limiting similarity predictions. Here we present and refute a potential challenge to limiting similarity. Robust continuous coexistence can arise if competition between species is modeled as a nonsmooth function of their differences—specifically, if the competition kernel (differential response of species' growth rates to changes in the density of other species along the trait axis) has a nondifferentiable sharp peak—a "kink"—at zero trait difference. The difference in predicted behavior stems from the fact that when the kernel is smooth, competition declines only to second order with species trait differences when those differences are small, creating strong competitive interactions between very similar species. Kinked kernels, on the other hand, decrease linearly, which dramatically relieves interspecific competition even between extremely similar species. We investigate what mechanisms lead to kinked kernels and find that they are created by discontinuities in resource utilization. We argue that such sudden jumps in the utilization of resources are unrealistic and all but logically inconsistent, and therefore one should expect kernels to be smooth in reality.

3.1 Introduction

The Darwinian view of life can be summarized as follows: 1) competition between similars is too strong for coexistence to happen, and the ensuing competitive exclusion favors the more fit type, thus driving natural selection and the evolution of all the marvelous adaptations on our planet; and 2) competition between sufficiently dissimilars can be reduced to a level where there is no competitive exclusion, leading to coexistence and the fantastic diversity of life we see around us. Darwin's insight does lead to some natural questions: what do species have to be different in to coexist, and just how much dissimilarity is sufficient to avoid competitive exclusion?

The first question was the main focus of early competition theory (**???**). The conclusion was that at equilibrium, no two species may consume the same resources. Later **?** noticed that, from a mathematical point of view, there is no essential difference between what we would call a "resource" and all other possible things that provide a negative feedback loop between growth rates and densities. These generalized resources (called limiting factors by Levin, and regulating factors by **?**, p. 288 and **?**, p. 146) are the things then that species have to utilize differently in order to coexist. Hence, traits associated with resource consumption (or, more generally, population regulation) are expected to differ among coexisting species: if bird populations are limited by seeds of various sizes, then differences in beak size would indicate specialization to different resources and therefore ecological differentiation.

The second question, how much interspecific dissimilarity is needed for coexistence, becomes important if there are infinitely many resource variables, as, e.g., in the case of a seed size continuum. The most important early result concerning this problem is by **?**, who demonstrated that limiting similarity (i.e., a tendency towards the spacing of phenotypes along the trait axis with exclusion zones in between) is the expected equilibrium behavior. However, their conclusions came into doubt when later work (**???**) demonstrated that not only are there no strict limits to similarity, but it is even possible for a continuum of species to stably coexist. These results lead to the paradoxical situation where, on the one hand, competitive exclusion seemed to be an irrelevant idea for ecology, but on the other hand nobody ever questioned the reality of Darwinian natural selection, which is strictly dependent on the ecological process of competitive exclusion between similar heritable phenotypes.

However, later it has been observed that while there are no formal limits to similarity, the more tightly packed a community is, the less robust it is against perturbations of model

parameters (**???**). In particular, it has been shown (**?**) that robustness (i.e., the volume in parameter space allowing for stable coexistence) always decays to zero with increasing similarity in *any* model of coexistence. Analogously, **?** proved an important theorem, demonstrating that if a continuum of species coexist, there always exists a perturbation of arbitrarily small amplitude that would destroy that coexistence. The extreme fragility of tightly packed communities leads to a reinterpretation of the old limiting similarity principle. Instead of asking how similar the species may be, we ask how robust any given coexistence pattern is. Since tightly packed species are so fragile, and random parameter variation is inevitable in a noisy environment, the default expectation for model behavior and empirical observations will still be limiting similarity — although the precise limits emerging will depend on model details. Thus, the apparent paradox of how natural selection could be a driving force in biology when there are no formal limits to similarity has been resolved by shifting the focus from the stability of coexistence to its robustness.

Here we show that there is another potential theoretical challenge to the expectation of limiting similarity. We demonstrate through numerical calculations that there are several cases where, though perturbations of arbitrarily small amplitude may still lead to the extinction of certain species (as is guaranteed by the Gyllenberg-Meszéna theorem), the general pattern of continuous coexistence is in fact quite robust. We will call situations where continuous coexistence is not entirely destroyed by perturbations *robust continuous coexistence*. What the models producing robust continuous coexistence have in common is that their competition kernels, defined as the differential response of the growth rate of the species with trait x to a change in the density of the species with trait y, is nondifferentiable whenever $x = y$, i.e., the kernel possesses a sharp peak or even a discontinuity at zero trait difference. This is in contrast with the classical practice of modeling the competition kernel as a strictly smooth function (and by smooth we will mean "differentiable at least once" throughout the chapter), usually of Gaussian form (but see **??**). We will say that such kernels possess a "kink" at the point of self-competition. We then further motivate our hypothesis that the property of possessing a kink is the key to robust continuous coexistence through two analytical arguments. The first one is based on a two-species coexistence scenario: we show that under this property of the competition kernel, limits to the similarity of two species disappear as long as certain (not very restrictive) conditions are satisfied. The second argument is based on the asymptotic properties of Fourier transforms, showing that models with smooth kernels tend to be more fragile than models with kinked ones. Finally, we discuss the mechanisms that lead to kinked kernels in the first place.

However, in light of these mechanisms, we argue that nonsmooth competition is unrealistic, i.e., it is not an accurate representation of competition that is expected to occur in nature. We base this argument on a demonstration that kinked kernels will not occur in the presence of intraspecific variation. Even in the absence of intraspecific variation, environmental variation would still lead to the smoothing out of kinked kernels. Therefore, we argue that one in fact should not expect kernels to be kinked, and therefore limiting similarity is still the expected behavior for stably coexisting species.

Competition kernels which are kinked according to our definition have been used in the context of the competition-colonization model (**??**), the competition-mortality tradeoff model (**?**), a model of seed size evolution (**?**), models of superinfection (**?**), the Lotka-Volterra competition model (**???**), and the tolerance-fecundity tradeoff model (**?**). Some of these studies (**???**) point out that sharply asymmetric competition (in which the better competitors have a much larger influence on the poorer competitor than vice versa) may lead to higher diversity and therefore tighter species packing along the trait axis, and **?** and **?** also emphasize the compromised realism of the assumption of sharp asymmetry. However, none of this prior work has studied the robustness of coexistence patterns predicted by these kernels, or identified the key property of the competition kernel influencing predicted patterns and their robustness. Our results here suggest that for considering the question of how much coexistence can be robustly generated by a given mechanism, the model of that mechanism should be constructed with care. In particular, although kinked kernels can provide a simpler, more analytically tractable description of competition mechanisms (as in, e.g., the competition-colonization tradeoff model), they lead to a vastly different answer to how much coexistence is to be expected. Note however that a key theme emerging from prior work is unchanged: some system-specific limits to the similarity of species along trait axes should be expected in practice, i.e., there should exist a minimum trait distance between stably coexisting species in any model, but this minimum distance will be different from model to model. Hence our work here provides development of the theory supporting the search for patterns of dispersion in trait-based community ecology (**?????**).

The chapter is structured as follows. After building the model framework and reviewing some of the better-known results emerging from it in Section 3.2, we go on to show examples of the model with kinked kernels (Section 3.3), which invariably produce robust continuous coexistence. Next, in Section 3.4 we give some mathematical arguments for why kinked kernels would have this property, but not smooth ones. Finally, in Section 3.5 we derive the conditions that lead to kinked kernels, and demonstrate that under realistic circumstances

one should always expect kernels to be smooth.

3.2 Background

3.2.1 Models of competition around equilibria

We wish to study the equilibrium patterns of competing organisms that vary in a single quantitative trait x. This trait parameter may assume any value within certain limits: $x \in [x_0, x_m] \subseteq \mathbb{R}$. We call the set of possible trait values x the *trait axis*. The canonical example for such a system is a community of birds with beak size x whose competition is mediated by the consumption of seeds of various sizes: this example is good to keep in mind, though our treatment will not be system specific. The most general continuous time, continuous density model within this framework reads

$$\frac{\mathrm{d}n(x)}{\mathrm{d}t} = n(x)\, r(n, E). \tag{3.1}$$

Here $n(x)$ is the abundance distribution of traits, $n(x)\,\mathrm{d}x$ measuring the number (or density) of individuals with trait values between x and $x + \mathrm{d}x$. While we write down differential equations to describe how $n(x)$ evolves as a function of time, we are primarily interested in $n(x)$ under equilibrium conditions — consequently, we simply write $n(x)$ instead of $n(x, t)$. The symbol r is the per-capita growth rate, which is a functional of the densities and all density-independent parameters, denoted by E (which could also depend on trait value). In principle, this equation could still produce arbitrarily complicated behavior. Therefore from here on we make the assumption that the system converges to some fixed point attractor. Then the per capita growth rates may be linearized around the fixed points. Denoting the equilibrium density distribution by n^*, we get

$$\begin{aligned}
\frac{\mathrm{d}n(x)}{\mathrm{d}t} &\approx n(x)\Big(\underbrace{r(n^*, E)}_{0} + \delta r(n, E)\Big) \\
&= n(x)\left(\int_{x_0}^{x_m} \frac{\delta r(x)}{\delta E(y)}\delta E(y)\,\mathrm{d}y + \int_{x_0}^{x_m} \frac{\delta r(x)}{\delta n(y)}\delta n(y)\,\mathrm{d}y\right),
\end{aligned} \tag{3.2}$$

where $r(x)$ is shorthand for $r(n(x), E(x))$ and the δ denotes functional differentiation (for those unfamiliar with functional derivatives, note that the expression

$$\delta r(x) = \int \frac{\delta r(x)}{\delta n(y)} \delta n(y) \, dy,$$

where x and y are continuous variables, is precisely analogous to the formula $dr_i = \sum_j (\partial r_i / \partial n_j) dn_j$ where i and j are discrete indices; see, e.g., **?** for the precise definition). Denoting the first term of the expansion by $c(x)$ and the functional derivative $\delta r(x)/\delta n(y)$ by $-a(x, y)$, this may be rewritten as

$$\frac{dn(x)}{dt} = n(x)\left(c(x) - \int_{x_0}^{x_m} a(x, y)\delta n(y) \, dy\right). \tag{3.3}$$

Using the fact that $\delta n(x) = n(x) - n^*(x)$, this dynamical equation can be brought to the usual Lotka–Volterra form:

$$\frac{dn(x)}{dt} = n(x)\left(c(x) - \int_{x_0}^{x_m} a(x, y)\big(n(x) - n^*(x)\big) \, dy\right), \tag{3.4}$$

or

$$\frac{dn(x)}{dt} = n(x)\left(\underbrace{c(x) + \int_{x_0}^{x_m} a(x, y)n^*(y) \, dy}_{r_0(x)} - \int_{x_0}^{x_m} a(x, y)n(y) \, dy\right), \tag{3.5}$$

and so

$$\frac{dn(x)}{dt} = n(x)\left(r_0(x) - \int_{x_0}^{x_m} a(x, y)n(y) \, dy\right), \tag{3.6}$$

where $r_0(x)$ is an effective density-independent growth term. (Note that the form of the equation preferred by most textbooks is recovered through the definitions $r(x) = r_0(x)$, $K(x) = r_0(x)/a(x, x)$, $\alpha(x, y) = a(x, y)/a(x, x)$). This equation applies around any fixed point equilibrium; the linearity of the approximation ensures equivalence with the Lotka–Volterra equations.

The function $a(x, y)$ is called the *competition kernel*. It measures the effect of a change in the abundance of species y on the growth rate of species x. In general it may be an arbitrary function of its arguments, but since we are interested in competitive systems, we shall make two assumptions. First, the kernel has to be nonnegative; this means that the growth of any one species necessarily inhibits the growth of the others and so there

are no mutualistic and/or exploitative interactions present. Second, the kernel should decrease with increasing $|x - y|$: competition is assumed to be stronger between more similar phenotypes. Without this assumption, being sufficiently different in phenotype would not confer an advantage and so there would not be any interesting coexistence patterns to analyze in the first place.

3.2.2 The fragility of continuous coexistence solutions

As mentioned in the Introduction, the original idea of strict limits to similarity had to be abandoned when it was demonstrated that even in the original Lotka–Volterra model (where the idea was first proposed) it is possible to have the stable coexistence of a continuum of species (**?**). However, such coexistence is extremely sensitive to perturbations of model parameters and is therefore not expected to occur under realistic circumstances. Let us investigate the original example of Roughgarden and its behavior under model perturbations. From Eq. (3.6), the equilibrium condition reads

$$r_0(x) = \int_{x_0}^{x_m} a(x, y)n(y) \, dy \tag{3.7}$$

for any species with positive density. Assuming $x_0 = -\infty$, $x_m = \infty$, and the functional forms

$$r_0(x) = \exp\left(-\frac{(x - x^*)^2}{2w^2}\right), \tag{3.8}$$

$$a(x, y) = \exp\left(-\frac{(x - y)^2}{2\sigma^2}\right) \tag{3.9}$$

for the parameters, it can be shown that the solution $n(x)$ will also assume the Gaussian form

$$n(x) = \frac{w}{\sigma \sqrt{w^2 - \sigma^2}} \exp\left(-\frac{(x - x^*)^2}{2(w^2 - \sigma^2)}\right) \tag{3.10}$$

as long as $w > \sigma$.

This solution is structurally unstable, i.e., a perturbation of arbitrarily small amplitude may destroy it (**?**). Fig. 3.1 shows an example where the continuous coexistence pattern collapses completely, even though the perturbation amplitude is small. Note that the spacing between surviving species is almost perfectly even, as expected in this model for the type of perturbation we employed (**?**).

It is instructive to look at these results in light of the Gyllenberg-Meszéna theorem (**?**). As a matter of fact, this theorem is a collection of several related results. But, for our purposes, we only need to distinguish between two cases. The first one concerns the equilibrium condition Eq. (3.7) in its full generality. It first assumes that, given the continuous parameters $r_0(x)$ and $a(x, y)$, an equilibrium solution $n(x)$ is produced whose support (i.e., values of x for which $n(x)$ is nonzero) includes a domain of continuous coexistence. Then the theorem states that there exists a positive function $\eta(x)$ such that for an arbitrarily small ε, if one replaces $r_0(x)$ by $r_0(x) + \varepsilon\eta(x)$, the resulting perturbed solution $\hat{n}(x)$ will not have the same support as $n(x)$. In other words, some species are bound to go extinct, no matter how small the disturbance is: continuous coexistence is, in this sense, fragile. Notice that the theorem does *not* say that continuous coexistence *as a whole* is going to collapse, merely that certain species will go extinct. However, a stronger version of the theorem, guaranteeing that an arbitrarily small perturbation can break down all continuous coexistence and lead to strict spacing can be proven for the special case of $a(x, y) = a(x - y)$, where $a(x - y)$ and $r_0(x)$ are analytic functions of their arguments.

This second, stronger theorem applies to the example in Fig. 3.1, since the parameters are all analytic. Therefore it is no surprise that continuous coexistence is completely destroyed. The next section will explore what happens if the parameters are not chosen to be analytic. It will be shown that spacing is still expected for kernels that are *smooth*, i.e., differentiable at least once: though technically speaking the stronger version of the Gyllenberg-Meszéna theorem does not apply, the results look as if it did. However, when the kernel becomes nondifferentiable at zero trait difference, the situation changes drastically.

3.3 Demonstrating robust continuous coexistence under kinked kernels

Fig. 3.2 presents several examples of smooth nonanalytic kernels (column 1) that support continuous coexistence (column 2). Our method for generating these solutions was to first choose a positive $a(x, y)$ and $n(x)$ arbitrarily, then use the equilibium condition Eq. (3.7) to obtain the corresponding $r_0(x)$ by performing the integration. Then the function $r_0(x)$ was perturbed and we obtained the solution to the perturbed problem by numerically integrating Eq. (3.6) (column 3). The four examples presented differ in whether the kernel is a function of trait diference only ($a(x, y) = a(x - y)$, rows 1 and 2, or $a(x, y) \neq a(x - y)$,

rows 3 and 4), and in whether the kernel is symmetric or not ($a(x, y) = a(y, x)$, rows 1 and 3, or $a(x, y) \neq a(y, x)$, rows 2 and 4).

In all cases, continuous coexistence is completely lost following the perturbation, and only a finite number of phenotypes persist, more-or-less evenly spaced out. The behavior of these models is therefore indistinguishable from the one we expect when the kernel $a(x, y) = a(x - y)$ is analytic (to which the strong version of the Gyllenberg-Meszéna theorem applies). We did not prove it mathematically, but based on our simulation results we will take it for granted that in all cases when the competition kernel is a smooth function of its arguments continuous coexistence collapses after perturbation and limiting similarity is recovered. In other words, a tightly packed community is extremely fragile to model perturbations, both with smooth and analytic kernels.

The situation is entirely different if the kernels are kinked (nondifferentiable at zero trait difference). Fig. 3.3 is analogous to Fig. 3.2, except that all kernels are kinked, which is evident from their graphs in column 1 (they all possess a sharp peak at each point where $x = y$). In these examples, though a few species do go extinct after perturbation, continuous coexistence itself is not eliminated: most regions on the trait axis still have arbitrarily similar species coexisting. This is exactly the situation we called robust continuous coexistence in the Introduction. Nondifferentiability at zero trait difference therefore has a tremendous impact on the robustness of the coexistence of similar species.

The perturbed densities in column 3 of Fig. 3.3 are not very different from their unperturbed counterparts (column 2), except in the direct vicinity of the perturbation. The effects of the perturbation therefore seem to be very local: beyond a certain distance, the coexistence pattern behaves as if no perturbation would have occurred at all. This distance depends on perturbation size, as Fig. 3.4 demonstrates: the larger the perturbation, the larger the exclusion zone in which species are driven extinct. Beyond that zone, however, coexistence is unaffected.

3.4 Kinked kernels and robust continuous coexistence

Why do kinked kernels lead to robust continuous coexistence while smooth kernels do not? We present two mathematical arguments why this is so: a two-species coexistence analysis and a multispecies one based on simple properties of Fourier transforms.

Consider two species that are extremely similar along the trait axis. The difference in their $r_0(x)$ values may then be expanded to linear order in the trait difference, neglecting

higher order terms. If the competition kernel is smooth, then the smallest nontrivial order of expansion of the kernel around zero trait difference is quadratic, since the kernel has a maximum there. Hence, to first order, the competitive effect of one species on itself is equal to its effect on the other and vice versa. Competition is therefore not reduced between the species: coexistence will in general not be possible (**??**). On the other hand, if the kernel is kinked, the linear-order decrease in competition is not zero anymore and so competition may immediately be reduced to tolerable levels where the two species can coexist, even for arbitrarily similar trait values. The abrupt decrease in competition in the case of kinked kernels brings about the possibility of the competitive coexistence of arbitrarily similar species. The precise, quantitative form of this argument is found in Appendix C.

Suggestive as it is, this result only applies for two competing species. We know and have seen in Sections 3.2 and 3.3 that smooth kernels do sometimes allow for continuous coexistence, so the limiting similarity condition obtained for the two-species case does not directly apply. However, the extreme fragility of such solutions signals that limiting similarity is still to be expected in all cases where the parameters have not been precisely fine-tuned. No such fine-tuning is required for retaining continuous coexistence in the case of kinked kernels. In the remainder of this section we demonstrate the extra fragility of continuous coexistence with smooth kernels via an argument based on Fourier transforms. This comes at a price though: only the $a(x, y) = a(x - y)$ homogeneous case may be treated in this manner.

For the special case $a(x, y) = a(x - y)$, the equilibrium condition Eq. (3.7) reads

$$r_0(x) = \int_{-\infty}^{\infty} a(x - y)n(y)\, dy, \tag{3.11}$$

where the limits of integration have been extended from minus to plus infinity for future convenience (since $r_0(x)$ can be arbitrarily small outside a relevant domain of trait values, this assumption is not really restrictive). Assume the equation has a positive solution $n_0(x)$. Now we perturb the left hand side with the arbitrary function $\eta(x)$, multiplied by the small parameter ε:

$$r_0(x) + \varepsilon\eta(x) = \int_{-\infty}^{\infty} a(x - y)n(y)\, dy. \tag{3.12}$$

This equation can be solved via Fourier transforms, invoking the convolution theorem.

Defining the transform of a function $f(x)$ as $\mathcal{F}(f) = \int_{-\infty}^{\infty} f(x)\exp(-i\omega x)\,dx$, we get

$$\mathcal{F}(r_0) + \varepsilon\mathcal{F}(\eta) = \mathcal{F}(a)\mathcal{F}(n), \tag{3.13}$$

which yields the solution

$$n(x) = \mathcal{F}^{-1}\left(\frac{\mathcal{F}(r_0)}{\mathcal{F}(a)}\right) + \varepsilon\mathcal{F}^{-1}\left(\frac{\mathcal{F}(\eta)}{\mathcal{F}(a)}\right) = n_0(x) + \varepsilon\mathcal{F}^{-1}\left(\frac{\mathcal{F}(\eta)}{\mathcal{F}(a)}\right). \tag{3.14}$$

The new solution is the sum of the unperturbed densities plus a perturbing term. As a side note, the solution is clearly unstable if the transform of the kernel is zero for any given frequency. This, however, will not happen if the kernel is chosen to be *positive definite*, i.e., $\iint f(x)a(x-y)f(y)\,dx\,dy > 0$ for all functions f, a simple consequence of which is that the Fourier transform of the kernel is strictly positive (**??**). Therefore we assume now that the kernel $a(x-y)$ is indeed positive definite.

The ratio $\mathcal{F}(\eta)/\mathcal{F}(a)$ is therefore finite for any given frequency, but might increase without bounds as frequencies go to infinity. If the Fourier transform of the kernel decays faster asymptotically than the transform of $\eta(x)$, then no matter how small ε is, there will always exist some frequency for which the ratio $\mathcal{F}(\eta)/\mathcal{F}(a)$ is large enough to make the solution $n(x)$ nonpositive for certain x values, destroying the original coexistence pattern.

We are going to use the following simple property of the Fourier transform (e.g., **?**). A function proportional to a Dirac delta has a transform which does not decay to zero asymptotically for large frequencies. A function with a finite jump (discontinuity) has a transform that decays asymptotically to zero as ω^{-1}. A continuous nondifferentiable function's transform decays as ω^{-2}, a function which is differentiable once has a transform decaying as ω^{-3}, and so on: the Fourier transform of a k-differentiable function decays asymptotically as ω^{k-2}.

Returning to the ratio $\mathcal{F}(\eta)/\mathcal{F}(a)$: due to the above property of the Fourier transform, if the kernel is differentiable k times, then the perturbing function $\eta(x)$ has to be differentiable $j > k$ times, otherwise the perturbing term in Eq. (3.14) will grow arbitrarily large, irrespective of the value of ε.

To give a specific example, let us define the perturbing function as

$$\eta(x) = \int_{-\infty}^{\infty} \frac{u(x-z)u(-z)}{\int_{-\infty}^{\infty} u(y)u(y)\,dy}\,dz, \tag{3.15}$$

48

where $u(x) = 1 - |x/\sigma|$ for $|x| \le \sigma$ and zero otherwise (the general shape of $u(x)$ is given in the top left corner of Fig. 3.2). It is easily seen that $\eta(x)$ is differentiable twice, therefore we expect its Fourier transform to decay asymptotically as ω^{-4}. This is indeed the case, since the transform of $\eta(x)$ is

$$\mathcal{F}(\eta) = \frac{3e^{-2i\omega\sigma}(e^{i\omega\sigma} - 1)^4}{2\sigma^3\omega^4}. \tag{3.16}$$

Now we choose a competition kernel that is differentiable more than twice, e.g., a Gaussian one:

$$a(x - y) = \exp\left(-\frac{(x - y)^2}{2\sigma^2}\right). \tag{3.17}$$

Its Fourier transform is also Gaussian:

$$\mathcal{F}(a) = \sigma\sqrt{2\pi}\exp\left(-\frac{\omega^2\sigma^2}{2}\right). \tag{3.18}$$

The ratio $\mathcal{F}(\eta)/\mathcal{F}(a)$ is

$$\frac{\mathcal{F}(\eta)}{\mathcal{F}(a)} = e^{\frac{1}{2}\sigma^2\omega^2}\frac{3e^{-2i\omega\sigma}(e^{i\omega\sigma} - 1)^4}{2\sqrt{2\pi}\sigma^4\omega^4}, \tag{3.19}$$

which clearly gets larger and larger for high frequencies. Therefore the solution cannot remain positive for all x: the perturbation will break the coexistence pattern, no matter how small ε is.

If, on the other hand, we assume a different form of the competition kernel, one that is kinked:

$$a(x - y) = \exp\left(-\frac{|x - y|}{\sigma}\right), \tag{3.20}$$

then $\eta(x)$ will never be able to break the coexistence pattern for ε sufficiently small. The Fourier transform of this kernel is

$$\mathcal{F}(a) = \frac{2\sigma}{1 + \sigma^2\omega^2}, \tag{3.21}$$

decaying asymptotically as ω^{-2}, as it should (since this kernel is continuous nondifferentiable); $\mathcal{F}(a)$ therefore decays more slowly than $\mathcal{F}(\eta)$. Their ratio is

$$\frac{\mathcal{F}(\eta)}{\mathcal{F}(a)} = \frac{3e^{-2i\omega\sigma}(e^{i\omega\sigma} - 1)^4(1 + \sigma^2\omega^2)}{4\sigma^4\omega^4}, \tag{3.22}$$

asymptotically decaying as ω^{-2}. It is well-behaved, its inverse Fourier transform will be

finite — and therefore there exists a sufficiently small ε such that the original coexistence pattern is unaffected.

Our result says that the more differentiable the competition kernel is, the larger the class of perturbations that can break the continuous coexistence pattern it generates. More specifically, if the kernel is differentiable k times, then a perturbation differentiable $j < k$ times will destroy the coexistence for any value of ε. Kinked kernels are nondifferentiable and so the patterns they generate cannot be broken for an arbitrarily small ε by differentiable perturbations: only nondifferentiable or discontinuous perturbations will be able to do that.

3.5 How do kinked competition kernels emerge?

3.5.1 Discontinuous utilization curves lead to kinked kernels

So far we have been discussing the impact of kinked kernels on the outcome of competition models. What biological factors would lead to such kernels in the first place is a question that remains to be answered. In this section we answer the question in the context of resource overlap models, i.e., we assume that if $u(x, z)$ is the rate at which a resource item of size z is consumed by a member of the species with trait x, then the kernel will read

$$a(x, y) = \int_{z_0}^{z_m} u(x, z) u(y, z) \, dz, \tag{3.23}$$

where z_0 and z_m are the maximum and minimum resource size, respectively (**???**). We also assume that the utilization function is bounded and only depends on the difference between resource type and trait: $u(x, z) = u(x - z)$. Then the competition kernel will also be a function of only the trait difference, since the amount of overlap depends only on how far the two traits are from each other, not on their absolute positions along the trait axis. (Appendix D generalizes the overlap picture to arbitrary ecological models, where it turns out that it is always possible to write the kernel as the overlap of two *different* functions, called the sensitivity and the impact; see also **?**).

With these assumptions we show that simple jump discontinuities in the resource utilization function are responsible for generating kinked kernels. The general analysis, not dependent on any of these assumptions about $a(x, y)$, is found in Appendix E, yielding very similar results and interpretation.

A kinked kernel is nondifferentiable at zero trait difference, therefore its second

50

derivative at that point is infinite. Our strategy is to take the second derivative of the kernel and determine the conditions under which it would be infinitely large. First we fix the trait value y to be zero without loss of generality, so that $a(x - y) = a(x)$ is a function of a single variable. The second derivative will read

$$a''(x) = \int_{z_0}^{z_m} u''(x - z)u(-z)\,dz,\tag{3.24}$$

where the prime denotes differentiation with respect to the argument. Now let us fix x to be zero as well:

$$a''(0) = \int_{z_0}^{z_m} u''(-z)u(-z)\,dz = -\int_{z_0}^{z_m} u''(z)u(z)\,dz\tag{3.25}$$

after a convenient change of variables $z \to -z$. Since in general the integral of the second derivative of a function is finite if the function is continuous but infinite if it possesses a jump discontinuity, we can already see that such discontinuities in u will make the kernel kinked. Let us assume now that the function u is continuous except at a point z^*. This means that u can be written as

$$u(z) = \alpha\Theta(z - z^*) + \eta(z),\tag{3.26}$$

where Θ is the Heaviside unit step function, α is a constant and $\eta(z)$ is a continuous function. Substituting this form into Eq. (3.25) we get

$$a''(0) = -\alpha \int_{z_0}^{z_m} \delta'(z - z^*)u(z)\,dz + \ldots,\tag{3.27}$$

where δ' is the derivative of the Dirac delta function, and the ellipsis means all other terms the derivative produces that have not been written out. (The derivative of a Dirac delta might seem like a strange construct, but not only is well defined, it also behaves in exactly the way one would intuitively expect, i.e., $\int \delta'(x - y)u(y)\,dy = -u'(x)$; see ? for the rigorous definition.) The integral of these other terms denoted by the ellipsis is necessarily finite and so they cannot contribute to the nonsmoothness of the kernel. Performing the integration with the help of the δ' function yields

$$a''(0) = -\alpha u'(z^*) + \ldots = -\alpha^2 \delta(0) + \ldots,\tag{3.28}$$

which is infinitely large. Note that if u has more than one discontinuity, $a''(0)$ will be a sum

of similar terms, i.e., each discontinuity contributes minus infinity times a constant squared to the expression above. Thus we have shown that the competition kernel is kinked if the utilization function has one or more discontinuities somewhere in its domain. Since we assumed u to be bounded, the converse will also be true (the most singular way a bounded function may behave is to be discontinuous, and the integral of a continuous function is differentiable). We therefore conclude that the competition kernel is kinked if and only if u has discontinuities. Finally, note that this result applies even if u is not a function of the difference of its arguments, and holds even if the kernel is not expressible via the overlap of utilization functions; see Appendix E for the generalization.

3.5.2 Mechanisms inhibiting discontinuous resource utilization

How is this result to be interpreted? A discontinuity in the resource utilization function means a species utilizing a certain resource is suddenly incapable of utilizing another, arbitrarily similar resource with similar efficiency. Expanding on the example of the competing bird species, one might imagine that each species has a box-like utilization curve: within a certain range σ of the beak size, all seeds are equally consumable, but outside of that limit, none at all ($u(x - z) = u_0$ if $|x - z| \leq \sigma$ and zero otherwise). Then, no matter how similar two species are, one will have access to seeds of certain sizes that the other does not, and vice versa (Fig. 3.5). Thinking of the various resources as the factors regulating the populations, this means that no matter how similar, the two species will still be independently regulated, which is the key to species coexistence in general (**??**). It follows that two species very similar along the trait axis are not really similar in the relevant sense of the word: no matter how close they are in their traits, their way of relating to the available regulating factors will be different, meaning that they are ecologically differentiated and thus can coexist.

This simple interpretation is not quite watertight because any discontinuity will lead to kinked kernels and therefore robust coexistence of arbitrarily similar species, not just those discontinuities that occur between some finite value and zero. Still, even if the jump occurs between two nonzero values, one can say that the species relate to arbitrarily similar resources in a qualitatively different way, bringing about their automatic ecological differentiation.

Natura non facit saltus — or does it? The question remains: what biological mechanisms would lead to sharp discontinuities in the resource utilization curves of species? Although one should not take the old Leibnitzian principle for granted (at least not in ecology), the

question raised by **?** is still a serious challenge: what qualitative difference could there be between two bird species which only differ in that one has a beak $1\mu m$ larger than the other, when clearly no one would even notice that there are two separate species to begin with? The question may be analyzed more clearly if, instead of asking whether nature exhibits jumps, we ask whether the kinds of *models* we use would exhibit them. Here we give two arguments supporting the assertion that sudden jumps will in fact never occur in the kinds of deterministic competition models we have been considering.

The first thing that has a smoothing effect is intraspecific variation in traits. Even if the utilization function of an individual with a given trait is discontinuous, one must not forget that not all individuals of a species are alike: as with all quantitative traits, there is some variation around a mean trait value. Let the "raw" utilization function be $u(x - z)$, assumed to be discontinuous, and let the trait distribution within a species be $p(\bar{x}, x)$, where $\bar{x}$ is the mean trait value. Then the species-level utilization function $u_s(\bar{x}, z)$ will be the sum of the contributions of all individuals to consuming the resources, i.e.,

$$u_s(\bar{x}, z) = \int_{x_0}^{x_m} p(\bar{x}, x)u(x - z)\,\mathrm{d}x. \tag{3.29}$$

This function is continuous even if the trait distribution $p(\bar{x}, x)$ is not, since the integral of a bounded discontinuous function is continuous. The only case when the original discontinuities in $u(x - z)$ are retained is when $p(\bar{x}, x) = \delta(\bar{x} - x)$, i.e., when all individuals are exactly the same. In reality, most quantitative traits follow a normal distribution (e.g., **?**), where the variance may depend on the mean trait $\bar{x}$:

$$p(\bar{x}, x) = \frac{1}{\sqrt{2\pi\sigma^2(\bar{x})}} \exp\left(-\frac{(\bar{x} - x)^2}{2\sigma^2(\bar{x})}\right). \tag{3.30}$$

The effective, species-level utilization function is then given by

$$u_s(\bar{x}, z) = \int_{x_0}^{x_m} \frac{u(x - z)}{\sqrt{2\pi\sigma^2(\bar{x})}} \exp\left(-\frac{(\bar{x} - x)^2}{2\sigma^2(\bar{x})}\right)\mathrm{d}x, \tag{3.31}$$

which is continuous even if $\sigma(\bar{x})$ is not.

The second smoothing mechanism comes from environmental variability. Even if all members of a given species are perfectly identical, there is an inherent randomness in their individual fates due to the unpredictability of their surroundings. Just as individuals of a

species are not exactly identical, no two seeds of the same size are identical either: one may be a little softer and thus may be opened by a bird with a slightly smaller beak, to give an example. Then, even if for the time being we do assume all individuals of the species to be identical, the discontinuity of the utilization curve will disappear, for the following reason. Let us denote the "raw" utilization function, which now becomes a function of the environment, by $u(x - z, E)$, where E specifies the state of the environment. Moreover, let us assume, as a worst-case scenario, that all individuals are perfectly identical: everyone has trait x. But, since each individual experiences a given environment, the species-level utilization curve will be the normalized sum of the raw curves over all individuals. Since continuous-density models inherently assume very large population sizes, the sum may be thought of as an integral over the probability distribution of E — which, by the logic of the previous paragraph, will smooth out any discontinuities in resource utilization.

Consequently, discontinuous utilization curves are not to be expected in any realistic ecological scenario. Since the emergence of kinked competition kernels is conditional on those discontinuities, it follows that in reality competition kernels are always smooth. kinked kernels emerge when model assumptions are too idealized or simplified. As we have seen, there is a major difference between the behavior of smooth versus nonsmooth models, which suggests siding with the more realistic smooth models when applying ecological theory.

3.6 Discussion

We have considered the effects of kinked competition kernels on species packing and coexistence along a trait axis. Kernels possess a "kink" if they are nondifferentiable when two species have the exact same trait value. It turns out that such kernels are able to produce patterns of continuous coexistence that are not entirely destroyed by model perturbations, in contrast to what one would expect based on limiting similarity arguments. The intuitive explanation for this behavior is the rapid decrease in competition between similar species: nondifferentiability at zero trait difference means that a small change in the trait of one of the species will lead to an immediate linear decrease in competition between them, as opposed to the much slower quadratic decrease of smooth kernels. The mechanism that produces kinked kernels to begin with is the sudden, discontinuous change in the resource utilization functions of the species. We also concluded that such discontinuities are unrealistic and that any real ecological situation would lead to continuous utilization

functions and therefore smooth competition kernels.

Our treatment relied heavily on the Lotka–Volterra equations. Though Lotka–Volterra models have mostly fallen out of favor and have been replaced by more mechanistic models in modern ecological literature, one must not forget that any model may be linearized and brought to a form equivalent to a Lotka–Volterra system near a fixed point equilibrium. Then, as long as the system does not exhibit cycles, chaos, or other complex dynamics, local analysis of the fixed points will lead to the understanding of the global behavior of the model. This justifies having restricted our attention to Lotka–Volterra-type equations.

The argument that kernels decreasing faster around zero niche difference will lead to more coexistence than smooth ones is the generalization of the intuitive argument given by **?**, who were comparing the diversity predicted by a restricted set of kernels. In particular, they were considering the class of kernels $a(x - y) \sim \exp(|x - y|^p)$, which is smooth for $p \geq 2$ but kinked for $0 < p < 2$. In their simulations 200 species were randomly thrown onto a niche axis with fully periodic boundary conditions, then their dynamics was simulated assuming Lotka–Volterra competition. What they found was that, for $0 < p < 2$, species thrown arbitrarily closely on the niche axis could stably coexist, while for $p > 2$ there were always zones of exclusion between prevailing species, i.e., limiting similarity was recovered. This result was interpreted in light of the fact that $p > 2$ kernels are more box-like than $0 < p < 2$ ones, and therefore competition between similars is stronger. The authors' main concern was the analysis of the limiting case $p = 2$ (Gaussian kernel), which lies on the borderline between box-like and peaked kernels. In our parlance, $p \geq 2$ kernels are a subcategory of smooth kernels, while $0 < p < 2$ ones are kinked. Work by the same authors determined that positive definiteness of the kernel is required for the stability of continuous coexistence solutions (**?**), and it so happens that for $p < 0 \leq 2$ the kernel is positive definite, but not for $p > 2$.

Similarly, **?** analyzed the existence and stability of fixed point solutions in the competition-mortality tradeoff model. They pointed out that the competition kernel's discontinuity allows for the coexistence of a continuum of species, but when the kernel is smoothed out, continuous coexistence is impossible. They correctly identified the discontinuity of the kernel as the key property generating continuous coexistence, and also argued that in reality the kernel should be smooth.

These results are all in agreement with ours, but are not the same. We were investigating robustness, not stability: what happens to a given solution if model parameters are perturbed? In the work of **?** robustness of continuous coexistence solutions with the

smooth kernel did not even come up, as they demonstrated that such a solution does not exist in the first place. However, they did not analyze the robustness of the continuous coexistence solution when the kernel is unsmoothed and therefore kinked. In light of our work, they would have found that continuous coexistence is robust (see also **?**). In the case of the work of **?**, they assigned the same r_0 value for all species and stuck to that choice, so the issue of robustness was not investigated. We can now say that they would have found robust continuous coexistence for $0 < p < 2$ kernels and unrobust one for $p = 2$, the Gaussian case. For $p > 2$ the fixed point is unstable and so the issue of robustness does not even arise.

The difference in behavior between smooth and kinked kernels is relevant in the context of the debate over the relative importance of stabilizing vs. equalizing mechanisms (**?**). **?** showed that the invasion growth rate of a species can be approximated as a sum of two terms, as long as the interactions within the community are purely competitive and all species but the invader are at their stationary equilibria. The first ("equalizing") term is always proportional to the difference (or ratio, in discrete time) of the intrinsic rates of growth, while the second ("stabilizing") term depends on the equilibrium densities of the resident species. Without stabilization, two species may only coexist if their intrinsic growth rates are exactly equal under all circumstances — a nongeneric scenario. However, as **?** pointed out, if the intrinsic rates are nearly equal, then even a very slight amount of stabilization will be enough to guarantee long-term coexistence. This seems to suggest that coexistence by virtue of species similarities, as opposed to differences, could lead to stable coexistence: although similar species would only have very weak stabilizing terms, their intrinsic growth rates will also be very similar and so the weak stabilization will still be enough to ensure a positive invasion growth rate for all species. This idea has spurred a body of literature on the coexistence and evolutionary emergence of similar species (**????**).

The concept that species with almost-equal intrinsic growth rates can coexist via relatively weak stabilization is surely uncontroversial. However, the situation is not that simple when the trait-dependence of the two terms is considered. We have seen in Section 3.4 (with the mathematical underpinning in Appendix C) that the equalizing term (difference in r_0) and the stabilizing, frequency dependent term do *not* approach zero at the same rate in general: the former is proportional to the difference in trait, while the latter is proportional to the square of the difference in trait. The stabilizing term is therefore incapable of overcoming differences in r_0 if the species are too similar — *except* when the competition kernel is kinked. For kinked kernels the stabilizing term changes linearly with

trait difference, just like the equalizing term, and so it can compensate for differences in r_0. In conclusion, only models with kinked kernels can allow for the robust coexistence of similar species; for instance, in the work of **?**, only transient coexistence of similar species was possible with a Gaussian competition kernel, but stable coexistence was observed when an extra term was added to the equations that rendered the kernel kinked.

Does the conclusion that models should be smooth mean one should avoid models possessing kinked kernels? As mentioned before, several well-known models exhibit this property, e.g., the hierarchical competition-colonization tradeoff model (**??**), the competition-mortality tradeoff model (**?**), a model of superinfection (**?**; in these three models the kernel is not even continuous), and the tolerance-fecundity tradeoff model (**??**). Despite their nonsmoothness, they do capture important features of the world. In particular, they drive attention to potential coexistence-enhancing tradeoffs which could operate in smooth models as well, although the precise amount of diversity predicted by the two approaches will be different. Smooth versions of these models, along with some consequences of the smoothing (in agreement with our results) are given in **?**. It turns out that the smoothed models are somewhat more inconvenient to handle, both analytically and numerically. Therefore even if nonsmooth models are less realistic, they could be good as a first proxy to assess the consequences of certain assumptions because they are simpler to solve. Perhaps the main lesson to be learned is not that kinked models should be eschewed, but rather that one should be careful not to push the simplifying assumptions too far: when a model like the competition-colonization tradeoff model produces arbitrarily tight species packing (**?**) and even robust continuous coexistence (**?**), we know that this result is just an artifact produced by the kernel and that in reality the kernel is smooth and no robust continuous coexistence is expected.

Of course it is possible to have kernels which, though not kinked in the technical sense, are "very peaked", meaning that their second derivative at zero trait difference is large. Continuous coexistence would be unrobust with these kernels, but still, we would expect their behavior to approach that of kinked kernels. Although we have not looked into the implications of such kernels in a rigorous way, both past results and common sense suggest that the more peaked the kernel is, the tighter species packing it will allow for. For instance, in the case of Gaussian kernels, tightness of packing depends on the competition width (**???**), which in turn is proportional to the kernel's second derivative at zero trait difference. In this way, one would expect the spacing between species to shrink as the kernel gets more and more peaked. Finally, in the limit where the second derivative of the kernel goes

to infinity, the nearest-neighbor distances shrink to zero, i.e., robust continuous coexistence is recovered. Thus, though kinked kernels are unrealistic, it might still be possible to have fairly tight species packing via kernels that are close to being kinked.

Needless to say, the theoretical expectation of limits to similarity may be violated in particular cases for several reasons. One obvious possibility is that the system has not yet reached its equilibrium and so some of the species are still on their way to extinction. Also, it might be that coexistence is maintained through multiple trait axes. If there are several important axes and we concentrate on only one of them, what we see is the projection of all species onto a single axis and depending on how traits map onto regulating factors the distribution of species expected along one trait axis may differ from a spaced pattern. Yet another reason why spacing could be obscured is that metacommunity processes may play a role as well: there is a constant stream of immigrants to a particular site, replenishing those species that are on their way to extinction (**?**). In this case the spatial scale at which the observation is carried out could be too small to see the effects of competition on community structure as a whole. Finally, it is certainly possible that the trait under consideration does not map onto any niche axis, i.e., a linear array of regulating entities. We usually think that the beak size of Darwin's finches corresponds to the size of the food they eat, and since we think of food of a certain size as providing potentially independent regulation from all the other types of food, we may justifiably claim that beak size as a trait is an indicator for niche differentiation. But in other cases such trait differences might not be indicative of adaptation to different regulating factors. The drought-tolerance of plant species coexisting in arid regions does not display limiting similarity, because drought acts as an environmental filtering agent and not as a regulating factor, let alone a whole continuum of them.

Despite these caveats, if spacing is always expected in competitive guilds then work aimed at discovering spacing patterns in data could lead to a better understanding of which trait differences allow for niche differentiation. Apart from the difficulties already mentioned, the problem of discerning limiting similarity from data is complicated by the fact that there are no universal, system-independent limits to similarity (**??**) and that even when one has limiting similarity the spacing between adjacent species need not be uniform (**??**). Discussion of the methodological tools needed to overcome these problems is beyond the scope of this chapter. Empirical as well as methodological research of limits to similarity, however, remains an important direction within community ecology (**?????**), and should remain so in the future.

3.7 Figures

Figure 3.1: Equilibrium patterns produced by a Gaussian competition kernel. The first panel shows the equation and the graph of the competition kernel used; $\Delta x = x - y$. The second panel gives the formula for $n(x)$ and the curves of $n(x)$ and $r_0(x)$ (which can be obtained by substituting the given forms of $a(x - y)$ and $n(x)$ into Eq. (3.7) and performing the integration). The third panel presents what happens to the equilibrium state when $r_0(x)$ is perturbed. We obtained the perturbed equilibrium $\hat{n}(x)$ by first adding a small perturbing function $\eta(x)$ to the original $r_0(x)$ to obtain the perturbed intrinsic rates $\hat{r}_0(x) = r_0(x) + \eta(x)$, then simulating the dynamics via Eq. (3.6) until it reached its stable equilibrium. The function $\Lambda(x)$ involved in the perturbation in panel 3 is defined as $400(1 - |x|)$ for $-1 < x < 1$ and zero otherwise. The argument is multiplied by 400 since this was the number of bins the trait axis was divided into in our simulations — this way the perturbation is effectively point-like, i.e., zero everywhere except at $x = 0.5$. In panels 2 and 3, $r_0(x)$ and $\hat{r}_0(x)$ have been scaled so they would fit on the same plot as the densities.

$$a(\Delta x) = e^{-\frac{\Delta x^2}{2 \cdot 0.06^2}}$$

$$n(x) = e^{-\frac{(x-0.5)^2}{2 \cdot 0.15^2}}$$

$$\eta(x) = 0.05\Lambda(x - 0.5)$$
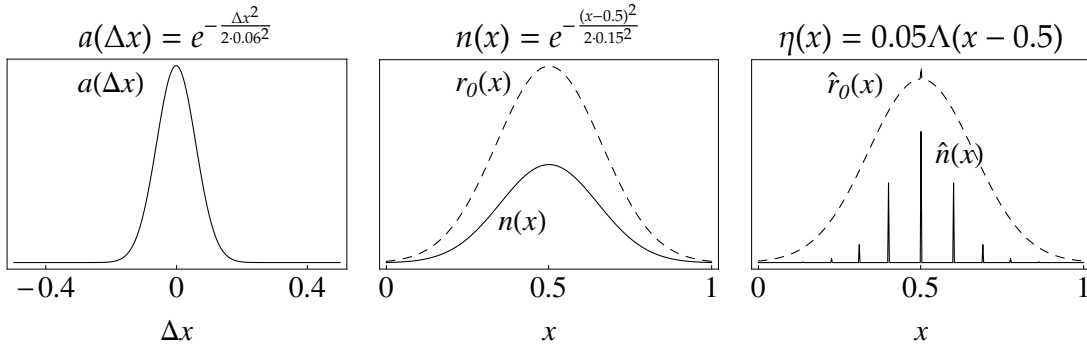
Figure 3.2: Equilibrium patterns produced by smooth nonanalytic competition kernels. Layout and notation and methods as in Fig. 3.1, with four rows instead of one; $u(x) = 1 - |x/0.1|$ if $|x| \leq 0.1$ and zero otherwise; $\Theta(x)$ is the Heaviside unit step function. The four rows present four different examples of continuous coexistence and the coexistence pattern obtained by slightly perturbing the intrinsic rates of growth. Continuous coexistence collapses in all cases following perturbation.

60
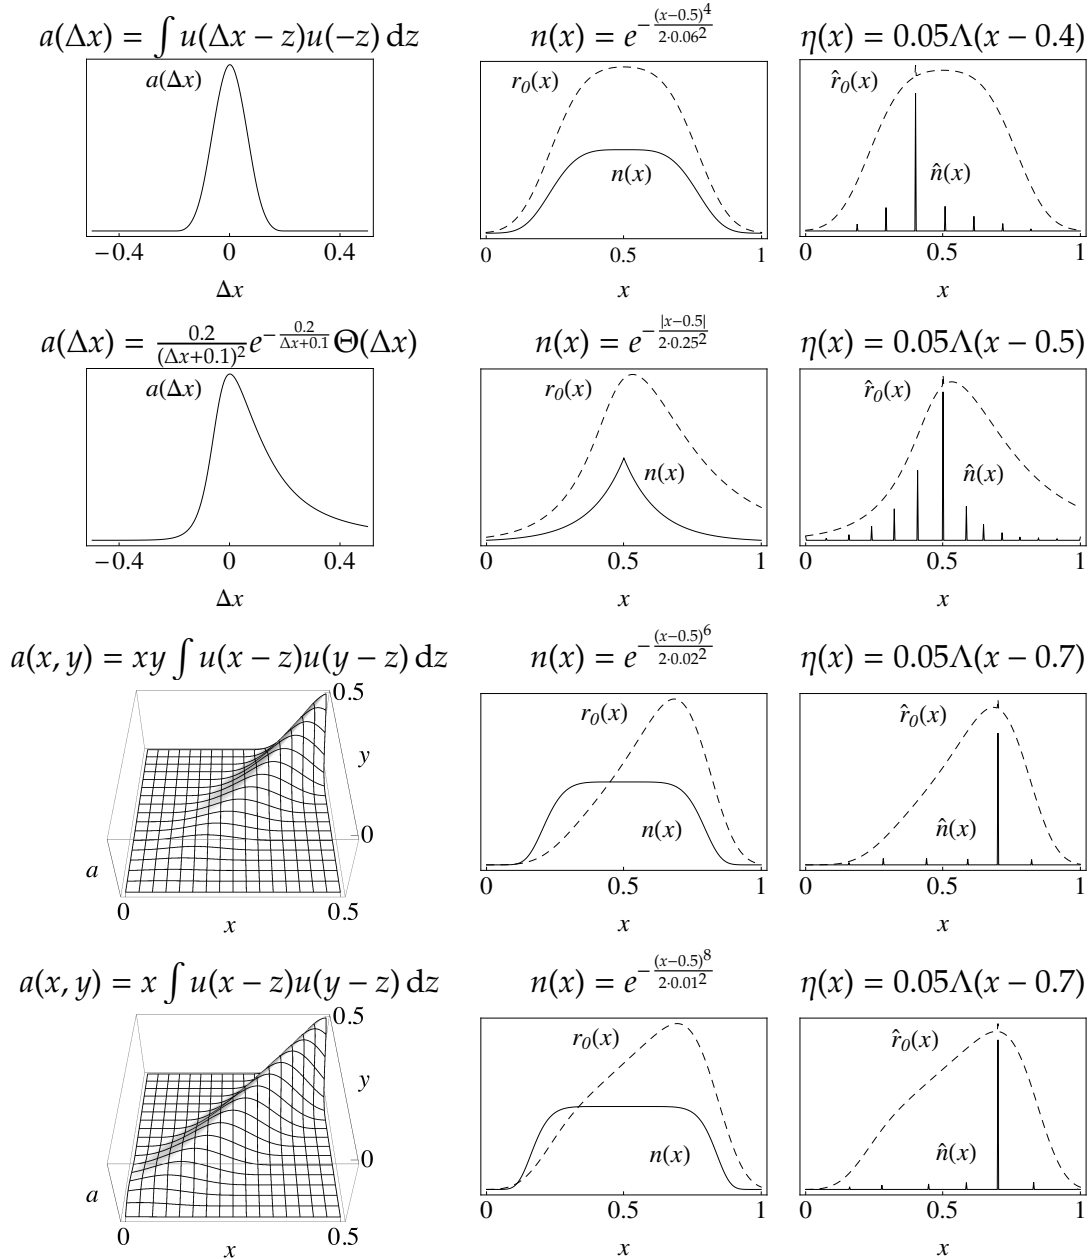
Figure 3.3: Equilibrium patterns produced by kinked competition kernels. Layout, methods, and notation as in Fig. 3.2. Although certain species go extinct following perturbation in all cases, continuous coexistence does not disappear.
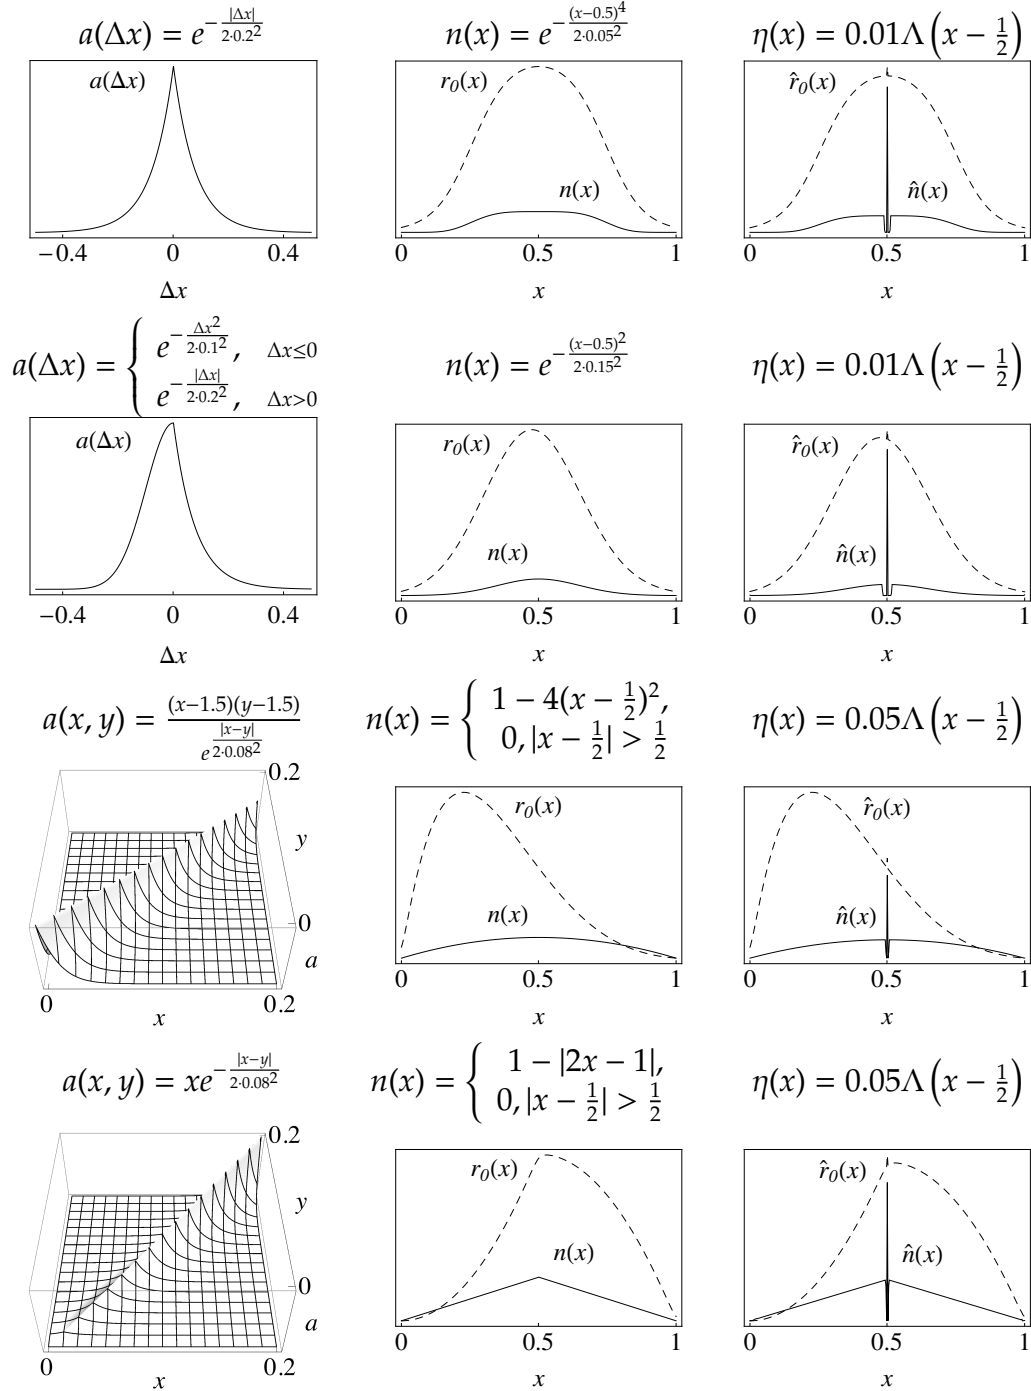
Figure 3.4: The effects of increasing perturbation size on a model with a kinked kernel. The kernel used is $a(\Delta x) = \exp(-|x|/(2 \cdot 0.1^2))$ (its general shape is given by the top left corner of Fig 3.3), and the unperturbed densities are $n(x) = \exp(-(x - 1/2)^{10}/(2 \cdot 0.0083^2))$. Notation is as in the previous figures. Panel A depicts the unperturbed solution. For sufficiently small perturbations (panel B) the equilibrium abundances are altered but no extinctions occur. For larger perturbations (panels C, D and E), some species go extinct, but beyond a well-defined exclusion zone coexistence is just like it was without the perturbation. As the perturbation size increases, the exclusion zone progressively increases until all but one single species are excluded (panel F). Note that this happens when the perturbation size is approximately 10^{10} larger than the original function, i.e., the perturbation is astronomically large compared to the original $r_0(x)$.
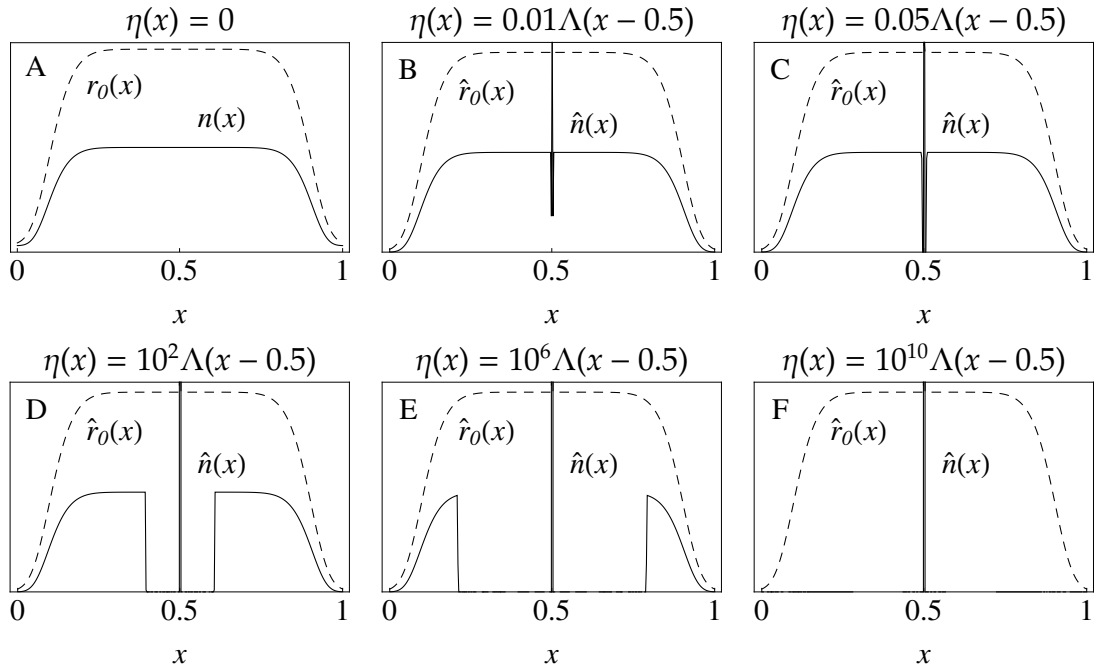
Figure 3.5: Utilization curves of two species with traits x_1 (solid line) and x_2 (dashed line), respectively. For the given box-like utilization function $u(x - z)$, no matter how similar the two species are, there will always be a range of resources (shaded in gray) that are utilized exclusively by only one of them. This leads to the independent regulation of the species and therefore to their coexistence, regardless of how close x_1 is to x_2.
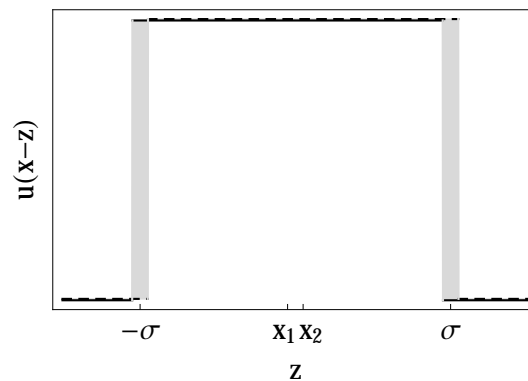
Chapter 4

Revising the tolerance-fecundity tradeoff; or On the consequences of discontinuous resource use for limiting similarity, species diversity, and trait dispersion

Abstract

The recently proposed tolerance-fecundity tradeoff model represents a step forward in the study of seed size diversity in plant communities. However, it uses a oversimplified picture of seed tolerance, with an infinitely sharp threshold: the probability that a seed tolerate a given stress level is either 1 or 0. This invites a revision of the model, presented here. We demonstrate that this simplification has large impacts on model behavior, including altering predictions regarding limiting similarity, raising expected diversity levels, and lessening expected spacing between species along the trait axis. Such dramatic impacts ultimately stem from the fact that a discontinuity in the probability of tolerating a site drastically reduces competition between similar species. This is one example of a class of models with a non-differentiable peak in the competition kernel, which we recently showed is produced by resource use unrealistically modeled as discontinuous, and affects fundamental predictions regarding limiting similarity. This paper illustrates those general results, and offers a revised model of the tolerance-fecundity tradeoff.

4.1 Introduction

In ecologists' never-ending quest to understand coexistence of competing species, tradeoffs are often assumed to be key (**???**). In the context of competition in plant communities mediated by seed size, the most commonly used approach is to postulate

an exchange between competitive ability and fecundity (or colonization ability): smaller seeds are cheap and thus produced in great numbers, whereas larger seeds provide for stronger competitive ability and thus the capacity to displace weaker individuals (**???**). This competition-colonization tradeoff idea has been utilized by several authors since at least the 70s (see **?** and their references). However, the assumption that larger seeds have a fixed competitive advantage over smaller seeds regardless of environmental conditions lacks empirical support (**???**). Motivated by this mismatch between theory and empirical observation, **?** proposed a novel type of tradeoff model of seed size diversity, according to which a tradeoff between fecundity (as measured by number of seeds) and tolerance to stressful conditions such as drought and shade guarantee coexistence of species with seeds of different sizes. Here, the advantage of the large seed over its lower-sized competitor is contingent upon the stress level of the particular site they land on: the big seed is favored only if the stress level is beyond the tolerance threshold of the small seed.

Although representing an improvement over competition-colonization models of seed size diversity that did not contain an empirically supported element of contingency to competitive advantage, the TFM as proposed by Muller-Landau is based on a highly simplified picture of seed tolerance, and effectively, species' use of habitat. In the model, the size of a seed is associated with a tolerance threshold that defines its ability to survive in a given patch: if the stress level in the patch is below that threshold, the seed can fully tolerate the conditions in that patch, and its survival is limited only by competition with other seeds; otherwise, the seed is simply unable to recruit. In other words, each species' seed tolerance switches from 100% to 0% as stress increases past a threshold value, with nothing in between. It is highly unlikely such an infinitely sharp transition would occur in nature, where immanent small amounts of variation in phenotype or environmental conditions experienced within a species would suffice to smooth out the transition.

Muller-Landau acknowledges this simplifying assumption and argues that it would likely be of little consequence for the fundamental coexistence-generating nature of the tradeoff described by the model (**?**). We agree. However, we point out that the assumption has potential consequences to the level of coexistence that is possible. In this scenario, no matter how similar the seed sizes of two given species, there will always be patches perfectly accessible to species A, but utterly out of reach for species B. In other words, niche overlap between two species is always limited, however similar they may be. Thus, though the assumption may not be essential to guarantee coexistence per se, it may have significant impacts on how much coexistence is allowed, and hence may impact how much

diversity can be explained in practice by the tolerance-fecundity tradeoff.

It is the purpose of this article to demonstrate the strong implications of the discontinuity in the tolerance function onto the tolerance-fecundity tradeoff model. First, we highlight the implications of this assumption for predictions regarding coexistence of a continuum of arbitrarily similar species (*continuous coexistence*, or *tight packing* in the real-world case of a finite number of species) and the related concept of limits to similarity, which itself has had a broad influence in ecology and evolution (**?????**). We then examine how this assumption bears on predicted levels of diversity and trait dispersion. To do so, we develop a revised model where the assumption is dropped, and examine how predictions are affected by this change.

The notion that continuous coexistence is not possible in nature began with the work of MacArthur and Levins (1967), who showed that, in a simple model of Lotka-Volterra competition, stable coexistence of a set of three species requires some minimum spacing between them on a trait axis. Despite later work showing that in fact continuous coexistence can be produced in that simple model (**?**), ultimately it was proven that, for a wide class of competition models, any possible continuous coexistence is not robust to small changes in parameters, and in this sense these models always predict system-specific limits to the similarity of stably and robustly coexisting species (**?**). Although technically possible given that the proof does not extend to all possible competition models, examples of models that actually predict robust tight packing have not been presented previously.

Here we show that Muller-Landau's tolerance-fecundity tradeoff model predicts robust continuous coexistence. Upon removal of the discontinuity in the model, however, the model no longer admits continuous coexistence at all. Hence the original version sees no fundamental limits to the similarity of species coexisting through this mechanism, whereas the revised version predicts such limits. Further, we find that the revised model predicts lower diversity and more detectable even spacing among species. Our demonstration that the shape of the tolerance function, a central feature of the model, has a profound and consistent influence on predicted levels of diversity and trait dispersion adds to Muller-Landau's contribution, and provides guidance to future empirical work regarding the link between this shape and the degree to which maintenance of observed diversity and dispersion in seed size can actually be attributed to the tolerance-fecundity tradeoff.

We note that the considerable differences in predictions found here are an example of more general results from our recent work on the relationship between continuity in resource use and robust continuous coexistence (**?**). We thus argue that our results extend

to any model with this offending property, of which Muller-Landau's model is by no means
the sole example. The general results we illustrate here using the tolerance-fecundity
tradeoff model as an example suggest caution when constructing models to study limiting
similarity, diversity, and patterns of trait dispersion.

4.2 The tolerance-fecundity tradeoff model

4.2.1 Original formulation

The dynamics of the tolerance-fecundity tradeoff model (**?**) are as follows. The
community is assumed to be saturated with adult individuals which pre-empt establishment
of propagules. Every time step, however, some individuals die, making sites available
to propagules and enabling competition among them – in fact, in this model competition
only occurs at the establishment phase. The number of seeds of a given species landing
on a site is assumed to be proportional to the product of the fraction of sites occupied
by that species, p_k, and its fecundity, f_k (i.e., no dispersal limitation is being considered –
Muller-Landau also examined chance dispersal limitation through a Poisson-distributed
seed rain and found no qualitative change in model behavior). Species are ordered by
fecundity level from f_1 to f_n, with species 1 being the least fecund (and most tolerant) and
species n being the most fecund (and least tolerant). In a fraction h_k of the sites, species
k can fully tolerate the stress level, and in the remaining fraction it cannot tolerate it at
all. Then given all the seeds that reach and tolerate the same site, a lottery decides which
seed eventually recruits into an adult (Fig. 4.1A). Given those conditions, the population
growth rate of species k can be written as

$$\frac{dp_k}{dt} = m\left(f_k \sum_{i=k}^{n} \frac{h_i - h_{i+1}}{\sum_{j=1}^{i} f_j p_j} - 1 \right) p_k. \tag{4.1}$$

The first term on the right-hand side corresponds to recruitment of new individuals,
and the second term reflects mortality. Both terms are proportional to species abundance p_k
and mortality rate m (assumed constant across species for simplicity). The recruitment term
is proportional to the species' fecundity f_k, and total recruitment is a sum over recruitment
in colonizable patches, which include all the patches colonizable by species coming further
down the hierarchy. For each fraction of patches colonizable by one species i ($i > k$) but
not the subsequent species, $h_i - h_{i+1}$, recruitment there is weighted by the total number of

seeds that could potentially colonize, $\Sigma f_j p_j$. Notice that $h_n < h_{n-1} < \cdots < h_1$ and $h_{k>n} = 0$ are assumed.

Note that Muller-Landau postulated an unambiguous association between a given fecundity value (defined as the number of seeds produced by each individual of a certain species) and a unique seed size. In other words, she assumed no noise in the relationship between seed size and species fecundity. This enabled her to focus on fecundity as the trait axis but draw from it conclusions about maintenance of seed size diversity. We will keep this simplified scenario, and regard the fecundity axis as the trait axis for the remainder of this article.

As just described, in Eq. 4.1 the recruitment term is a sum over the prevalence of different types of sites, categorized by the number of species that can tolerate them. The fact that the recruitment term can even be organized this way is the telltale sign of the simplified picture of seed tolerance used in the model. Note that no matter how similar two species i and $i + 1$ are in their seed size and fecundity, species i has unfettered access to a proportion of the sites that species $i + 1$ has no access whatever, namely $h_i - h_{i+1}$ of them. If seeds have some likelihood of tolerance of a given stress level other than 100% or 0%, then that must be factored into a species' probability to colonize a given site, and subsequently sites cannot be sorted by who can or cannot tolerate them as above; instead the stress level at each site must be considered.

4.2.2   Revised formulation

We now describe our revised version of the tolerance-fecundity tradeoff model. We rewrite the model from first principles, maintaining the same basic dynamics as Muller-Landau's model, except that we now allow for species' tolerance of a site to be a continuous function of the stress level at that site (Fig. 4.1B).

We first note that for mathematical convenience we consider a pool of species forming a continuum, with all possible fecundity values in the range $[f_0, f_m]$ represented, because we will later study the potential for coexistence of arbitrarily similar species ('continuous coexistence'). Although such continuum is a mathematical abstraction that can't occur in nature, where only a finite number of species are available, we want to know how tightly packed that finite number of species can be, or correspondingly whether there are limits to the similarity of a finite set of coexisting species. Considering the possibility of continuous coexistence enables us to see whether such limits exist.

Let us divide our total area into sites (each colonizable by at most one individual),

characterized by their stress level σ. Let c be the total number of sites. Let $n(f)$ be the number of individuals with fecundity f. Then

$$\frac{dn}{dt}(f) = \underbrace{(\text{\# of seeds that recruit})}_{fn(f)\int Q(\sigma,f)\,d\sigma} - \underbrace{(\text{death})}_{m(f)\,n(f)}. \tag{4.2}$$

Here $m(f)$ is the mortality rate, and $Q(\sigma, f)$ is the probability that a given seed from an individual of fecundity f germinates in a site of stress level σ. Following Muller-Landau, we do not include displacement terms: once an individual occupies a site, it cannot be displaced by another until it dies. This reflects the assumption that competition occurs only at the seed stage.

To determine $Q(\sigma, f)$, we need to examine the probability of a given seed to arrive at an available site, and then consider its probability of successfully recruiting. Let $T(\sigma, f)$ denote the probability that a seed produced by an individual of fecundity f arriving at a site of stress σ successfully recruits there. Let $c(\sigma)$ be the number of sites of stress level σ (thus $c = \int c(\sigma)d\sigma$). The total number of colonizable sites of stress σ is equal to $c(\sigma)$ minus the number of occupied sites, or $(c(\sigma) - \int n(\sigma, f)\,df)$, where $n(\sigma, f)$ is the number of individuals of fecundity f occupying sites of stress σ (note that $n(f) = \int n(\sigma, f)d\sigma$). Therefore, the a priori probability of a given seed landing on any of the empty sites of stress level σ per unit time is $(c(\sigma) - \int n(\sigma, f)\,df)/c$. However, the seed has to survive the stress level of the site, so this has to be weighted by $T(\sigma, f)$. Additionally, we need to consider the probability that, once this seed gets there and is able to survive the local stress level, it wins over the competition with rival seeds. This probability is just the inverse of the total number of seeds arriving at the site and able to recruit there as well, namely $(\int f\,n(f)T(\sigma, f)\,df)/c$. We therefore arrive at the stress-dependent probability of colonization:

$$Q(\sigma, f) = \frac{T(\sigma, f)\Big(c(\sigma) - \int_{f_0}^{f_m} n(\sigma, x)\,dx\Big)/c}{\Big(\int_{f_0}^{f_m} y\,n(y)T(\sigma, y)\,dy\Big)/c} = \frac{T(\sigma, f)\Big(c(\sigma) - \int_{f_0}^{f_m} n(\sigma, x)\,dx\Big)}{\int_{f_0}^{f_m} y\,n(y)T(\sigma, y)\,dy},$$

where f_0 and f_m are, respectively, the minimum and maximum fecundity within the community.

We should have an intuitive idea of the functional form of the function $T(\sigma, f)$: the species with fecundity f should be able to tolerate stress levels below a certain threshold $s(f)$ fairly well, but should be less successful at levels above that threshold. $T(\sigma, f)$ should

69

tend to one at low σ and zero at high σ, and the transition between these extremes should center at $s(f)$.

Putting together what we have so far, our model reads

$$\frac{dn}{dt}(f) = \left(f \int_{\sigma_0}^{\sigma_m} \frac{T(\sigma, f)(c(\sigma) - \int_{f_0}^{f_m} n(\sigma, x)\, dx)}{\int_{f_0}^{f_m} y\, n(y) T(\sigma, y)\, dy} \, d\sigma - m(f) \right) n(f), \tag{4.3}$$

where σ_0 and σ_m are respectively the lowest and highest stress levels in the patch. It is convenient to convert the absolute population abundances to proportions. This is done by defining $p(f) = n(f)/c$, so that $\int p(f) \leq 1$. Eq. 4.3 becomes

$$\frac{dp}{dt}(f) = \left(f \int_{\sigma_0}^{\sigma_m} \frac{T(\sigma, f)(c(\sigma)/c - \int_{f_0}^{f_m} p(\sigma, x)\, dx)}{\int_{f_0}^{f_m} y\, p(y) T(\sigma, y)\, dy} \, d\sigma - m(f) \right) p(f). \tag{4.4}$$

The model can be simplified by assuming saturation: each site is occupied and will only become available through deaths. In this case, the proportion of sites of stress σ being made available at any moment, $c(\sigma)/c - \int p(\sigma, x)\, dx$, is expressible as the number of deaths that occurred, $\int p(\sigma, x) m(x)\, dx$. In the saturated case, therefore, we have

$$\frac{dp}{dt}(f) = \left(f \int_{\sigma_0}^{\sigma_m} \frac{T(\sigma, f) \int p(\sigma, x) m(x) dx}{\int_{f_0}^{f_m} y\, p(y) T(\sigma, y)\, dy} \, d\sigma - m(f) \right) p(f). \tag{4.5}$$

We can further simplify things by assuming uniform mortality across species: $m(f) = m$. We then get

$$\frac{dp}{dt}(f) = m \left(f \int_{\sigma_0}^{\sigma_m} \frac{r(\sigma) T(\sigma, f)}{\int_{f_0}^{f_m} y\, p(y) T(\sigma, y)\, dy} \, d\sigma - 1 \right) p(f), \tag{4.6}$$

where $r(\sigma) = \int p(\sigma, f)\, df$ is the proportion of occupied sites of stress level σ within the full habitat (which in the saturated case is the total proportion of sites of level σ).

In words, Eq. 4.6 tells us that the dynamics of a species labelled by its fecundity f is determined by a recruitment term and a mortality term, both of which are proportional to the mortality given our saturation assumption. Mortality is being assumed a species-independent parameter, and recruitment is the sum of successful recruitment in each class of sites (defined by their stress level σ). Eq. 4.6 is our revised tolerance-fecundity tradeoff

model.

In Appendix F, we show that Muller-Landau's original formulation of the model can be obtained from our revised formulation if to each species f there is an associated stress threshold, $s(f)$, and $T(\sigma, f)$ is defined as the unit step function $\Theta(s - \sigma)$, defined as 1 if $\sigma \leq s$ and 0 if $\sigma > s$. Importantly, this characterizes a discontinuous transition in stress tolerance (Fig. 4.1C). Clearly, this state of affairs is a caricature of nature: site colonization by plants mediated by seed size is not truly expected to rely on stress thresholds that determine absolutely the fate of the seed (**?**).

If, conversely, we define $T(\sigma, f)$ as a continuously varying function, then sites of increasing stress levels present correspondingly increasing levels of difficulty to a seed. Habitat (resource) partitioning between two species becomes increasingly similar with species similarity (Fig. 4.1D). Competition is presumably much stronger in this scenario. In order to gauge the consequences of relaxing Muller-Landau's assumption to important predictions in the model such as levels of species diversity and patterns of species trait distribution, below we compare results from simulations of the model using both the step-like $T(\sigma, f)$ and an alternative continuous form.

4.3 Comparisons between continuous and discontinuous tolerance function

4.3.1 Tight packing

If, as in **?**, $T(\sigma, f)$ is chosen to be a step-like function as above, then the model is greatly simplified, and the fixed points and their stability can be determined analytically. Of particular interest, we can find solutions $p(f)$ where species of arbitrarily similar traits can coexist – a scenario which here will be called *tight packing*, and is also known as *continuous coexistence* (strictly speaking, tight packing as referred to here should be defined as the *potential* for continuous coexistence, as a finite number of species does never truly make up a continuum). In equilibrium, $dp(f)/dt$ in Eq. 4.6 must vanish for all f. For tight packing, $p(f) \neq 0$ for all f in $[f_0, f_m]$. Thus, for tight packing, the expression in parenthesis in Eq. 4.6 must vanish identically. If we define the quantity $h(f)$, representing the proportion of land whose stress level is equal to or less than the threshold $s(f)$ of species f (see Appendix F),

$$h(f) = \int_0^{s(f)} r(\sigma) \, d\sigma,$$

then a necessary condition for tight packing is (see details in Appendix G)

$$fh''(f) + 2h'(f) + p(f) = 0, \tag{4.7}$$

where the prime is standard shorthand notation for differentiation with respect to a function's argument. Mathematically, the tradeoff consists in imposing that $h(f)$ be a monotonically decreasing function. There are many acceptable functions $h(f)$ for which the solution to this equation $p(f)$ is of the tight packing kind. Thus, tight packing is supported in this model. Although we do not perform formal stability analysis here, we note that all tight packing solutions we tried proved dynamically stable in our simulations.

Previous work has shown that for a wide class of competition models, tight packing, when present, is fragile to changes in model parameters (**??**). Fragile states are not expected in nature since fluctuations in parameters are inevitable in natural systems. To our knowledge, robust tight packing is unheard of in the theoretical literature; the slightest change in parameters suffices to destroy tight packing entirely, in the sense that after such perturbation no continuous range of species remains. Nevertheless, we show with simulations that this is not the case here. Given tight packing solutions, we perturb $h(f)$ at either a single fecundity value or many, and observe that the effects of such perturbations remain close to the perturbations (single-point case depicted in Fig. 4.2). Tight packing proves robust to parametrization changes.

In contrast, if the tolerance function loses its discontinuity, the model's predictions regarding tight packing changes qualitatively. In fact, we provide proof in Appendix G that there is no tight packing solution to Eq. 4.6 if the tolerance function is analytic [1] – a mathematical result that should apply generally to any smooth tolerance function and could be argued to hold for more general conditions as well, such as non-uniform mortality (see Appendix G).

4.3.2 Species diversity

When tight packing is not allowed (or not robust), we say there is *limiting similarity* to coexisting species. We now ask how the model with a continuous tolerance function differs from Muller-Landau's discontinuous version in terms of diversity when limiting similarity occurs in both formulations. To find out, we perform the following test: starting with 100 species of equal initial abundances uniformly distributed in the fecundity range

---

[1]An analytic function is a differentiable function that converges to a power series of its argument.

[0,1] and a given parametrization $h(f)$, we simulate the model dynamics until equilibrium is reached. We then take species richness and evenness (defined as the ratio between the obtained Shannon diversity index and its maximum possible value given the number of species present) and compare these diversity metrics between the discontinuous and the continuous case. This test is replicated many times, each for a different parametrization.

The parameters $h(f)$ in each run are determined randomly as follows: we draw 100 real numbers between 0 and 1 from a uniform distribution, spline-interpolate them (cubic spline), calculate their partial cumulative sums, normalize these by the total cumulative sum, and then subtract them from 1 – thus generating a cascade of numbers in decreasing order, ending in 0. Each value is then assigned to each of the species in order, as the proportion of sites below their stress threshold. This set of steps was taken to ensure that the parametrization would be random and smooth at the same time, with the intent of reproducing a typical real-world case and avoiding results that depend on any special parametrization. The dynamics is completely deterministic, meaning that a certain set of initial species abundances and site-stress parametrization $h(f)$ uniquely determines the outcome. Thus by averaging out the parametrization, we get an overall comparison of diversity levels across model types. We start with a species-rich state with evenly spaced species at equal abundances to represent an initial community in a tight packing state.

Results are shown in Fig. 4.3A and 4.3B. A very clear and strong distinction in diversity levels between the two models is observed under both metrics: the model with continuous tolerance function almost invariably allows for considerably less diversity in the equilibrium community.

We note that in our simulations species are not allowed to evolve and there is no immigration. The addition of new types can bring in better competitors that could drive out hitherto coexisting inferior types (**?**). Thus it is likely that, should mutation and immigration be implemented, diversity would come out even lower in the continuous formulation, while the discontinuous formulation would show less change, as that scenario typically sees pockets of tight packing where all species can coexist. The results reported here should thus be seen as conservative estimates of the disparate diversity levels across these two formulations of the tolerance-fecundity tradeoff.

It should be noted that the discontinuous case is a limit of the continuous case. For instance, our expression of choice for the continuous tolerance function used in the tests, $T(\sigma, f) = \frac{1}{2}(1 + \text{erf}[\nu(s(f) - \sigma)])$, approaches the step function as $\nu \to \infty$. Thus, it behooves us to check how diversity levels in the continuous case behave as the limit

is approached. Fig. 4.3C shows a progression of species richness in the equilibrium community as a function of the parameter v. The curve rises to no limit, again confirming the tendency towards higher diversity as the tolerance function becomes steeper and nears the discontinuity.

4.3.3 Trait dispersion

Finally, we examine how the models differ regarding the distribution of species along the trait axis. Recently, community ecologists have been interested in finding evidence for the role of species differences in allowing for competitive coexistence, and employ several different metrics of trait dispersion to that purpose (**????**). Here we focus on one of such metrics, the coefficient of variation in adjacent-neighbor trait distances (**?**). For each of the simulated runs, we take the differences in fecundity between consecutive (adjacent) neighbors on the trait axis. The coefficient of variation is then defined as the ratio between the standard deviation and the mean of such distances. A value significantly lower than expected by the null hypothesis implies overdispersion – species are more separated than expected by chance – and the opposite indicates clustering – species clump together.

The test is perfomed as follows: for each of one hundred simulation runs, we set up model parametrization with a random small perturbation around $h(f) = exp(-f)$ and run the continuous and discontinuous versions to equilibrium (this is a different approach to parametrization than used in the tests above, but was necessary because completely random $h(f)$ tends to yield very few species in the continuous case, which brings in power issues). We then take the corresponding coefficient of variation of adjacent-neighbor distances along the fecundity axis (henceforth referred to simply as CVAND for brevity) for each and establish the probability that these values would be obtained by chance (p-value) by comparing with a pool of null CVANDs. We end up with a distribution of one hundred p-values for each formulation of the model, which are summarized in the box plots on Fig. 4.3D.

The pool of null CVANDs is generated and used to arrive at p-values as follows: for every number between 3 and 99, we randomly draw that number of species from the pool of 100 species of the initial community used in the simulations, and take the CVAND. This is repeated 100,000 times, thus obtaining a pool of null CVANDs for all richness values between 3 and 99. Then, for each run of each model, we compare the resulting CVAND with the corresponding pool of null CVANDs (i.e., that with the same number of surviving species). The proportion of times (out of 100,000) that the null CVANDs are lower than

74

the observed value is the p-value, representing the probability that randomly generated CVANDs be lower than the observed one. As low CVANDs indicate overdispersion, this is in effect the probability that a random assemblage would be at least as much overdispersed as the observed assemblage.

Upon comparing the box plots for the discontinuous and continuous versions (Fig. 4.3D), the distinction in trait dispersion becomes clear. Although both tend towards overdispersion in that both result in median p values less than 0.5, the pull of the continuous model is much stronger. In fact, 83% of the runs in the continuous model are within the $\alpha = 0.05$ significance level for overdispersion, compared to 4% in the discontinuous case. Results do vary quantitatively with different parameter implementations, but importantly the two formulations consistently lead to quite disparate trait dispersion regardless of how we parametrized the models. We should note that while different metrics could exist that may be better suited to capture the spread in fecundity among surviving species than the coefficient of variation of adjacent-neighbor distances, our test suffices to capture the dramatic differences in trait dispersion across the two model formulations under analysis here.

In order to further ensure generality of our results, we performed the tests described above for different initial conditions (namely, more species and unevenly spaced initial species). Indeed, although outcomes vary quantitatively, the continuous model invariably produces less diverse communities and shows stronger tendency for overdispersion than the continuous model.

4.4 Relation to competition kernel and generalization to other models

Our comparisons indicate that a discontinuity in the tolerance function has a strong, qualitative impact on model predictions regarding the distribution of species along the fecundity (or seed size) axis. This phenomenon is in fact an example of a more general result recently shown by us (**?**) regarding the influence of a non-differentiable sharp peak in the competition kernel on model predictions of competitive outcomes.

The *competition kernel* is the function that determines the level of competition between two species given their traits. It can be defined as $a(f_1, f_2) = \frac{\delta R(f_1)}{\delta p(f_2)}$, where $R(f) = \frac{1}{p(f)} \frac{dp}{dt}$ is the per capita growth rate of species f. (Note that alternative definitions are also in use; see, e.g., **?**. The definition we employ is the most pertinent to our analysis, and has the added benefit of, when applied to Lotka–Volterra models, being equivalent to the collection of all

the competition coefficients.) In Appendix H we show that the competition kernel in our tolerance-fecundity tradeoff model (Eq. 4.6) has a point of non-differentiability – a *kink* (Fig. 4.1E, 4.1F) – at its peak if and only if $T(\sigma, f)$ has a discontinuity. In our recent work (**?**), we show that coexistence of arbitrarily similar species is robust to changes in parameters when the competition-kernel is kinked, and nonexistent or fragile when the competition kernel is smooth (no kinks). The results shown here for tight packing in the TFM are in good accord with that generalization. Above and beyond this, here we have shown for the TFM additional model behavior that is strongly affected by a discontinuous tolerance function, namely species diversity and patterns of spacing between species on the trait axis.

When the kernel is expressible in terms of the overlap of resource utilization functions (**?**), kinked competition kernels arise from discontinuities in those functions (**?**). More generally, competition kernels can be written in terms of functions describing species' sensitivity to and impact on *regulating factors* (i.e. factors that influence and are in turn influenced by the population sizes of species; **?**), and we show in **?** that kinked competition kernels arise from discontinuities in those functions. Clearly, the tolerance function $T(s, f)$ is related to these functions, and a discontinuity in the former must be tantamount to a discontinuity in the latter. Informally it is apparent that the tolerance function is reflective of resource use, where the resources can be seen as patches of a given stress level.

To further make the point that the impacts of non-differentiability in the competition kernel extend beyond the particulars of any given model, in Appendix I we provide the same analysis shown above for a different model of seed size diversity which turns out to have a kinked kernel: the hierarchical competition-colonization model (**???**). In this model, the non-differentiability arises from the assumption that the hierarchy is absolutely strict, with a clearly dominant competitor arising out of any two species considered, no matter how similar they are in colonization ability. Results of the analysis are much alike: when the model's original non-differentiability is removed, much lower levels of diversity are obtained, tight packing is no longer supported, and tendencies for species overdispersion are much stronger. We note that **?** showed previously that predicted diversity dramatically decreases when strictly hierarchical competition is relaxed for a similar mechanism, namely the competition-mortality tradeoff. Their finding can be now understood in the context of the severe consequences of a kinked kernel to model behavior.

4.5 Discussion

We have examined the importance of relieving Muller-Landau's tolerance-fecundity tradeoff model of the discontinuity in resource use that it previously contained. Our work on this model can be viewed as a case study on the consequences of modeling *resource use* (interpreted broadly as described above) as discontinuous. We can understand the general consequences of discontinuous resource use analogously to our understanding of the importance of the tolerance function: if resource use is discontinuous, then two species, no matter how close in traits, stand apart in their use of resources. Hence competition among similar species is relatively lax compared to models with continuous resource utilization, which explains the greater permissivity to coexistence in the former.

We would like to impress upon the reader the unnaturalness of the discontinuity in tolerance level. For one, any small variation in intraspecific trait expression or environmental conditions, which are inevitable in any biological system, suffices for this theoretical scheme to break down (**?**). More importantly, the infinitely sharp transition is hopelessly unrealistic because ultimately in nature there is no such thing as a continuous change in the cause bringing about a discontinuous change in the effect.

Muller-Landau acknowledged the discontinuous tolerance regime as unrealistic, but argued that this assumption was not central to her demonstration that this tradeoff can generate coexistence. Nevertheless, one cannot ignore the implications of this assumption for the amount of diversity that the mechanism can explain. All else being equal, the smoother the tolerance function is, the lower the number of species that will coexist in this model. As is the case with any model of coexistence, it remains an empirical question whether conditions in nature are such that this tradeoff actually explains observed diversity. Our work here highlights the importance of the tolerance function for assessing the existence and potential importance of the tolerance-fecundity tradeoff.

We do not purport to have purged Muller-Landau's model of all its unrealistic simplifications. Many others remain, such as no explicit spatial structure, uniform mortality, no dynamics to stress levels, and no stochasticity. One must keep in mind that any model is a limited representation of reality and must perforce rely on helpful, if simplistic, assumptions. Of course the trouble is that there is no general way to know beforehand the degree to which any given assumption impacts the quality of the model as a representation of the real world. Here we showed that discontinuous resource use has a very strong impact in the tolerance-fecundity tradeoff, and connected this result with the more general impact

of modeling resource use as discontinuous on predictions regarding limits to similarity, species diversity, and patterns of trait dispersion.

In addition to cautioning against modeling resource use discontinuously, our study presented and at the same time cast aside a potential challenge to the widely influential theory of limiting similarity. We showed that a specific mechanism of coexistence seemingly indicates that tight packing of species can be robust, but eventually rooted this unexpected prediction to one of the simplifying and unrealistic assumptions of the model.

4.6 Figures

Figure 4.1: **A**: Cartoon representation of the distribution of seeds given a gradient in habitat stress according to Muller-Landau's assumption of all-or-nothing tolerance. Three seed sizes are represented. Density in seed numbers across habitat represents tolerance levels. The smallest seeds (light grey) occur with uniform density (constant tolerance) through sites with increasing stress levels, up to a threshold – marked by the gray line – beyond which no small seeds are found (zero tolerance). Intermediate-sized seeds (red), occur with uniform density until their own stress threshold is reached – red line – beyond which no medium-sized seeds occur. Large seeds (green) occur uniformly throughout the patch, as its particular threshold is never reached. **B**: Analogous representation in the case where the tolerance function varies continuously. Here, all seeds occur in increasingly lower densities as stress levels increase, and there is no sharp threshold separating sites where seeds occur from sites where they are absent. **C**: When tolerance curves of different species are overlaid, it becomes clear that given any pair of species, no matter how similar, there is always a range of stress levels (resources) that is at once perfectly accessible to the more tolerant of the pair and absolutely out of reach to the other. This range is indicated in the figure by the black bar. **D**: When the tolerance function is continuous, that range disappears. Instead, the way arbitrarily similar species use resources becomes arbitrarily similar, thus making for higher competition than in the previous case. **E**: Illustration of the competition kernel $a(f, f')$ when the tolerance function is discontinuous, showing an infinitely sharp ridge at $f = f'$. A 2D projection of this surface would show a curve with a kink at its peak. **F**: When the tolerance function is continuous – here modeled as $T(\sigma, f) = \frac{1}{2}(1 + \text{erf}[v(s(f) - \sigma)])$ – the ridge is absent, and the surface is smooth across its domain.

Caption on previous page
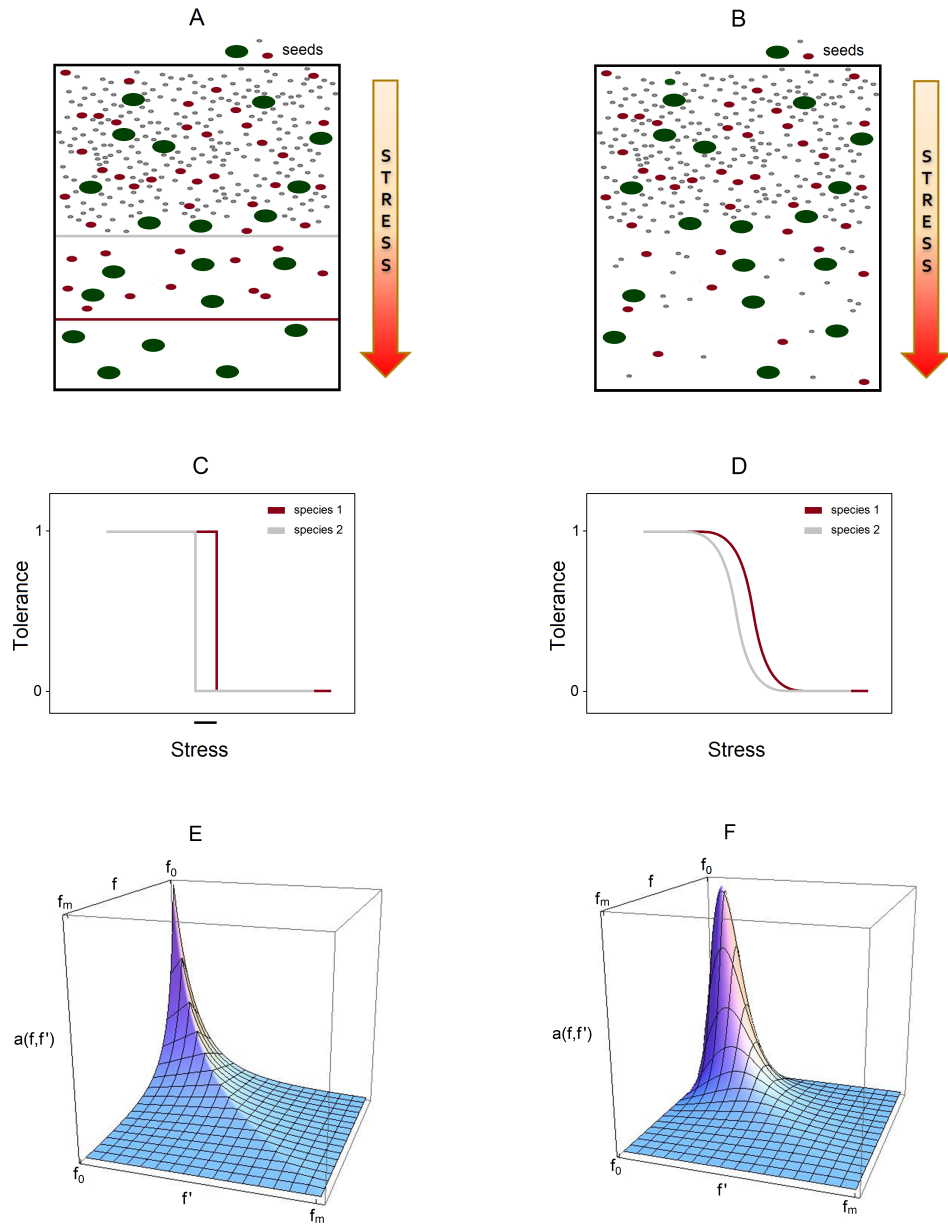
Figure 4.2: Robust tight packing. The gray solid line shows one possible solution of Eq. 4.7, $p(f) = 3f^2$ (gray line), obtained when $h(f) = 1 - f^3/4$ (black solid line). The dotted line shows what happens to the community upon a small local disturbance at $h(f = 5)$. The abundance of species $f = 5$ is greatly affected, and so are those of its immediate neighbors, which go extinct. The remainder of the community, however, is left intact. In the end, tight packing survives the disturbance. This is in striking contrast with known results for a great variety of models (**??**).
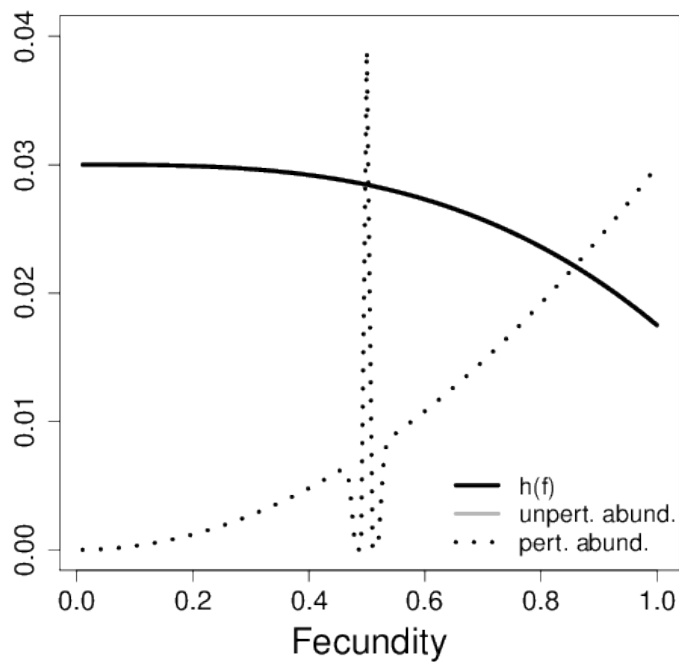
Figure 4.3: **A**: histogram with the distribution of richness R ratios between runs of the discontinuous and continuous versions of the model ($R_{discont}/R_{cont}$)for 1,000 runs. If the continuity of the tolerance function does not significantly affect diversity outcomes, we should obtain a distribution centered around $R_{discont}/R_{cont} = 1$. Results, however, are strongly skewed towards $R_{discont}/R_{cont} > 1$, indicating that the discontinuous formalization consistently leads to a higher species count in equilibrium (see main text for explanations). **B**: histogram showing the distribution of evenness E ratios ($E_{discont}/E_{cont}$) for the same runs as in A. Again, results deviate from the null hypothesis ($E_{discont}/E_{cont} = 1$); the discontinuous version of the model tends to produce more evenly distributed species abundances in equilibrium. **C**: trend in species richness with the steepness of the continuous tolerance function. As the steepness parameter ν in the function $T(\sigma, f) = \frac{1}{2}(1 + \text{erf}[\nu(\sigma(f) - \sigma)])$ increases while all other parameters are kept fixed, so does the number of species in equilibrium after the community reaches equilibrium, starting with 100 species uniformly distributed in the [0,1] fecundity range. This indicates that diversity grows as the continuous formulation approaches its discontinuous limit. Inset shows how the tolerance function looks like for two examples, $\nu = 5$ and $\nu = 20$. **D**: comparison of the distribution of coefficients of variation between adjacent-neighbor distances (CVAND) throughout 499 runs of each version of the model, parametrized as described in the main text. The CVAND gives an indication of how closely surviving species distribute themselves along the trait axis in equilibrium. We see a striking distinction between results for each model: although both tend towards overdispersion (median below 0.5), the pull of the continuous model is much stronger. In addition to a much narrower scatter, the continuous version yields a median p-value of 0.004, in contrast with 0.286 in the discontinuous case. 83% of the runs in the continuous model are within the $\alpha = 0.05$ significance level for overdispersion, compared to 4% in the discontinuous case.

Caption on previous page
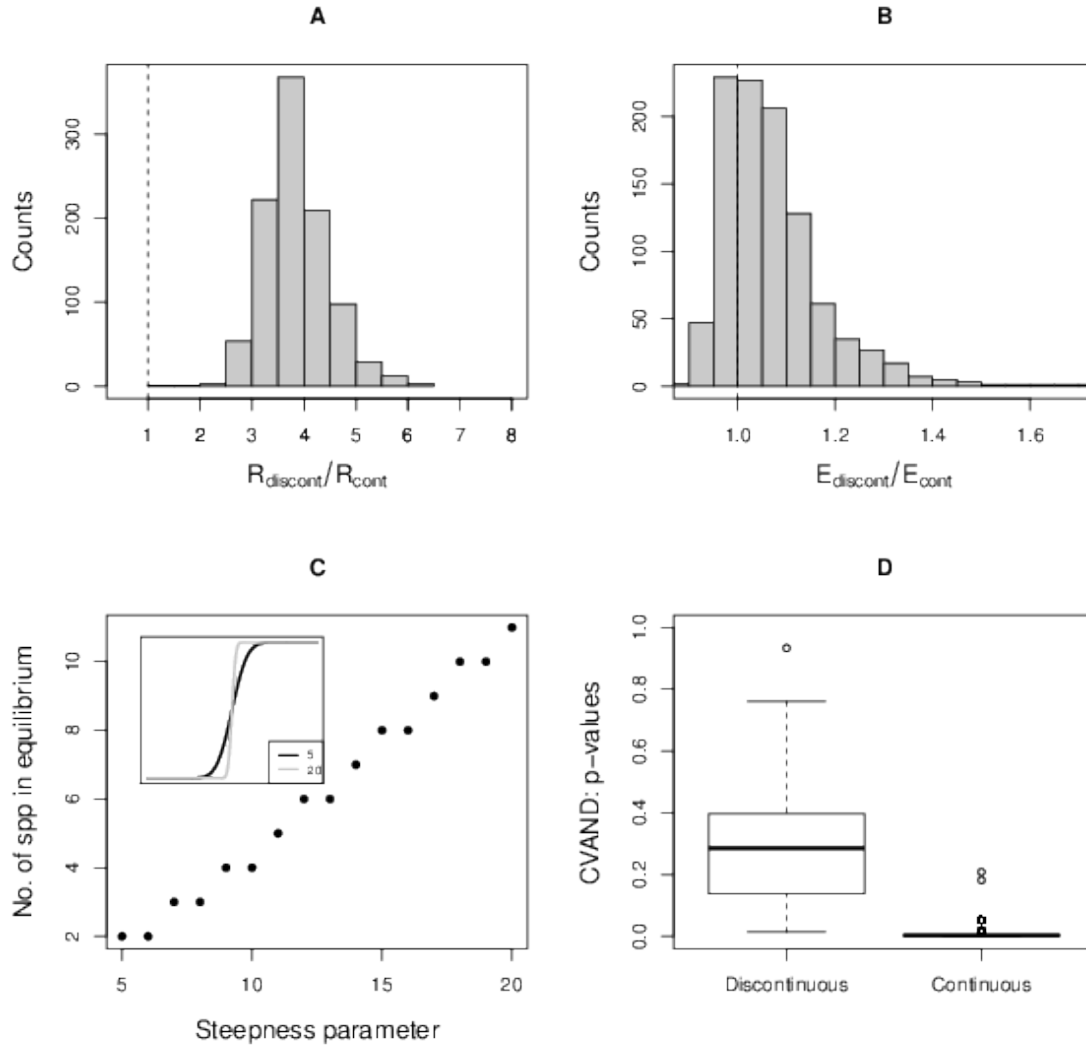
Chapter 5

Niche differentiation does not guarantee higher species diversity, longer lifetimes, or lower extinction rates

Abstract

Niche differentiation is normally regarded as a promoter of species coexistence in competitive systems, as it stabilizes species abundances. In a stochastic context, one might expect lower extinction rates and higher species richness and life expectancy in niche-differentiated communities than in neutral assemblages. Here we compare stochastic niche dynamics to neutrality in simulated assemblages subject to immigration from a regional pool, and find that the effect can be more complex. Trait variation that lessens competition between species will not necessarily give all immigrating species their own niche to occupy. Niche differentiation protects certain species from local extinction, but expedites exclusion of others. As a result, niches may lessen extinction rates and raise richness and mean local persistence times, or they may have the opposite effect, depending on the number of niches available to species. In particular, in a species assemblage similar to the tree community on Barro Colorado Island, extinction rates may actually be higher than neutral expectations if the number of niches is not sufficiently high. The impact of niche differentiation may also be similar to that caused by asymmetries in regional abundances and intrinsic growth rates, and is lessened by higher immigration rates.

5.1 Introduction

Niche differentiation is widely considered a prime force behind species coexistence, and thus instrumental in maintaining biodiversity. Niche differentiation stabilizes communities

by guaranteeing positive growth rates of rare species **??**. Without such stabilization, coexistence is only temporary, and biodiversity can only be maintained if gains from immigration or speciation compensate losses to competitive exclusion or drift as posited in neutral theory **?**.

Views on the implications of niche differentiation for biodiversity maintenance mostly come from deterministic community models with no immigration. However, stochastic formulations, which model the influence of variability whose deterministic details are unknown and unrelated to species differences, are more relevant for natural systems. In such settings, coexistence is defined in a probabilistic sense **?**. Furthermore, many if not most communities in nature are subject to propagule pressure from regional pools. In such open-community scenarios, even species that are not stabilized will colonize the community and may persist for a substantial period of time, and extinction is not an absorbing state, since re-colonization is possible in the future. One measure of diversity maintenance in this context is the mean species local persistence time – the average time between introduction through immigration and extirpation through drift or competitive exclusion.

Species local persistence times have been studied as indicators of ecological processes **??**. In particular, species local persistence times are well understood theoretically in the context of neutral dynamics, and have been used as tests of neutrality **????**. A study of a nearly neutral model showed that even slight fitness differences may greatly shorten local persistence times **?**. In contrast, niche model studies suggest that niche differentiation may prolong local persistence times considerably **???**. However, it can be argued that the niche scenarios proposed in these studies are too extreme to be applicable to most natural systems.

Pigolotti and Cencini (2013) **?** showed that niche stabilization increased species local persistence times compared with neutral assemblages. They used a simplified Lotka-Volterra competition model where species interactions are identical except that competitive suppression is stronger within than across species **?**. In the deterministic form of this model, all species stably coexist and in this sense occupy their own niche. Given this extreme degree of stabilization, the higher local persistence times are not surprising.

However, a one-to-one ratio between species and available niches is unlikely in highly diverse systems such as tropical forests, where it has been argued that the niches-to-species ratio can be low **??**. Even in lower diversity systems, there is no clear evidence as to whether all co-occurring species differ enough in traits so as to mutually invade one another. In

these situations is not clear that the average time between introduction and extinction of invaders in systems with niche redundancy is longer than in systems lacking niches altogether. Nor is it logically necessary that such a partially stabilized system will be more diverse than a similar system where richness is maintained only through mass effects.

The classical MacArthur-Levins model (**?**) allows for different niches-to-species ratios by positing that competition depends on phenotypic differences, a feature observed in nature **???**. For a deterministic, closed community, this model predicts limits to similarity between coexisting species, and exclusion of species that do not meet such limits. Simulations of this model have revealed transient states characterized by clumps of similar species separated by gaps **?**. This pattern suggests that some species are excluded faster than others. In particular, those with unfavorable characters might actually be excluded faster than in a game of drift, while those optimally adapted to a specific niche may persist indefinitely. This raises the possibility that partial stabilization, when the number of niches available is less than the number of species in the pool, may result in lower diversity and shorter local persistence times than neutral dynamics.

Here we perform a simulation-based study of the stochastic MacArthur-Levins model to compare species richness, the distribution of local persistence times, and extinction rates between assemblages differing by the presence or absence of stabilization. We consider how the niches-to-species ratio affects this comparison, as well as the effect of other factors governing species input and exclusion such as regional species abundances, immigration rate, and differences between species intrinsic growth rates. We structure our simulations to maximize relevance for the potential influence of niche differentiation on the dynamics of the Barro Colorado Island tree community, for which a recent paper found species input rates compatible with neutrality **?**.

5.2 Methods

Our study is essentially a comparison between species assemblages with and without niche differentiation. Our assemblages represent local communities of fixed size subject to a birth-death process, and surrounded by a large fixed pool of species, from which the local community receives immigrants. Communities start as a random draw from the pool, and in each subsequent step one individual is randomly chosen to die and is replaced by a new recruit. The probability that the death occurs on species i is proportional to $N_i \sum_j A_{ij} N_j$, where the competition coefficient A_{ij} quantifies the impact of species j on species i. The

proportionality factor is calculated by normalizing the sum of these probabilities across species to 1. The new recruit can either be an immigrant with probability m or a local birth event with probability $1 - m$. The probability that species i is chosen at this stage is proportional to $N_i(1 - m)r_i + mp_i$, where r_i is species i's intrinsic fitness, and p_i is its relative abundance in the regional pool. The deterministic immigration-free version of this model is the Lotka-Volterra model[1] $1/N_i\mathrm{d}N_i/\mathrm{d}t = r_i - \sum_j A_{ij}N_j$, which was used by MacArthur and Levins to describe competition along a trait axis **?**. Niche stabilization is only possible if intraspecific competition exceeds interspecific competition (**?**). In our niche model this is achieved by setting $A_{ij}(x_i, x_j) = \exp[-((x_i - x_j)/w)^\rho]$, where x_i and x_j are the trait values of species i and j, w sets the scale for the decline of competition between species with increasingly different trait values, and ρ sets the speed of exclusion dynamics. In contrast, neutral competition is achieved by setting $A_{ij} = 1$ between all species.

We simulate a local community with 21,000 individuals, approximately the number of trees larger than 10 cm diameter at breast height in the 50-hectare plot on Barro Colorado Island **?**. This local community is surrounded by a fixed species pool **?**, which we populate with 400 species with trait value x_i drawn uniformly from 0 to 1. We start with a random draw of individuals from the regional pool, and then simulate dynamics for 5.0×10^7 steps to allow for a stationary state to be reached (i.e. to allow for species abundance distributions averaged across a set of runs to settle to relatively stationary values). We then run for another 5.0×10^7 steps while keeping track of all introduction and extirpation events, for a total of 1.00×10^8 steps, corresponding to roughly 5,000 community turnovers. At the end of a run, each species average local persistence time is computed by calculating the average number of steps that species persisted between being introduced through immigration and being extirpated due to competitive exclusion or drift.

By analogy with common statistical terminology for null and alternative hypotheses, we will refer to communities lacking niche stabilization as H_0 assemblages, and those with niche stabilization as H_1 assemblages. Aiming to isolate the effect of niche stabilization from other drivers of dynamics, in our "baseline" H_0 and H_1 assemblages all species have identical regional abundances and intrinsic growth rates, and we use a low immigration rate ($m = 0.01$, about an order of magnitude lower than has been estimated for the 50ha Barro Colorado Island plot **?**). In our baseline H_1 case we consider a scenario with substantially

---

[1]The correspondence with our model is fairly intuitive, but for a more detailed understanding the reader can consider the analogy with stochastic models of chemical reactions and the corresponding differential equation models **?**.

fewer niches than species. Specifically, we set w so that about 13 species would stably coexist in the absence of stochasticity and immigration within this trait range (i.e. 13 niches. Specifically we use $w = 0.063$). We set $\rho = 4$ because it gives relatively fast niche sorting. The classical model studied by MacArthur and Levins (1967) with an exponent 2 leads to much slower dynamics, which are not ideally suited for contrasting with neutrality (see Fig. J.3). The actual rate of exclusion dynamics in nature is unknown.

Other factors causing differences between species might mimic the effects of partial niche differentiation. Regional abundances affect the frequency of immigration by different species and hence are likely to modulate local persistence times. If regional abundances vary across species, some will be more prone to extinction than others due to different probabilities of being rescued from low abundance by immigrants. Differences in intrinsic growth rates can have similar effects, as some species would recover from low abundance faster than others. Finally, higher immigration increases local persistence times, magnifies the effects of uneven regional abundances, and dampens the effects of internal dynamics thus bringing H_0 and H_1 closer.

To examine the influence of uneven regional abundances we populate the regional pool according to a neutral metacommunity species abundance distribution using R package untb **?**. We used a regional pool size of 150,000 individuals, and $\theta = 50$, the approximate value for which the neutral model provides the best fit to the BCI species abundance distribution. Regional richness is approximately 400 species as in our base case, but here abundances are approximately log-series distributed. We consider the impact of raising the immigration rate tenfold to $m = 0.1$, maintaining the uneven abundance distribution in the pool. To test for the effect of differences in intrinsic growth rates, we consider two alternatives to the baseline case where r_i differs among species. In the first, representing a scenario where metabolic costs or other physiological limitations cause lower fitness of species with extreme trait values, we take $r_i = x_i(1 - x_i)$. In the second, we consider influences on the growth rates unrelated to the trait at hand, and hence draw the r_i at random from a uniform distribution, $r_i \sim \text{runif}(0, 1)$.

Finally, given that the true number of niches in real communities like BCI is largely unknown, and in order to more fully explore the influence of niche number on dynamics, we consider a set of scenarios with niche number varying from 4 to 380, leading to niches-to-species ratios ranging between 0.01 and 0.95. For these scenarios we used the neutral metacommunity abundance distribution as described above, an immigration rate equal to that estimated for BCI ($m = 0.08$ **?**), and randomly drawn r_i across species.

5.3 Results

Local persistence times in the H_0 assemblage show no relation to species trait values (Fig. 5.1A) and do not correlate with species abundances observed at the end of the simulation (correlation index $\rho_{\log(N),\log(T)} = 0.008$), while in the H_1 both species abundances and local persistence times form a distinctive pattern along the trait axis (Fig. 5.1B), and correlate strongly with each other ($\rho_{\log(N),\log(T)} = 0.86$, see also Fig. J.1). The average local persistence time across all species in the H_1 assemblage, $\bar{T}_1 = 85.4$ community turnovers, is higher than in the H_0 assemblage, $\bar{T}_0 = 22.1$ community turnovers. On the other hand, it also has a much higher coefficient of variation, $CV[T_1] = 4.9$ against $CV[T_0] = 0.4$, indicating much higher local persistence time variation in the partially stabilized assemblage.

This is confirmed by the histograms in Fig. 5.1C. The partially stabilized community is characterized by a strong positive skew and the presence of 8 resident species, defined as species that were present throughout the tracking period. The local persistence times of the resident species are many orders of magnitude higher than other species, and therefore heavily drive the higher mean local persistence times in the H_1: if residents are discounted, the mean local persistence time in the H_1 drops by an order of magnitude to 8.3 turnovers, lower than $\bar{T}_0$. Compared with neutral dynamics, partial stabilization leads to higher overall mean local persistence times but lower local persistence time of transient species, as well as lower species richness (Fig. 5.1D). The estimated extinction rate, calculated as the number of extinction events divided by the time interval spanned by the tracking period, is actually higher than neutral (Fig. 5.1D).

It should be noted that the niche axis in our model is finite, and hence has edge effects. Our simulations showed that these drastically influence the local persistence time distribution compared with a circular axis (see Fig. J.2). This is because the positions of the niches on a finite axis are set, whereas on a circular axis they constantly shift through time, thus mitigating the effects of stabilization. We used a finite axis as it is more likely to represent real systems. Also, for the reader interested in the classical MacArthur-Levins model, in Fig. J.3 we show that the local persistence time distribution in that model is much closer to neutrality, which is not surprising given the slower (weaker) exclusion dynamics under $\rho = 2$.

When abundances in the pool are log-series distributed, the local persistence times distribution shifts towards the extremes (Fig. 5.2A, 5.2E): compared with a pool with identical abundances, species now generally have much shorter local persistence times,

although now in the H_0 a number of them persist throughout the entire tracking period, which did not occur with the homogeneous pool, and in the H_1 that number is higher than before. These new and additional residents arise because of their comparatively high regional pool abundance (see Fig. J.4). As a result, H_0 and H_1 now have more similar local persistence time statistics and extinction rate than in the baseline case (cf. Fig. 5.2I, 5.1D).

When immigration is increased in this uneven abundance scenario from $m = 0.01$ to $m = 0.1$, most local persistence times increase, and many species achieve residence status in both H_0 and H_1 assemblages (Fig. 5.2B, 5.2F, cf. Fig. 5.2A, 5.2E. See also Fig. J.5). A few species now have actually lower local persistence times, presumably due to regional rarity. Overall, regional abundances become stronger determinants of local persistence times compared with the lower immigration scenario. As a result, the ratio of richness, community average local persistence time, and extinction rates are all closer to 1, compared with the lower immigration case (cf. Fig. 5.2J, 5.2I). It should be noted that although high immigration may make those statistics rather indistinguishable, niche differentiation still influences which species are residents (see Fig. J.5).

The introduction of a parabolic shape to species intrinsic growth rates r_i has the anticipated effect on the H_0 assemblage. It dramatically shifts most of the H_0 local persistence time distribution to the left, and promotes a very small number of species (those with the highest r_i) to resident status (Fig. 5.2C). The effects on the H_1 assemblage are more complicated, with local persistence times actually increasing for several species (Fig. 5.2G), and correlating negatively with r_i ($\rho_{r,\log(T)} = -0.2$, see also Fig. J.6). This is presumably due to lower competitive suppression between species with low intrinsic growth rate near the edges of the trait axis. However, the qualitative shape of the H_1 local persistence time distribution is unchanged, indicating that the r_i have a small effect on H_1 exclusion dynamics relative to competition. The H_0 and H_1 local persistence time distributions become less distinguishable under parabolic r_i compared with identical r_i across species, and the extinction rate of H_1 is now lower than H_0 (cf. Figs. 5.2K and 5.1D). Random variation in r_i lowers local persistence times for almost all species in both the H_0 and H_1 assemblages, but again the effect is much stronger on the assemblage lacking stabilization (Fig. 5.2D, 5.2H). Now H_1 has an even greater lead in local persistence times over H_0 than in the baseline, and the extinction rate in H_0 increases so much that it is now comparable to H_1 (Fig. 5.2L).

Fig. 5.3 shows the community average local persistence time and extinction rate for different niche numbers in a scenario geared to include the complex elements of a real

species assemblage like the BCI tree community. As the number of niches increases, the community average local persistence time increases, while average extinction rate decreases. Critically, both straddle neutral levels depending on the niche-to-species ratio. Further examination of the local persistence time distributions (see Fig. J.7) reveals that this is because the presence of few niches shortens the local persistence times of many species, but as more niches are added more species are promoted to resident status.

5.4 Discussion

Our main results can be understood as follows. Niche differentiation causes some species to persist all but indefinitely, thus raising community average local persistence times; however, when there are fewer niches than species, many species fall in exclusion zones and are extirpated more quickly than expected from drift. Among those species, local persistence times can be much lower than in a neutral community. As a result, average extinction rates in partially stabilized communities may surpass the neutral assemblage. And since the number of niches is limited and many species are quickly excluded, richness can also be lower than in a neutral community. As stabilization increases towards the limit of one niche per species, mean local persistence times increase together with the proportion of residents, and extinction rates decrease. Niche stabilization may therefore result in a net increase in local persistence times compared to neutrality if there are enough niches relative to the number of species in the pool, but a net decrease if that ratio is too low.

Asymmetries in regional abundances affect local persistence times because they cause asymmetries in the rescue effect. While local species that are common in the pool are less prone to extinction due to an enhanced rescue effect and therefore may last longer, the opposite holds for the large majority of species that are regionally rare. Similarly, differences between species intrinsic growth rates accelerate competitive exclusion, and increase local persistence times only for the few species with the highest growth rates.

Overall, asymmetries in regional abundances or intrinsic growth rates had a bigger impact in the assemblages lacking stabilization entirely. This showcases the robustness of stabilized assemblages against other dynamical forces, in contrast with the susceptibility of neutral dynamics. However, it should be noted that species differences, whether related to partial niche stabilization or to pool abundances and intrinsic growth rates, have qualitatively similar consequences: some species die fast while others may last much longer. Furthermore, immigration tends to blur the dynamical differences between stabilized and

non-stabilized communities. Thus it may not be trivial to distinguish between assemblages with and without niche stabilization using local persistence times and diversity alone.

Our findings are compatible with real-life observations of "core" species which persist for long times and "occasional" species whose local populations may regularly undergo local extinction and recolonization events **?**. Our results suggest that core and occasional species could fall into these categories on the basis of their stance in competitive interactions, or else because of their regional relative abundance or intrinsic growth rate.

Our results also bear on the recent finding that observed species input on BCI is compatible with neutrality **?**. We find extinction rates are often higher under niche differentiation than neutrality. When species richness is at equilibrium, species input rates should match extinction rates. Our results then suggest that the observed species input could also be compatible with niche differentiation, as the link between species input and diversity, being mediated by species local persistence times, is not unique to a specific assembly process. However, if species local persistence times can be measured and related to their traits, abundances, and regional prevalence, it may help reveal whether niche differentiation is playing a role in coexistence.

Our study provides a refined understanding of the sense in which niche differentiation promotes coexistence, highlighting the new insights into biodiversity maintenance that can be gained by building on the framework of neutral models **?**. Compared to neutral dynamics, niche differentiation stabilizes species abundances and may indefinitely prolong local persistence times of certain species, but it does not guarantee higher species diversity, longer local persistence times, or lower extinction rates.

5.5 Figures

Figure 5.1: **A**, **B**: Species abundances and local persistence times plotted against relative trait values in baseline H_0 assemblage (**A**) and H_1 assemblage (**B**). Stems show $\log N / \log N_{max}$, and red lines show $\log T / \log T_{max}$, where N are species abundances and T are species local persistence times. **C**: Number of species by mean local persistence time in the baseline assemblages with no stabilization (H_0, blue bars) and partial stabilization (H_1, red bars). **D**: Comparison of species richness, average local persistence time, and extinction rates between baseline H_1 and H_0 assemblages. Bars show ratio H_1/H_0.
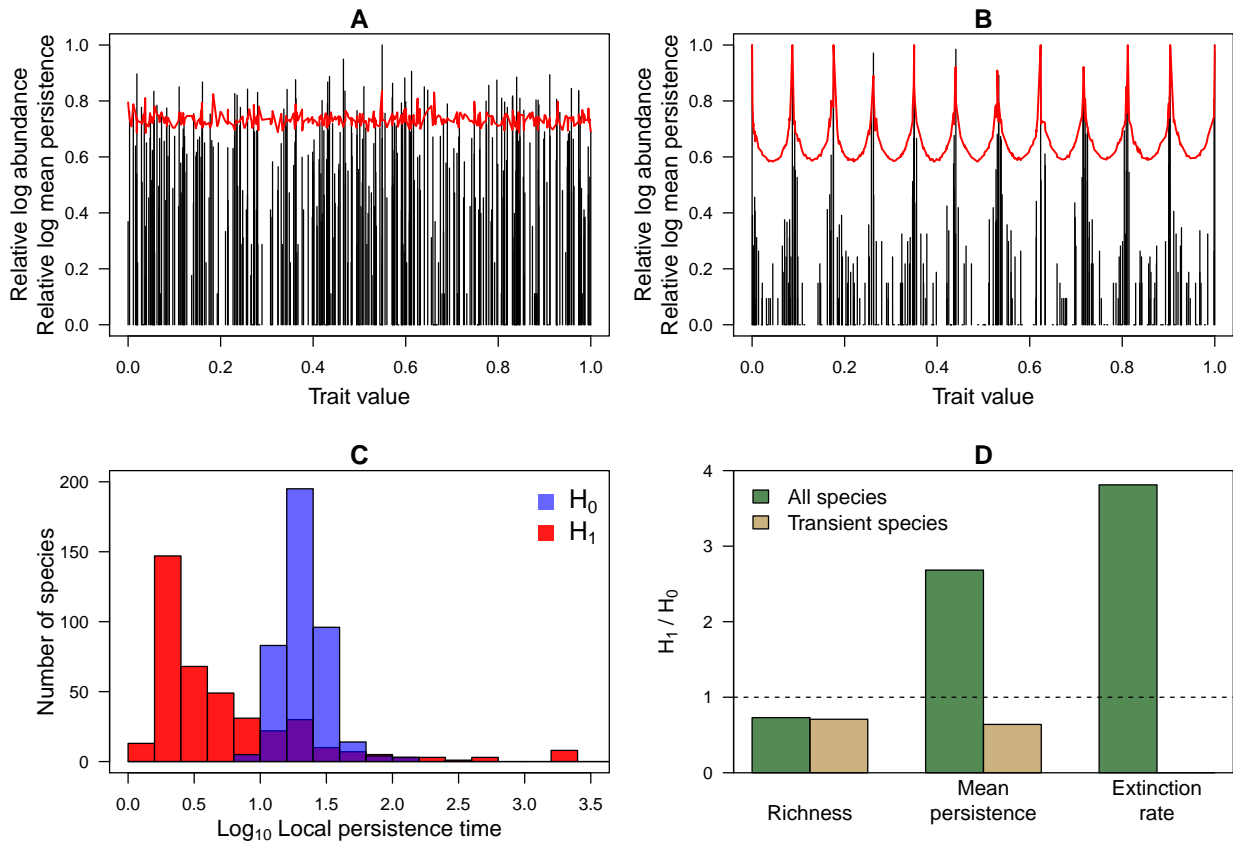
Figure 5.2: Comparison of distribution of local persistence times between baseline and variant treatments for H_0 assemblages (top row) and H_1 assemblages (middle row). Baseline (light colored bars): identical pool abundances, identical r_i, immigration rate $m = 0.01$. Variants (dark colored bars): **A**, **E**: log-series distributed pool abundances; **B**, **F**: log-series distributed pool abundances and higher immigration rate ($m = 0.1$); **C**, **G**: parabolic distribution of intrinsic growth rates r_i; **D**, **H**: random distribution of r_i. Bottom row, **I-L**: ratio between richness, mean local persistence time, and extinction rate of H_1 and H_0 assemblage in each of the variants.
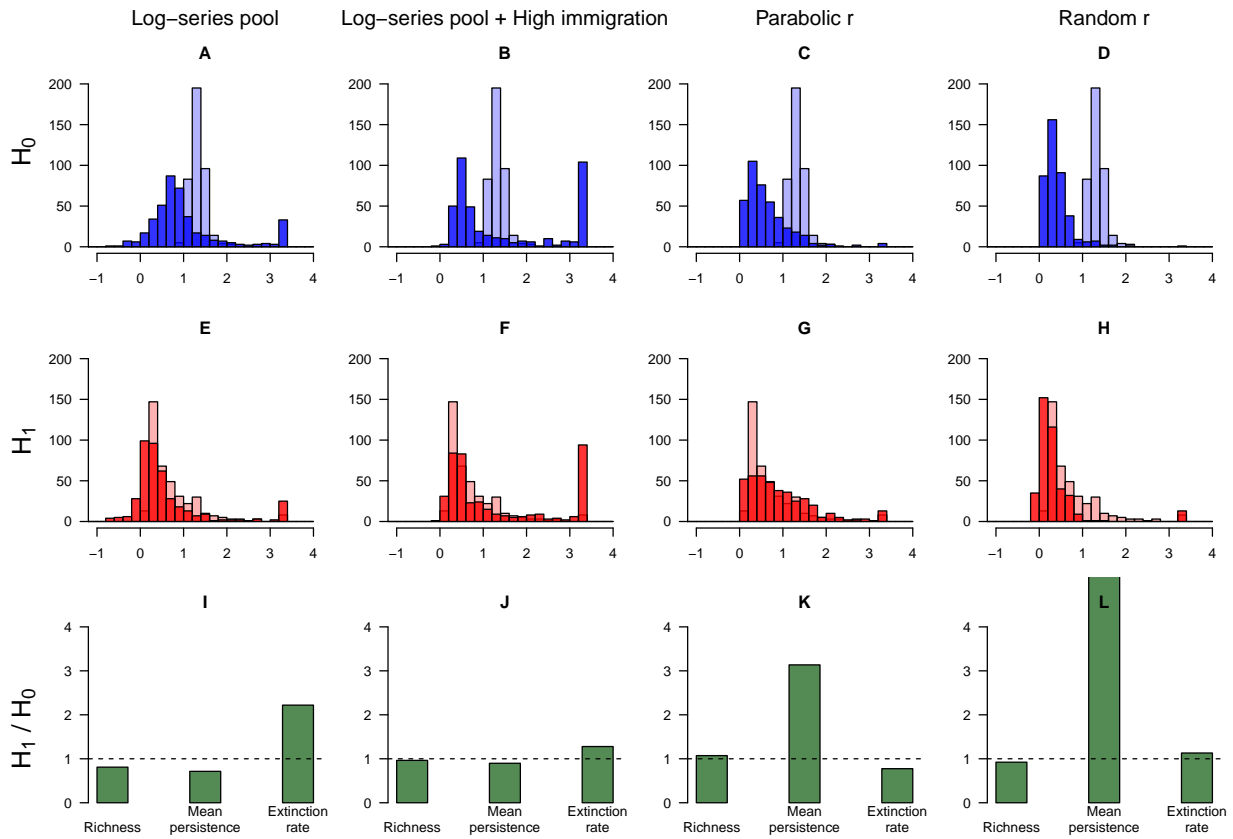
Figure 5.3: Comparison of community-wide mean local persistence times (in units of community turnovers) and extinction rates (in units of extinction events per community turnover) between neutral assemblages and partially differentiated assemblages with increasing number of niches and fixed regional diversity. Community size 21,000 individuals, fixed pool with 150,000 individuals and circa 400 species, and immigration rate $m = 0.08$ were chosen to approximate parameters observed or estimated for the BCI tree community. Regional abundances follow neutral (approx. log-series) distribution. Intrinsic growth rates drawn at random from a uniform distribution.
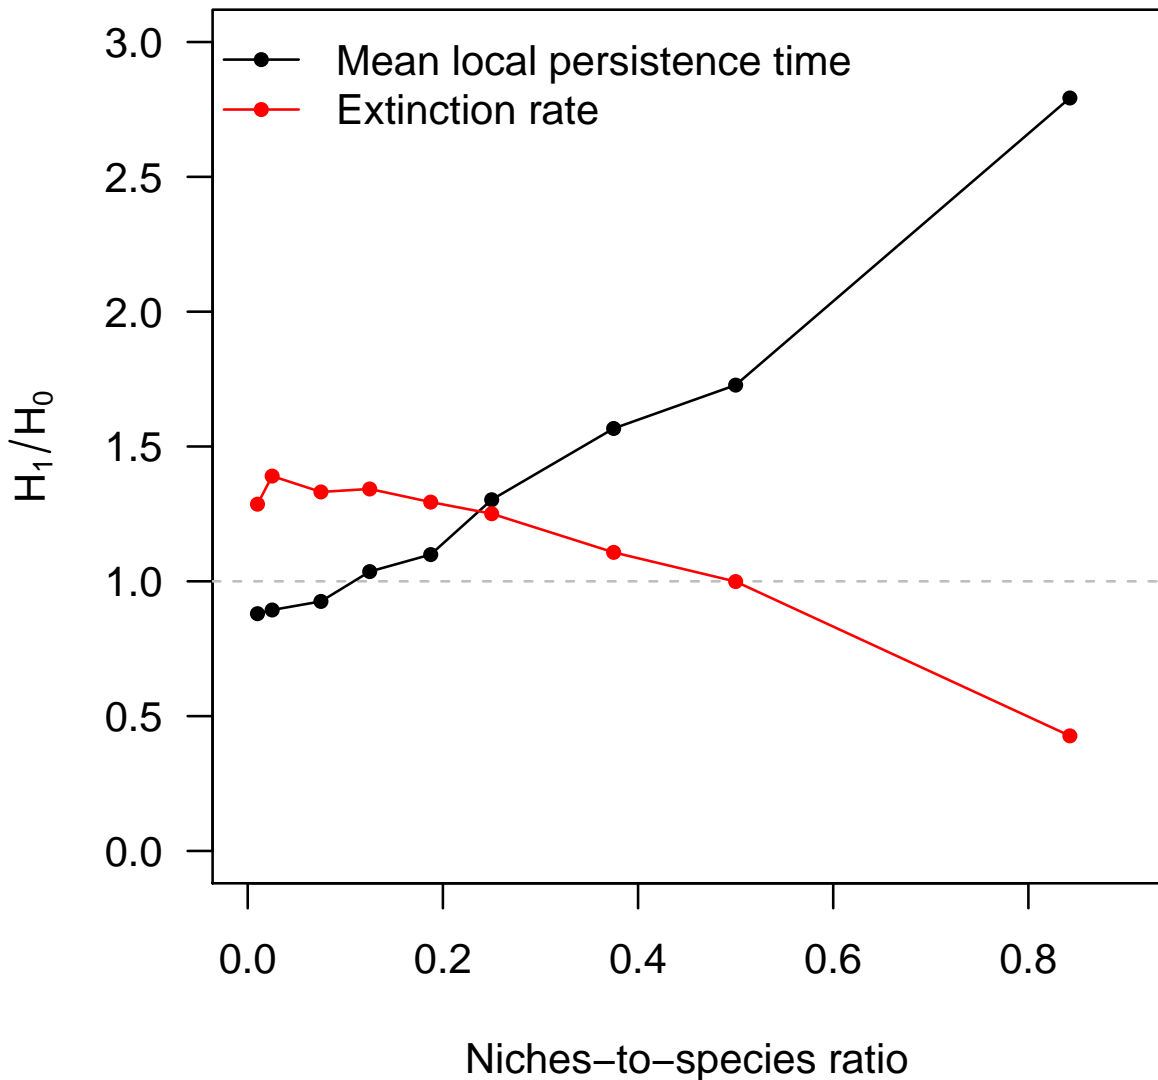
Chapter 6

Trait pattern under stochastic niche assembly: are species clusters a general phenomenon?

Abstract

Niche differentiation explains the maintenance of biodiversity, but remains difficult to demonstrate. Inference based on trait patterns is commonly used, as traits are easy to measure and directly tied to ecological strategy. However, classical trait pattern theory ignores stochasticity and immigration, and assumes competitive interactions are determined strictly by trait similarity, which makes it difficult to extrapolate predictions to natural systems. Recent literature suggests that niche differentiation may drive the spontaneous formation of species clusters, but it is not known whether the phenomenon is general and under what conditions it is expected to appear. We propose a metric to quantify clustering, and show that clustering occurs across different niche mechanisms, and is more likely to appear when competition is primarily determined by trait similarity, niche sorting is fast, regional diversity is high, and immigration is low. Classical predictions of functional overdispersion and even spacing were not supported in our simulations.

6.1 Introduction

It is widely held that nature's prodigious levels of biodiversity would not be possible without niche differentiation. Although speciation provides new species, those can be lost to extinction and ecological drift without stabilizing forces. As an important provider of such forces, niche differentiation is thought to be a major driver of community structure. Yet, this can be difficult to verify empirically. Experimental approaches are the golden standard

since they can isolate niche differentiation from other community assembly processes (**?**) and directly test competition between pairs of species (**?**). However, experiments can be logistically difficult for large or long-lived organisms. Inference-based methods that suggest process based on observed patterns can be a valuable aid, and several have been used (**?**). Notable among those are trait-based approaches, e.g. (**??**). Traits have been extensively used to infer niche differentiation and other community assembly processes (**??**). Trait-based approaches have several advantages over strictly taxonomic approaches because they are easily generalizable and have explicit ties to ecological strategy and performance (**??**). Furthermore, inference based on trait patterns does not require censuses spanning extended periods of time, which is particularly handy in the study of long-lived organisms such as trees.

However, we have previously argued that the theory underlying many trait-based inference approaches needs updating (**?**). Studies of pattern on niche axes are based primarily on intuitive hypotheses linking niche differentiation and functional overdispersion (**??**), or simplified models of competition that do not capture population dynamics (e.g. **?**), or where competition is determined purely by trait similarity (**????**), which has been criticized (**??**). Classical expectations of limiting similarity have been questioned in the past (**?**), and lag behind recent findings from modeling studies (**?**). In addition, they also typically ignore the influence of stochasticity and immigration (with some notable exceptions, see **???**); it is not clear that predictions based on deterministic immigration-free models bear out in stochastic open communities, which hampers extrapolation to many natural systems.

Recently, ecologists have considered the possibility that community assembly under niche differentiation can produce clusters of similar species (**??**), a phenomenon that has been hailed as potentially bridging neutral theory with classical niche theory (**?**). Clusters arise in classical deterministic niche models as long-lived transients (**?**), and can persist indefinitely if species similar in one trait can be stably maintained by differences on other niche axes (**??**), environmental fluctuations (**?**), and evolution (**?**). We suggested in a recent paper that clusters also appear under stochastic niche assembly with immigration (**?**), but clusters have been strikingly absent from stochastic niche simulations by other authors (**???**), and have not been tested across models with distinct niche mechanisms. It remains unclear whether clustering is a general phenomenon, and therefore whether it can be used to infer niche differentiation (**?**).

Furthermore, as clusters have only recently been associated with niche differentiation, the literature is sparse on tools for detecting them in ecological data. Empirical assays have

so far been either based on subjective examination of data (**?**) or relied on metrics that depend on parameter choices such as binning the trait axis (**?**), or presence-absence metrics that require time series data (**?**), which is often unavailable for long-lived organisms.

Here we first propose a parameter-free abundance-weighted metric to identify species clusters, and then ask whether clustering is a general feature of stochastic niche dynamics under immigration. We validate our metric by testing it on a model where clusters are known to arise, the MacArthur-Levins model (**????**). In this model competitive interactions are set directly by trait similarity, thus setting the stage for niche differentiation. We then test how the clustering pattern is affected by immigration rate, regional diversity, abiotic filtering, and departures from the exclusive dependence on trait similarity. Finally, we assess whether clustering appears across a suite of models comprising different niche differentiation mechanisms: a) differences in preferred resource, b) differences in preferred abiotic conditions, and c) differences in life history strategies. We also use our niche scenarios to test two common predictions regarding niche differentiation: that species under niche assembly will be more different than expected by chance (functional overdispersion, **?**), and that species will be evenly dispersed along the trait axis (even spacing, **?**).

6.2 Methods

Metrics

Clustering metric: the gap statistic

We propose an abundance-weighted metric that identifies the number of clusters in the data and quantifies the degree to which the data is clustered. Our metric is based on the gap statistic method (**?**), coupled with the k-means cluster-finding algorithm (**?**). For a given community to be tested, the metric produces an index quantifying the degree of clustering in the data, compares the index against a set of neutral communities used as the null hypothesis, and computes the standard score. For details about the metric, see Appendix K.

In our simulations we also measure the average time interval between introduction and extirpation of each species (i.e. lifetimes) to look at whether and how the clustering pattern relates to impacts of niche differentiation on species dynamics.

Other metrics

In addition to our new metric, we also test our niche scenarios for pattern with three other metrics that have been used in the literature: Rao's quadratic index (**?**), functional dispersion (**?**), and the coefficient of variation in trait differences between adjacent species (**?**). See Appendix K for a detailed description.

Models and simulation design

In Appendix L we provide a detailed description of our simulation design and niche models, which we now summarize.

We simulate local communities of fixed size undergoing competitive dynamics subject to a stochastic birth-death process and surrounded by a large fixed pool of species, from which the local community receives immigrants. We set community size and immigration rate to match observations and estimates for the tropical tree community on Barro Colorado Island, in Panama. Communities start as a random draw from the pool, and in each subsequent step one individual is randomly chosen to die and is replaced by a new recruit. A proportion of recruits fulfilling vacancies left by deaths are immigrants from the pool, and the remainder are local birth events. The particular model used dictates the probabilities that each species is selected for each death and each recruitment event.

MacArthur-Levins scenarios

We consider a variety of scenarios spanning several different niche mechanisms. We start with variations of the Lotka-Volterra competition model studied by MacArthur and Levins (**?**), a very influential model where clusters have been shown to occur (**?**). This model is conceptually important because the strength of competition is determined entirely by species similarity—see Discussion section below. Our different MacArthur-Levins scenarios are obtained by setting different values of the parameters regulating immigration, regional diversity, environmental filtering, speed of competitive sorting, and shape of the competition matrix.

Niche differentiation by preferred resource

The second category of models we test is the Rosenzweig-MacArthur model of resource-consumer interactions (**?**). The model assumes a linear array of substitutable resources

depended upon by a set of consumers. Resources follow the same stochastic birth-death process as consumers, except there is no resource immigration, so extirpated resource populations cannot be replenished. A consumer's trait value defines its preferred resource, and its consumption of all other resources declines as a function of their difference to that consumer's preferred resource. Abrams showed that resource depletion or exclusion affects competition-similarity relationships (**?**), as well as on coexistence outcomes (**?**). We therefore consider two scenarios of this model. In the first, most resources are never severely depleted by consumption, whereas in the second resource depletion is severe and many resources are extirpated.

Niche differentiation by preferred abiotic environment

The third type of niche model we examine was introduced in **?** based on Tilman's stochastic niche model (**?**), and also studied by **?**. It assumes a linear array of sites forming a gradient of fixed abiotic conditions. A species' trait determines its optimal environment, and its ability to thrive in other local environments declines with the difference between those environments and its optimum. We explore two dispersal scenarios in this model: local dispersal, where the probability of arrival of propagules to a site quickly drops with the distance to the parent, and global dispersal, where the probability of arrival is independent of distance to the parent. The latter corresponds to a case where the physical distances are small relative to the dispersal ability of the individuals.

Niche differentiation by life history strategy

We examine two niche models where species differ by life history strategy. Unlike the previous mechanisms, here there is a competitive hierarchy in the sense that some species are better competitors than others. In both models, the species trait is fecundity, or number of propagules produced by an individual. In the competition-colonization tradeoff model (**??**), species trade off fecundity with competitive ability, defined as the ability of propagules to displace individuals of other species. In Muller-Landau's tolerance-fecundity tradeoff model (**?**), species trade off fecundity with propagule stress tolerance, and the environment varies in stress levels.

Null model: neutral dynamics

We quantify clustering pattern in our niche models by comparing them to an appropriate set of neutral communities, where complete equivalence is assumed for all species. These neutral communities undergo the same zero-sum stochastic birth-death-immigration process as the niche models, and birth and death probabilities are simply proportional to species abundances. Each niche community is compared against a corresponding set of 1,000 neutral communities with identical regional diversity, local community size, and immigration rate. All simulations and statistical manipulations are performed using the R language (**?**).

6.3   Results

The stem plot in Fig. 6.1A represents the outcome of a simulation of our baseline scenario of the MacArthur-Levins niche model. Species are arranged by trait value on the x-axis, and their abundances are indicated on the y-axis. The plot shows unmistakable structure, with the apparent formation of about 13 clusters around dominant (i.e. high-abundance) species, which in turn are evenly dispersed along the trait axis. The gap statistic has a clear spike at $k = 13$ (Fig. 6.1D), supporting the visual identification of 13 clusters. In the deterministic immigration-free formulation of this model, there are 13 evenly spaced stably coexisting species (Fig. 6.1G). The correspondence with the deterministic case suggests interpreting each cluster as a group of species vying for the same niche.

The average standard score of the gap statistic across ten runs of our baseline MacArthur-Levins scenario was > 5, indicating unambiguous clustering (Fig. 6.2A). The standard score of the gap statistic clearly decreases with immigration (Fig. 6.2B). Increasing regional diversity while maintaining the number of niches available also has a clear effect: the average standard score is higher for scenarios with more diverse pools (Fig. 6.2C).

In addition to the baseline, clustering was also clearly evident regardless of the presence or absence of edge effects and abiotic filtering (Fig. 6.2A). However, this was not the case of scenarios where competitive niche sorting was slow and scenarios where competition coefficients deviated from exclusive dependence on species trait similarity. The clustering metric failed to distinguish those cases from the neutral nulls.

The center and right columns in Fig. 6.1 show two of those scenarios in more detail (see Appendix L for a detailed description of the scenarios). Compared with the baseline case, the stem plot of the scenario with Gaussian competition coefficients (Fig. 6.1B) shows less

clear-cut distinction between species abundances by trait, while the scenario with random w (Fig. 6.1C) shows strong abundance structure, as in Fig. 6.1A, but less regularity on the trait axis. Our deterministic immigration-free results show that stably coexisting species in the scenario with random w are irregularly spaced (Fig. 6.1I), whereas in the scenario with Gaussian competition species are as evenly spaced as in the baseline scenario (Fig. 6.1H). These results suggest that the scenario with Gaussian competition differs from the first in the speed of competitive sorting, whereas the scenario with random w differs by asymmetries in species interactions. The gap statistic was generally much lower than in the baseline scenario. Indeed, whereas the baseline case shows a distinct peak soaring above the neutral expectations, both the slow-sorting case and the case with randomness in the competition coefficients failed to express strong evidence of clustering (compare Figs. 6.1E and 6.1F to 6.1D). For stem plots of deterministic and stochastic runs of all other scenarios, as well as their respective gap statistic curves, see Appendix M.

Simulation outcomes of the other niche models are shown in the stem plots in Fig. 6.3, and the gap statistic results for these models are shown in Fig. 6.4. We see considerable variation in the standard score across models with different niche mechanisms, as well as across different scenarios within the same niche mechanism. Significant clustering was sometimes observed in the Rosenzweig-MacArthur model, but only in the scenario with high resource depletion. In the Schwilk-Ackerly model of competition for local environmental conditions, clustering was strong in the scenario with local dispersal but did not occur in the scenario with global dispersal. The competition-colonization communities are strongly clustered, but the abundance structure in Fig. 6.3 suggests that the dominant cluster may be driving the result. Indeed, all 10 runs were estimated to have only either 2 or 3 clusters. This is in contrast with the deterministic results where a higher number of species stably coexist (see Appendix M). The tolerance-fecundity communities consistently fit a single cluster (see caption of Fig. 6.4). The abundance structure (Fig. 6.3) suggests that this is because our metric "sees" one big cluster on the left side of the axis. Indeed, if the test is restricted to the left side, more clusters appear and clustering is much enhanced (mean standard score 1.86 across 10 runs, standard deviation 0.7).

Immigration links species lifetimes to regional abundances. The relationship is very strong in the neutral scenario, where lifetimes are almost completely predicted by regional abundances (Fig. 6.5A). This is less so in the niche scenarios, particularly when niche sorting is fast (Figs. 6.5B and 6.5C). Specifically, in the niche cases some species had higher lifetimes than could be expected based on their regional abundances. Both in

neutral and niche cases, species that persisted throughout the entire tracking period in the simulation ("residents") tended to have high regional abundance, revealing an obvious influence from immigration; however, in the fast niche sorting cases some residents did not have particularly high regional abundances (Fig. 6.5C. See also Appendix M for all other scenarios). Species identified as centers of their clusters (medoids) tend to be residents, and in the fast sorting cases their lifetimes may exceed expectations based on regional abundance. Finally, note that there is a tendency for cases identified by the gap statistic as strongly clustered to also be cases with a looser relationship between regional abundances and lifetimes and vice versa, indicating strong impact of competition on lifetimes (see Appendix M).

It should be noted that these niche scenarios do not express the type of trait structure that motivates commonly used metrics in the empirical literature. Species in these scenarios do not show a tendency to overdisperse in functional space in the sense of Rao's quadratic index (**?**) or the related functional dispersion index introduced in (**?**). In fact, not only did our scenarios largely fail to express significantly high Rao and functional dispersal indices, some actually showed significantly low values compared with the neutral scenarios (Fig. 6.6A). This occurred with our tradeoff scenarios, and the MacArthur-Levins scenario where environmental filtering towards an optimal trait value acts on top of niche sorting. While the latter result is perhaps not surprising, the outcome of the hierarchy models indicates that species mostly cluster around a single dominant niche. In that sense, communities formed under these hierarchy models are not entirely distinguishable from a case of pure filtering from the point of view of these metrics.

Species are also not evenly dispersed along the trait axis, contrary to the common expectation that has motivated metrics of even spacing such as the variance in distances between nearest neighbors or adjacent neighbors on the trait axis (**??**). We tested the coefficient of variation (CV) in distances between adjacent neighbors on the trait axis, and found that it is indistinguishable from that of the neutral communities (Fig. 6.6A). Metrics of even spacing that do not account for abundances cannot distinguish between dominant species and rare species that may be present due to a recent introduction from the regional pool. As a result, immigration severely limits the ability of this metric to detect a pattern. One possibility for bringing abundance into account is to subset only the most abundant species. In Fig. 6.6B we illustrate the idea using an example run of the baseline MacArthur-Levins scenario. The CV is calculated in turns across progressively larger subsets, starting with only the community's 5 most abundant species and then adding the

other species in order of abundance until the entire community is sampled. Notice how the CV ranges from lower to higher than null (neutral) expectations as more species are sampled, but falls back to null expectations when the entire community is taken.

6.4   Discussion

Our metric successfully identified clusters in the MacArthur-Levins model. While the fact that species identified as medoids tended to have high regional abundance signals the influence of immigration, the fact that medoids tended to be residents signals both their dominant status in the competitive dynamics and the success of the clustering metric in finding meaningful patterns in the community, as it indicates that species around which others are clustering are those whose lifetimes are most enhanced due to niche differentiation.

Pattern was stronger at low immigration, which is intuitive: because the arrival of immigrants to the local community is independent of competitive interactions and any other internal processes, higher propagule pressure is expected to lead to noisier outcomes and therefore weaker pattern. The positive effect of regional diversity is also easy to understand: other things being equal, clusters with more species are more distinctive. Note that our case with the lowest regional diversity, which had a species-to-niches ratio of only about three species per niche, was not identified as distinctive from its neutral counterpart, even though the competitive interactions driving this community were of the exact same nature as those in the cases that were distinctively clustered. This suggests that clusters in nature could be more easily found in extremely diverse systems such as tropical forests and coral reefs than in relatively species-poor systems such as temperate forests and grasslands.

? argued based on a proof from **?** that a community supersaturated with species in an unstable equilibrium will always undergo transient clustering if the matrix formed by the competition coefficients is cyclic, i.e. its rows are rotations of each other. Specifically, they show that a long time after the disturbance species abundances will converge to the dominant eigenvector of the community matrix, which for any cyclic matrix is periodic. Our results across variants of the MacArthur-Levins model can be understood under the light of **?**'s (**?**) argument. Similarity-based competition on a circular axis meets the cyclic prescription exactly, whereas similarity-based competition on a finite axis meets it approximately if the axis is sufficiently long (i.e. if interactions are not dominated by edge

effects[1]). In contrast, the cyclicity is more severely broken, and with it the clustering pattern, when the parameter w is no longer held constant across all competitive interactions, as was the case in two of our MacArthur-Levins scenarios. Environmental filtering that acts only on species intrinsic fitness does not affect the matrix, and accordingly had no deleterious effect on clustering.

The agreement between our results and predictions based on a deterministic analysis indicates that stochastic assembly with immigration does not affect the nature of the transient. Rather, it appears that stochastic immigration maintains the system in a permanent transient state of the deterministic formulation. On the other hand, clusters were not seen in our scenarios with slow niche sorting. This suggests that within stochastic assembly the relative speed of sorting is as important as the architecture of the competition matrix. Clusters were absent not because the matrix is not cyclic but because the pace at which they form is so slow that stochastic immigration can easily overwhelm it. Indeed, clusters occur when immigration is turned down in these slow-sorting scenarios (Appendix M).

The lack of clustering in the low-depletion scenario of the Rosenzweig-MacArthur consumer-resource model reveals the connection between this scenario and the MacArthur-Levins model with Gaussian competition kernel (see Appendix L). Because in our consumer-resource model the gains in consumer population upon consumption are proportional to the depletion of the resource population, the competition coefficients arising from overlap in resource consumption form a positive-definite matrix (**?**), which has been linked to stable coexistence among species with arbitrarily similar traits and strong sensitivity to edge effects (**?**). Our results show that it also leads to slow niche sorting and therefore high sensitivity of pattern to immigration.

On the other hand, clustering appeared in some of our runs with high resource depletion. This could be because the extirpation of most resources forces species to compete for the same few surviving resources and thus clusters are formed based on specialization to specific resources (see Appendix M for stem plots of consumers and resources in both scenarios). Our finding that resource depletion has a strong effect on trait

---

[1]Note that our clustering results were stronger in the baseline MacArthur-Levins scenario than in the scenario with a circular axis, in apparent contradiction with this argument. The paradox is dissipated by considering that edge effects break the otherwise perfect symmetry of the axis and set the positions of the niches, thereby speeding up competitive sorting relative to a circular axis. This highlights our finding that even if the deterministic immigration-free model predicts clusters, when immigration occurs their appearance is affected by the relative rates of internal dynamics and immigration.

pattern is consistent with Abrams' results regarding the effect of depletion on competitive interactions and coexistence (**??**). However, we note that even though Abrams criticizes the assumption of similarity-based competition as simplistic for ignoring resource depletion, at least with regards to pattern our high depletion scenario is comparable to our baseline MacArthur-Levins scenario, which is based purely on similarity (discounting edge effects, which did not affect clustering in our simulations).

The negative results in the global-dispersal scenario of the Schwilk-Ackerly model of competition for suitable environments can also be understood via its connection with the MacArthur-Levins model. Fitness of species to specific environments are modeled as a Gaussian function of the difference between the local environment and the ideal environment for that species. Species compete based on their attempts to colonize the same spot. In the simpler case of no dispersal limitation, the competition coefficients result from overlaps in suitability to environmental conditions. **?** showed that coefficients originating from overlap of Gaussian resource use curves are themselves Gaussian functions of species similarity. We conclude that the global-dispersion version of the Schwilk-Ackerly model is directly analogous to the MacArthur-Levins model of Gaussian competition, and thus inherits its slow niche sorting and low ability to form pattern under high immigration.

In contrast, we found that the picture changes drastically when dispersal limitation is considered. Dispersal limitation lowers the rescue of populations that are ill-adapted to local environments by propagules from individuals located in more propitious areas. The slow niche sorting characteristic of this model is then compensated by what amounts to a dramatic reduction in "internal" immigration. This enhances local exclusion and therefore pattern formation.

It should be noted that Schwilk and Ackerly, along with **?**, found very strong pattern in this model, but marked by even spacing rather than clusters. There is no contradiction with our results because these authors used extremely low immigration rates. The MacArthur-Levins model, of which this model is an analogue, produces even spacing when immigration is absent. Our stochastic simulations will also produce even spacing for sufficiently low immigration rate. **?**, tested both very high and very low immigration rates. As expected, even spacing was observed at the low end. The picture forming here is that stochastic niche pattern falls on a spectrum with even spacing under little to no immigration, no pattern when high immigration overwhelms local dynamics, and clustering in between.

The hierarchy models were characterized by striking asymmetries on the niche axis,

as evidenced by the stem plots and the low Rao and Functional Dispersion indices. Despite these asymmetries, both models showed strong clustering, even though their niche mechanism is fundamentally different from the MacArthur-Levins model. In the competition-colonization tradeoff, the vast difference in size and abundances between the leftmost cluster and those on its right calls into question whether the outcome expresses coexistence at all, as low-abundance species are vulnerable to extinction due to drift, even if they are being stabilized by niche differences. As for the tolerance-fecundity tradeoff, clustering was present but results were complicated by the fact that the model predicts coexistence—and therefore pattern—only among species with high competitive ability.

Functional overdispersion and even spacing between species were notably absent from our results, in stark contrast to widely held expectations regarding the effects of niche differentiation. We note that the coefficient of variation has also been used to detect clustering (**?**)—a clustered community is supposed to have significantly high CV—but it did not identify clustering in any of our scenarios. We adapted the CV metric to take abundance structure into consideration and saw that the clusters—and not the species—may be evenly spaced. We propose that this abundance-weighted extension of the CV metric can be further developed into a metric that could simultaneously detect even spacing between niches and clustering between species, as may be the case when regional diversity far exceeds the number of niches available.

The relationship between clustering and competition based on similarity can be linked back to MacArthur and Levins's results (**?**) regarding invasion of a two-species community by a third species with an intermediate trait value. Coexistence between the three species is only possible if they are separated by a minimum distance on the trait axis (limiting similarity), but when that condition is not met, the combined competitive suppression caused by the two residents is lower on an invader closer to one resident than on an invader in the middle. Extending this argument to multispecies cases, the total competition felt by a species is lower near one of the niche-differentiated residents than in the center of the interspersing gaps. Therefore in immigration-free assembly with a community initially supersaturated with species, transient clusters spontaneously form because competitive exclusion is slower in the immediate vicinity of niche-differentiated residents than in the center of the interspersing gaps. In our stochastic models, this transient state is made permanent by immigration and mass effects, but can be overwhelmed if the species sorting is too slow.

The appearance of clusters under niche mechanisms that are not expressly based on

trait similarity is potentially due to the fact that even in these other models competitive interactions do tend to decline with trait difference, if less symmetrically than in the MacArthur-Levins model. Some degree of similarity-based competition is the quintessential element of niche differentiation, and stabilization is probably impossible without it. Niche models that predict coexistence can be expected to contain some form of similarity-based competition, and therefore potential for clustering. However, our study shows that whether clusters will arise depends on many other factors such as immigration, speed of sorting, the species-to-niches ratio, and randomness across pairwise species interactions. Our results indicate that depending on these factors niche pattern may fall somewhere on a spectrum between even spacing on one end and no pattern on the other, with clustering in between.

Although our results were obtained for niche differentiation along a single axis, our conclusions hold in general for multidimensional niche space. If competition originates primarily from species similarity in an n-dimensional space, species under stochastic assembly with immigration will form n-dimensional clusters (results not shown). Our metric can quantify clusters in any dimension, so long as a distance function is provided.

Our study is limited by the assumption of a neutral regional pool. In real systems, the pool is likely to be under its own biotic and abiotic filtering and may display pattern which will be inherited to some degree by the local community. In addition, all our niche scenarios except for the Schwilk-Ackerly model with local dispersal are spatially implicit and therefore ignore dispersal limitation within the local community. Thus our immigration rate may overestimate propagule pressure in real systems, where most immigrants would originate from a smaller pool of nearby individuals. Particularly in the case of the life history tradeoff models presented here, dispersal limitation may have a considerable impact on coexistence and the formation of pattern. Combined with the fact that immigration tends to dilute pattern, we conclude that our results are probably a conservative estimate of the potential for clustering in the scenarios tested.

In sum, our manipulations of the MacArthur-Levins model and examination across different competition mechanisms found that clustering is a general phenomenon in stochastic niche assembly under immigration, emerging from competitive dynamics dominated by species similarity. We provided a tool to find and quantify clusters in data. Our study also delimited the circumstances under which clusters are more likely to be observed, namely low immigration, high regional diversity, and fast competitive sorting. Our findings highlight the importance of tying pattern predictions to modeling assays that come as close as possible to realistic scenarios.

6.5 Figures

Figure 6.1: **A-C**: Outcomes of stochastic assembly with immigration of three MacArthur-Levins scenarios. **D-F**: Gap statistic as a function of number of clusters fitted to the stochastic data. The null expectation of the gap statistic for any given number of clusters is zero, as marked by the dashed grey line. The peak of the curve is the test statistic returned by the metric. **G-I**: Outcome of deterministic immigration-free simulations of the same scenarios. Every species from the regional pool are initially present, and competitive dynamics proceeds until only stably coexisting species remain. Left: baseline scenario. Center: scenario with slow niche sorting ($\rho = 2$). Right: scenario with random w (see Methods for details.)
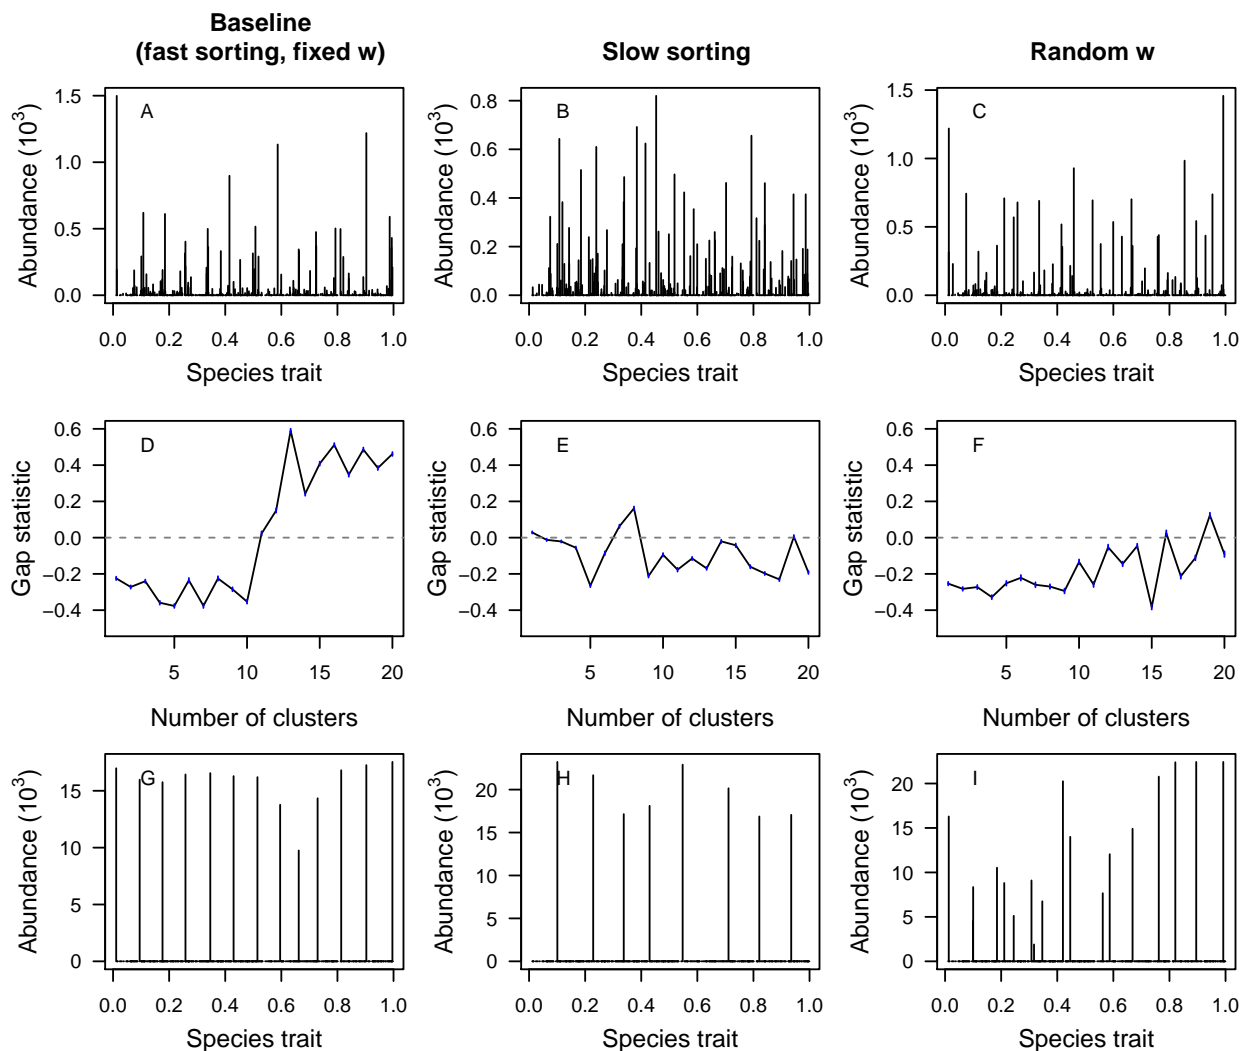
Figure 6.2: Average standard score of the gap statistic across 10 runs of each MacArthur-Levins scenario. Error bars show standard error of the mean. **A**: Comparison of the standard score of the baseline case against variants with slower competitive sorting or reduced symmetries in the competitive interactions. See Methods for full description of each scenario. **B**: Comparison among variations of the baseline scenario with different immigration rates. Labels show immigration rate used in simulation. **C**: Comparison among cases with different regional diversity. Labels show neutral diversity parameter θ used to generate regional pool. Red lines mark a deviation from the neutral mean of 1.645 standard deviations, which for normally distributed null values corresponds to the 95% percentile.
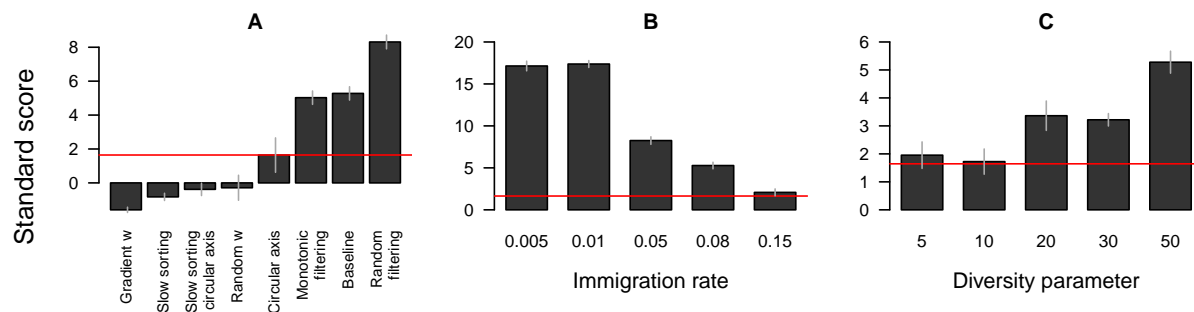
Figure 6.3: Stem plots representing communities generated from models with various niche mechanisms. **Left**: Niche differentiation by preferred resource. In the scenario shown in the top panel, parameters are set so that resources are not severely depleted by consumption; in the bottom scenario, consumption causes extensive resource depletion and extinction. **Center**: Niche differentiation by preferred abiotic environment. In the top scenario, species are not limited by dispersal: the offspring of every individual compete for every vacancy. In the bottom scenario, dispersal declines with distance to the vacated site. **Right**: Niche differentiation by life history strategy. Top: Competition-colonization tradeoff. Numbers in red show abundances of species exceeding 1,000 individuals (marked by red dots). Bottom: Tolerance-fecundity tradeoff.
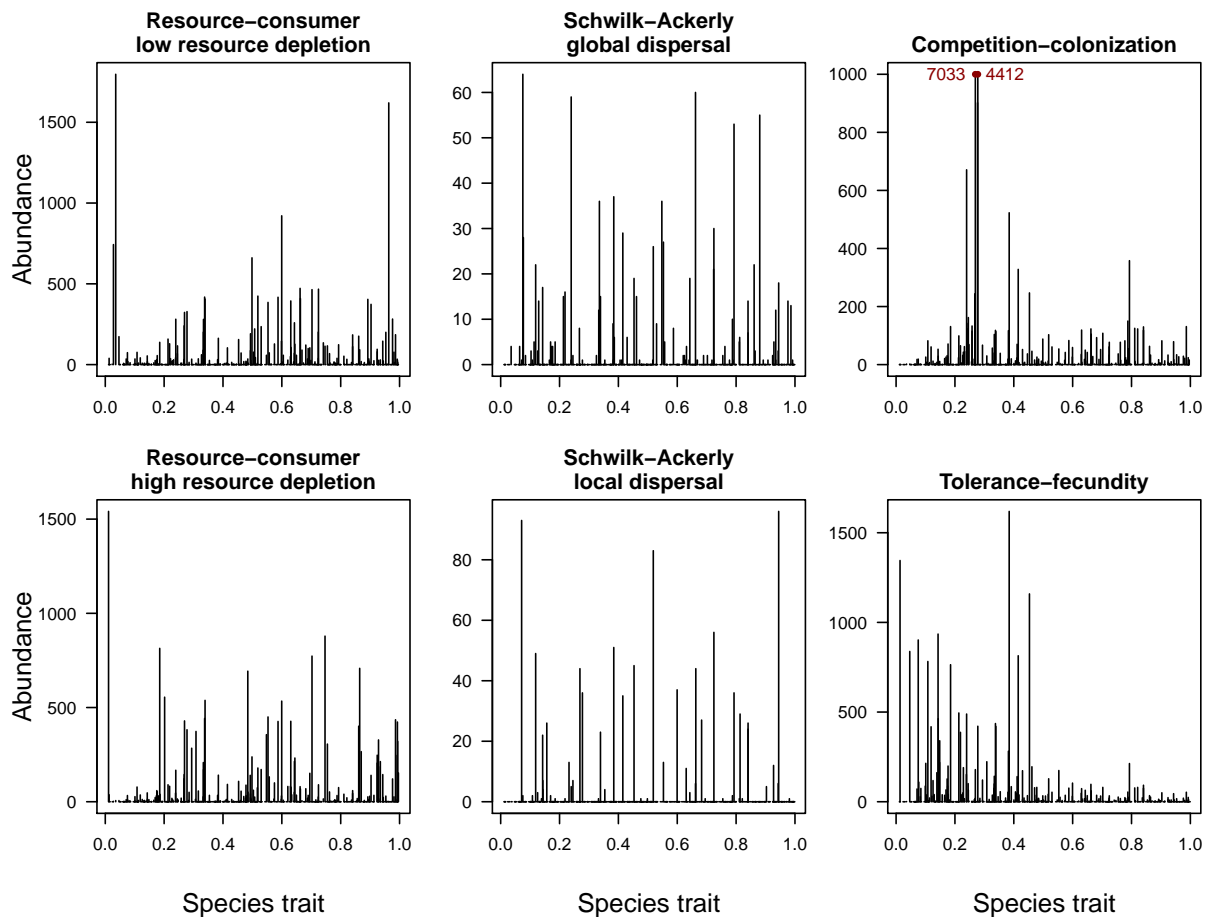
Figure 6.4: Average standard scores of the gap statistic of communities generated from competition models with several niche mechanisms. Clustering was observed in all of the mechanisms tested, although not in every scenario; within the same type of niche mechanism, clustering varied significantly depending on factors such as resource depletion and dispersal limitation. The gap statistic of communities estimated to have a single cluster was set to zero, so scenarios that fit a single cluster appear as negative bars. This was the case of the tolerance-fecundity tradeoff and the pure filtering scenario, characterized by purifying selection with no niche differences ($A_{ij} = 1$).
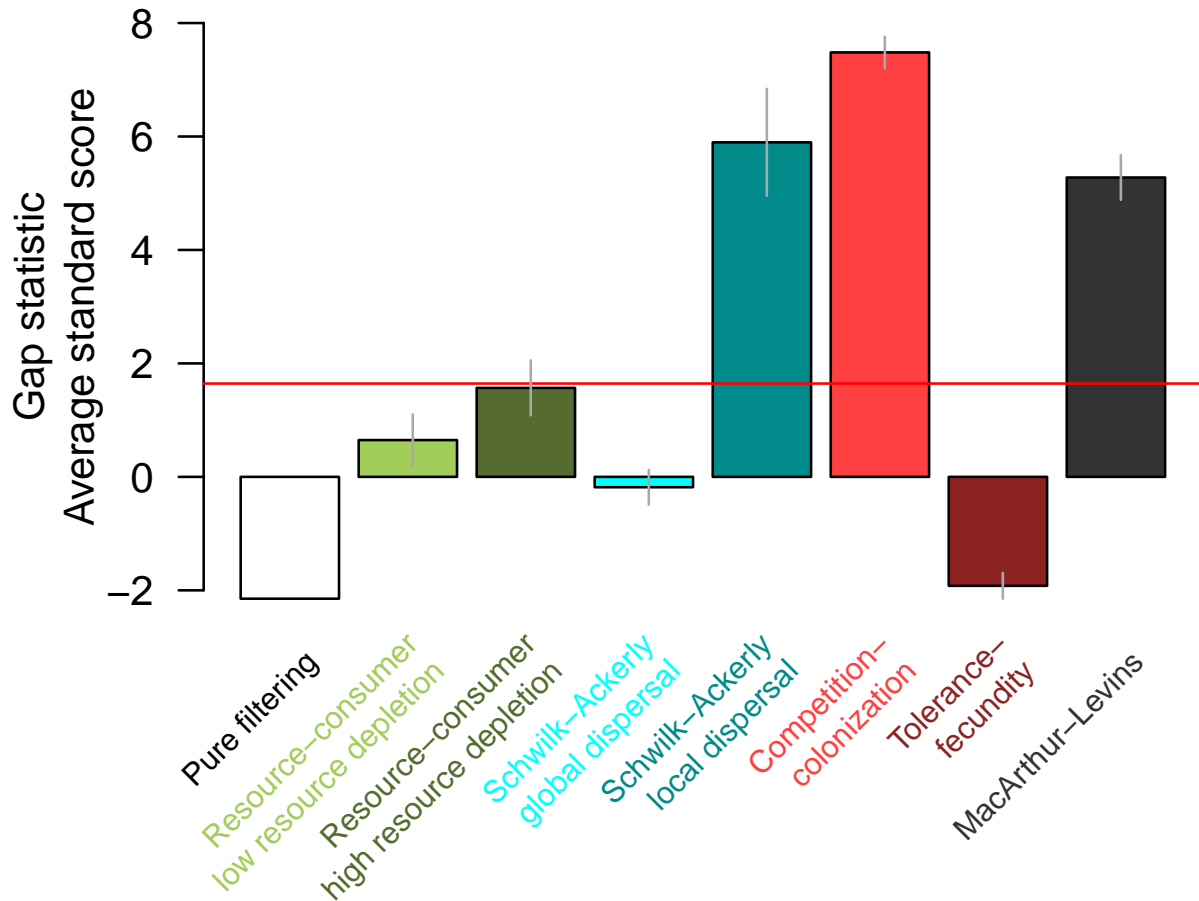
Figure 6.5: Immigration links species lifetimes to regional abundances. Red dots indicate species identified as cluster centers (medoids). The relationship is very strong in the neutral scenario, but less so in the niche scenarios, particularly when niche sorting is fast. Both in neutral and niche cases, species that persisted throughout the entire tracking period in the simulation ("residents") tended to have high regional abundance; however, in the fast sorting niche case many residents are not of particularly high regional abundance. Niche scenarios shown are the MacArthur-Levins baseline and the MacArthur-Levins with Gaussian competition coefficients.

Figure 6.6: **A**: Standard scores of the Rao quadratic index, the functional dispersion index (FDis), and the coefficient of variation in distances between adjacent neighbors on the trait axis (CV). The red lines mark the threshold $z = \pm 1.96$, corresponding to the 95% confi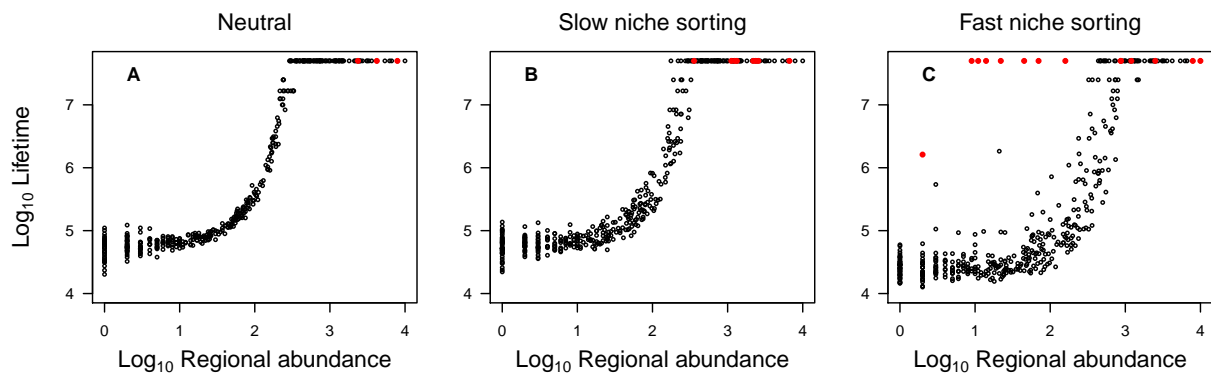dence interval for the standard score against normally distributed null values. Rao and FDis were not significantly high in any of our niche scenarios, and were low in the hierarchy models as well as the similarity-based models with environmental filtering. CV was not significantly high or low in any scenario. **B**: Coefficient of variation of the trait differences between adjacent species in subsets with the n most abundant species, with n between 5 and the total number of species observed, S. Pink band shows the 95% confidence interval from the set of 1,000 neutral runs. The CV is clearly lower than neutral when n is similar to the estimated number of clusters $\hat{k}$ (indicated by the red dot), and clearly higher than neutral when n is intermediate between $\hat{k}$ and S, but falls to null expectations as n approaches S.

114
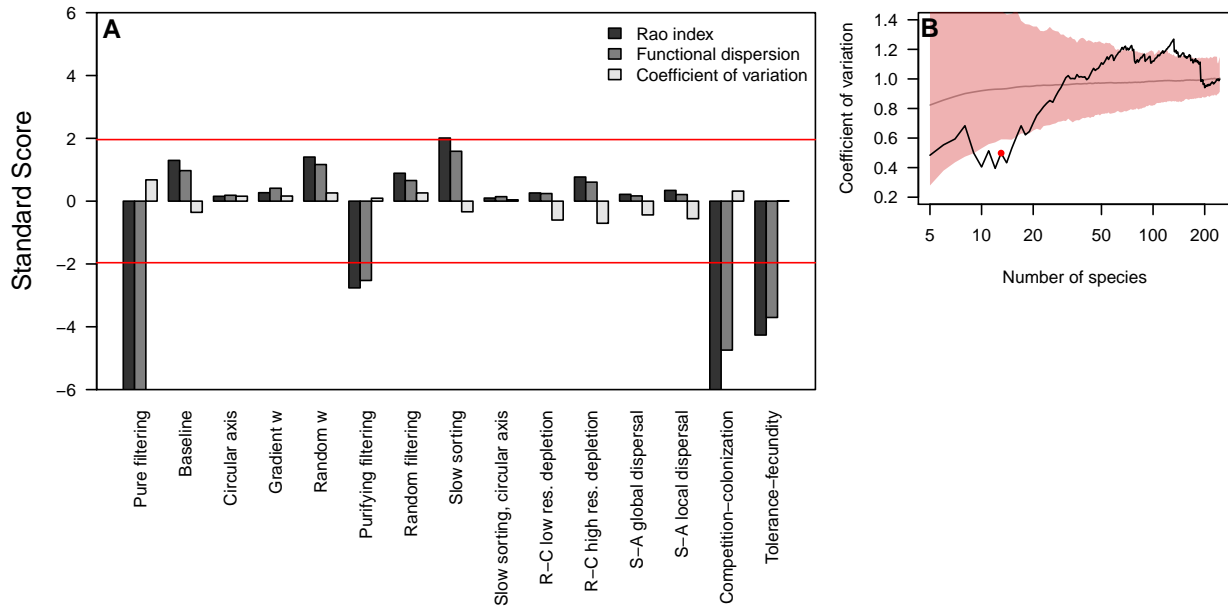
Chapter 7

Can clustering in genotype space reveal niches?

Abstract

Community ecology lacks the success enjoyed by population genetics to quantify the relative roles of deterministic and stochastic processes. Jeraldo et al. (2012) proposed that clustered patterns of abundance in genotype space provide evidence of selection in microbial communities, since no such clustering would arise in the absence of selection. We critique this test for its unrealistic null hypothesis. We show mathematically and with simulations that point mutations alone lead to clustering in genotype space by causing correlations between abundances of similar genotypes. We also show potential deviations from the mutation-only pattern caused by immigration from a source pool. Clustered patterns in genotype space may still be revealing of selection if analyzed quantitatively, but only if neutral and selective regimes can be distinguished once mutation and immigration are included in the null model.

7.1 Introduction

Population geneticists have at their disposal various techniques for detecting selection at the gene level, including null-hypothesis tests against neutral models calibrated with data from a large number of genes (e.g. **??**), and cross-species tests such as the McDonald-Kreitman and the dN/dS test (**?**). However, assessing the relative roles of selection and neutral forces in community assembly has proven a challenge. **?** have contributed a step forward by offering a test for selection's influence on community assembly in microbial communities based on patterns in genotype space. This method, if reliable, would be very

useful, as it needs only a combination of abundance and genomic data that are widely available in microbial ecology studies.

? argued that clustering of operational taxonomic units (OTUs) in genotype space can be used to infer selection. Jeraldo et al. refer to selection in general as "niche stabilization". This use of the phrase differs from the current predominant usage in community ecology (e.g. **?**), where niche stabilization is specifically reserved for stable coexistence promoted by species differences. Still, clear signatures of selection in the broader sense meant by the authors would be of great interest to ecologists. Jeraldo et al. found that within the gastrointestinal microbiomes of domesticated vertebrates, rare OTUs tend to be close in genotype space to abundant ones, and argued that this is qualitatively different from the pattern expected under neutrality, where rare and abundant OTUs should be randomly located in genotype space.

This paper has already generated some interest, having typically been cited as existing evidence that selective forces influence the assembly of gut microbial communities (e.g. **?**). The test was introduced in the context of gut microbiomes and applied to a gene known to be under selection due to its functional importance (**?**), but it would apply in general to genes of unknown fitness effects in any microbial community with high genetic diversity. Furthermore, it tells whether selection on that gene is influencing the assembly of OTUs into a community.

Here we argue that a more thorough consideration of the appropriate null model for abundance patterns in genotype space is needed before Jeraldo et al.'s approach can be accepted. In particular, we argue that Jeraldo et al. did not consider the potential influences of mutation and immigration in the neutral scenario. We show that point mutations alone are expected to lead to clustering in genotype space, and that immigration may lead to the local community mirroring pattern existent in the source. We argue that the presence of clustering in genotype space is typically not enough to infer selection. Instead it must be inferred by demonstrating quantitative departures from patterns expected from mutation and immigration alone. We also discuss the particular role mutations and immigration from a source pool might play in the gut microbiome system to which Jeraldo et al. applied their test, and show how they should be factored into a more informative null model for the test.

7.2 The Jeraldo et al. test

The Jeraldo et al. test looks at operational taxonomic units (OTUs) in a microbial community. An OTU is generally the taxonomic level of sampling selected to be used in a study, and in this context is essentially a "species" concept for microbes based on genetic similarity. For the purposes of describing the qualitative properties of abundance pattern in genotype space identified by Jeraldo et al., the species/OTU label is unnecessary. Henceforth we will refer to genotypes in a community, with no regard for which species/OTU they belong.

Consider a community of asexual individuals, each characterized by its genotype within a certain genetic sequence that is polymorphic. Ranking genotypes by abundance, let the 5% most abundant ones be called "modal", and call all others "rare". For each of the rare genotypes, find the modal genotype with the shortest Hamming distance to it, i.e. the smallest number of discordant loci. Jeraldo et al. (2012) proposed that the distribution of such shortest distances between rare and abundant genotypes can be used to infer the influence of selection among genotypes. (Note that Jeraldo et al. frame the test in terms of OTUs.)

Jeraldo et al. reasoned that in the case of selection, the community will consist of a few modal genotypes (those being most selected for), and many more rarer genotypes having arisen through point mutation from the modal ones, with the likelihood of a particular genotype being present declining with increasing number of allelic differences (Hamming distance) due to decreasing fitness. Put simply, assuming similar genotypes have similar fitnesses, the fittest genotypes would be most abundant, followed by their immediate neighbors in genotype space, and so on. They then showed that static communities constructed that way show a monotonic distribution of shortest Hamming distances between rare and modal genotypes, with shorter distances prevailing over larger ones.

On the other hand, Jeraldo et al. claimed that in the absence of selection, rare and modal genotypes would be randomly distributed in genotype space, resulting in a peak in the distribution of distances to nearest modal genotype at an intermediate Hamming distance. Essentially, Jeraldo et al. theorized that the curve in the neutral scenario would track the number of neighbors by Hamming distance H to a given genotype, which peaks at $H = 3/4$. They confirmed their expectations by measuring the Hamming distance curve on randomly constructed communities.

Jeraldo et al. then applied their test to microbiomes collected from the guts of swine,

cattle, and chickens, and rejected neutrality because the observed distributions of distances to nearest modal OTU were a better fit to a monotonic distribution than a distribution with a peak.

7.3 Limitations of Jeraldo et al.'s null hypothesis

The authors' null prediction of a curve with a peak reflects their assumption that the abundances of different genotypes would be entirely uncorrelated in the absence of selection. However, this assumption can be violated in a number of ways, most notably by the effects of mutations and immigration from a source with a nonrandom abundance distribution.

7.3.1 The effect of mutations

Consider a community of microbes where individuals die and reproduce asexually with a chance of mutation under a purely neutral regime. Individuals differ with respect to their genotype in a genetic segment of specific length. Suppose there can only be point mutations, that is, mutations that replace a single base nucleotide with another nucleotide. In Appendix N we prove that in the long term genotype abundances are correlated, and the correlation declines monotonically with the Hamming distance between them. This means that genotypes with short Hamming distance will tend to have similar abundances compared to pairs with long distance. Intuitively, those are precisely the conditions necessary for the formation of clusters in genotype space that Jeraldo et al. associate with selection (Fig. 7.1).

To show that these correlations do in fact lead to this patterning in Jeraldo et al.'s metric, we simulated neutral communities with point mutations. Our simulated neutral communities have 1,000 individuals, each characterized by a genotype consisting of 40 loci. We draw the genotypes of the initial individuals from a uniform probability distribution, and run zero-sum neutral dynamics (i.e. community size is kept fixed) for 2 million simulation steps, corresponding to 2,000 turnovers. At each simulation step, one individual is selected at random to die, and another is selected to reproduce. It then either sires a clone or a mutant with one point mutation (i.e. a substitution in a single nucleotide). We set the mutation rate $\mu = 0.1$, so every newborn has 10% chance of carrying a point mutation.

As expected, we obtained a monotonic distribution of distances to nearest modal genotype, qualitatively similar to Jeraldo et al.'s expectation for the selection scenario

(Fig. 7.2). In contrast, in simulated neutral communities that are not constrained by point mutations, i.e. where each genotype can mutate with equal probability into any other in a single step, we see a mode at a Hamming distance of 3/4, matching Jeraldo et al.'s expectations exactly.

We used a very high mutation rate to save computation time, but we checked that our results are qualitatively the same for mutation rates several orders of magnitude lower, $\mu = 10^{-3}$ and $\mu = 10^{-4}$. Estimates for actual mutation rates in bacteria hover around 10^{-11} substitutions per site per generation (**?**), which would require years of simulation time with our code. However, given that our analytical result for the correlations applies to any mutation rate, we see no reason to expect any qualitative difference in this distribution for lower mutation rates. We note that organizing genotypes into OTUs would not eliminate correlations and hence would not qualitatively change our results.

So when structure caused by neutral mutations is considered, Jeraldo et al.'s test is unable to infer selection based upon a simple qualitative assessment of the shape of the curve. But can it do so upon a quantitative assessment? The answer depends on the array of possibilities that exist for the strength of selection, mutation rates, and community size. In Appendix P, we carry out an example analysis using simulations to illustrate how this question might be answered for specific systems in future studies.

7.3.2 Immigration

To test how immigration can affect the neutral pattern caused by mutations, we added immigration to our simulations of neutral dynamics with mutation and compared with the mutation-only scenario. We simulated local communities with 1,000 individuals receiving immigrants from a fixed source pool. We considered two different types of sources: one lacking any structure, where genotype abundances are drawn at random from a uniform distribution, and one that is shaped by mutation and selection. The structured source was generated through simulation as described above, with ten times the size of the local community (see Appendix P for a description of the selection case considered). For the unstructured source we simply chose genotypes randomly for immigrating individuals when an immigration event was to occur. The intensity of immigration is controlled by the immigration parameter m, which sets the probability that a death event is replaced by an immigrant from the pool instead of a local birth. We tested three different values for the immigration parameter, $m = 0.001$, $m = 0.01$, $m = 0.1$, representing low, medium, and high contribution of immigration relative to mutation (set at $\mu = 0.1$ for all simulations).

Results are shown in Fig. 7.3. When immigration is low compared to mutation, the effect on the local Hamming distance curve is minimal. For higher immigration rates, the effect of immigration depends on the abundance structure in the pool. Immigration from a pool without structure adds a mode and causes the monotonic part to be steeper and narrower. In contrast, immigration from a pool shaped by mutation and selection adds no mode, but deforms the monotonic part if the immigration rate is sufficiently high.

In summary, structure in the source pool, or lack thereof, will influence the local community Hamming distance curve if immigration is substantial enough compared with mutation, and hence must be factored into the null model to detect selection occurring at a local scale.

7.3.3   The gut microbiome

The results above draw a general picture for a process-based null model for the Jeraldo et al. test that combines mutation and immigration. In this section, we discuss the appropriate null model for the specific case of the gut microbiome, to which the test was applied.

Assuming about 10^{13} microbial cells in a cow's gut (**?**), a mutation rate of 10^{-11} mutations per site per cell cycle (**?**), and a gene with 1 kbp, we can expect about 10^5 new mutations in that gene per generation in the gut. It thus seems likely that mutations would matter for the dynamics (for comparison, our simulated communities had an average of 10^2 mutations per generation).

Ingested microbes have been shown to impact gut flora composition (**?**). It is thus conceivable that immigration may affect local structure as well. The main modes of transmission for gut-living microbes are presumably maternal sources (**?**) and the fecal-oral route. Grazers such as cattle regularly ingest soil microbes, but experiments with mice indicate that xenomicrobiota from soil and the guts of other species are poor competitors compared with autochthonous microbiota **?**, and are thus unlikely to significantly affect gut community structure within a host species. Thus one can assume that in addition to drift, selection, and local mutations, structure in the gut microbiome of an individual is influenced by the legacy of maternal sources and ingestion of autochthonous microbes, and to a lesser extent, xenomicrobes.

There are thus a few possible null models for gut microbiomes, depending in part on the scale at which one wishes to infer selection. If one is interested in selection within the gut of an individual host, the null model would be a neutral community with initial structure

based on the gut microbiome of the individual's mother (due to the prominence of maternal sources early in life) and dynamics driven by drift, mutation, and immigration from ingestion of autochthonous gut microbes. Such a null model would require information about the composition and relative abundance of genotypes in the host population. In the absence of that data, one might model the pool as structured by mutation. The distribution of Hamming distances is expected to be monotonic (Fig. 7.3). On the other hand, if one aims to infer selection at the level of the host population, immigration can probably be ignored, as recruitment of xenomicrobes is infrequent compared to local mutations. The null model would be a community with dynamics driven by drift and mutation.

7.4   Discussion

If selection favors one or a few genotypes over others, and if similar genotypes have similar fitnesses, one expects the abundance of a genotype to be negatively associated with its Hamming distance to the fittest genotypes. Genotype space should be characterized by regions of high abundance near the fittest genotypes and regions of low abundance away from them. This idea has been used to infer selection among OTUs in microbiomes, but the implementation was problematic because it ignored the fact that neutral dynamics under point mutation can create similar abundance patterns. Jeraldo et al.'s null hypothesis is based on oversimplified genotype-blind dynamics. Once point mutations are accounted for, the test fails to distinguish selective from neutral regimes based solely on a qualitative assessment. A quantitative approach is required.

We have shown how mutation and immigration contribute to neutral predictions to be used as a null hypothesis to the test proposed by Jeraldo et al. If immigration events are rare compared to mutation events, mutation dominates structure. The neutral model with mutation described here can then be used as a process-based null model for the Jeraldo et al. test. Given the low values of mutation rates in actual systems, simulations of that model will need to use parallel computing, as serial computing is impractical.

When immigration is significant, the null prediction will depend on the immigration rate and the abundance structure in the source pool. If the pool has weak or no structure, then immigration adds noise and brings the curve towards the random null hypothesis used by Jeraldo et al. But if the pool itself is shaped by selection, then the shape of the local Hamming distance curve is qualitatively similar to the mutation-only case, but influenced quantitatively by the structure in the source pool.

Jeraldo et al. were not specific about whether they aimed to detect selection at the scale of the individual host or the host population, and made no mention of a potential source pool. Their null hypothesis corresponds to a case where mutations are so slow as to be non-existent in the local community, and where immigration comes from an unstructured pool. Their data should be reanalyzed in the context of a mutation-only null model in order to gain evidence for selection at the level of the host population, or in the context of a null model structured by mutation and ingestion of autochthonous microbes for detecting selection at the level of an individual host.

Issues with the test's null hypothesis notwithstanding, its limitations in scope must be noted. The test looks for selection, not niches: in taking abundance as a proxy for fitness, the test is blind to niches that only allow for a low carrying capacity, such as specializing on a scarce resource. Also, the selective regime that the authors test for is somewhat specific. The fitness landscape is assumed smooth, which is not realistic when fitness does not correlate strongly with Hamming distance. The test is therefore better suited to sequences in which mutations have cumulative small effects on fitness. Finally, the test can only be applied if there is enough variation present for clustering patterns to possibly emerge. This was mentioned by Jeraldo et al., but bears emphasizing, since it means the test can only be successful with highly variable genes.

One further use of our analytical work in Appendix N is as a precursor for an alternative metric based on the correlations themselves. We offer a partial derivation, where correlations are obtained as a function of the mean and variance of observed genotype abundances. To complete it, one must find predictions for the mean and variance as a function of population size and mutation rate. That could be complemented with simulations to establish variability in the correlations and carry out statistical analysis. Preliminary investigation of this suggests that a test based on correlations would have at least as much power, if not more, as the test proposed by Jeraldo et al.

Our analytic results mirror previous findings of clustering of organisms in a spatial context (**?**). Our model also has parallels with quasispecies models, which describe populations of asexual replicators at a high mutation rate and are thus of interest in virology (**?**). For a comparison between these types of models and conventional population genetics models, see **?**. Finally, there are obvious parallels between genotypes in genotype space and species in multidimensional trait space, and between point mutations and sympatric speciation with a phylogenetic signal. This suggests that the Jeraldo et al. test can also in principle be applied to detect clustering of species in trait space caused by

habitat filtering and niche processes (**?**).

7.5 Figures

Figure 7.1: Conceptual illustration of genotype space: clustered versus random. Circles represent genotypes, with size indicating genotype abundance. Positions of the circles mark their placement in a multidimensional genotype space, here represented in two dimensions. **A**: Neutral scenario envisioned by Jeraldo et al. : abundance of genotypes unrelated to their location in genotype space. **B**: Clustering pattern: the immediate neighborhood of particularly abundant genotypes consists of other relatively abundant genotypes, whereas distant regions are populated with rare genotypes. Jeraldo et al. associate this pattern with selection, but it can also appear due to mutation.
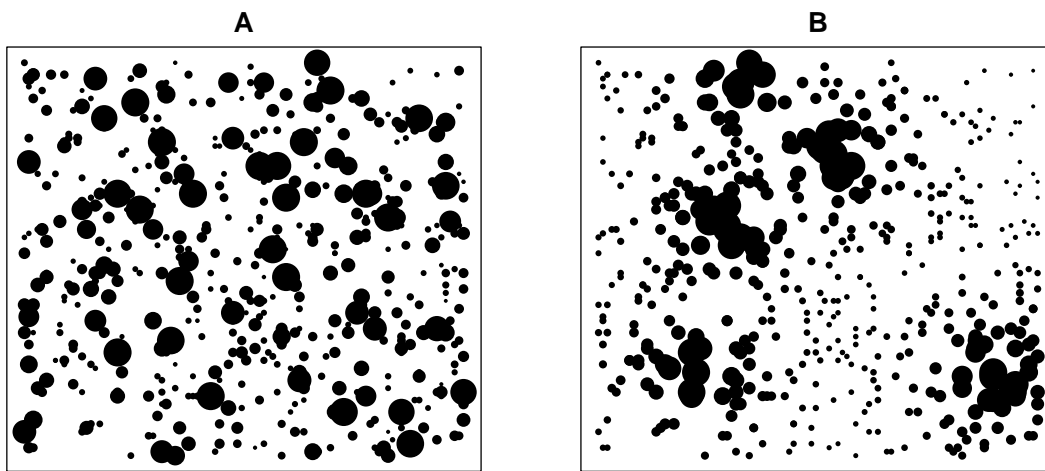
Figure 7.2: Distribution of distances to nearest modal genotype observed in neutral communities driven by drift and mutation. (normalized Hamming distance is defined as z/L, where z is the number of discordant loci and L is the total number of loci.) Numbers on the y axis reflect the frequency of rare-modal pairs. Black curve corresponds to neutral scenario with restriction to point mutations, and grey curve corresponds to neutral scenario with mutations of unconstrained size. Curves and error bars represent mean and standard error of the mean across 127 (point mutation) and 163 (mutation of unconstrained size) runs. Community size for all runs is 1,000, and mutation rate is 0.1 per genome per generation, for a genome with 40 loci.
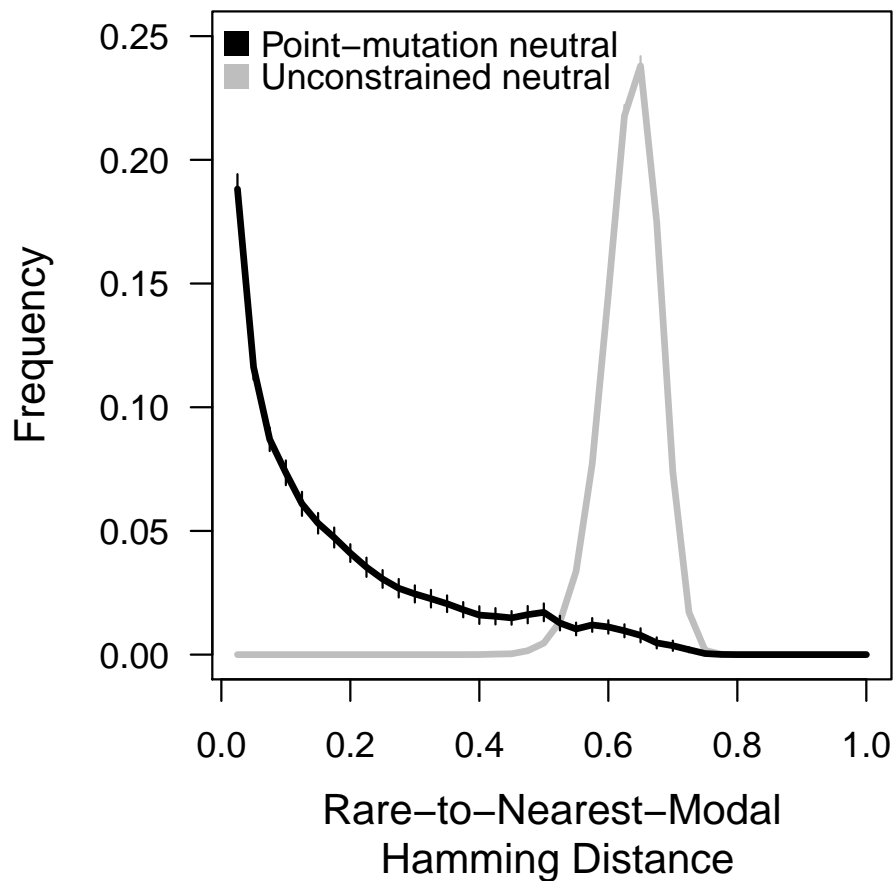
Figure 7.3: Comparison between neutral communities with and without immigration, for three different immigration rates *m* (columns) and immigration from source pools with different abundance structures (rows). Grey curves correspond to communities subject only to point mutation at fixed rate $\mu = 0.1$, while black curves correspond to communities under both mutation and immigration. The last column shows the source pools of the corresponding rows.
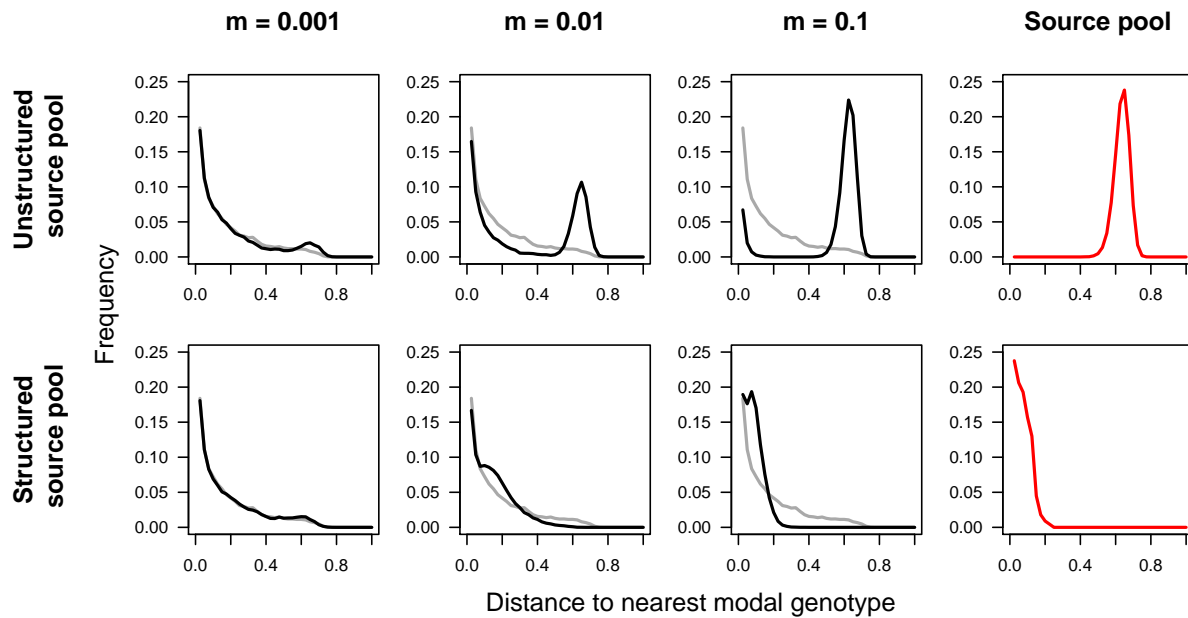
Chapter 8

Dissertation summary and closing thoughts

8.1 Dissertation summary

Despite being a central tenet of ecological theory, niche differentiation can be difficult to verify, and attempts at empirical measurement have met with mixed results (**??**). Chapter 1 promoted inference based on trait pattern as being generalizable across systems, directly tied to ecological strategy, and less data-intensive than other methods, but argued that recent findings from dynamical niche models require a new round of revisions in expectations of pattern. Key themes included the need for developing theory across models and the need for metrics better suited to capture pattern the models are telling us to look for.

Chapter 2 investigated the consequences for model behavior of the problematic assumption of discontinuity in ecological strategy. This assumption, while hopelessly unrealistic, has found its way into a variety of models spanning different niche mechanisms (**???????**). We showed that such discontinuity may raise coexistence to arbitrary levels and therefore drastically affects predicted pattern. In Chapter 3, we offered revised formulations of two important niche models that implicitly make this assumption as originally presented, namely the tolerance-fecundity tradeoff (**?**) and the competition-colonization tradeoff (**?**). These revised formulations are essential for studying trait pattern on these models.

An important step in linking pattern and process is to gain a clear picture of the dynamical influence of the process. Niches promote coexistence in deterministic models (**?**), but what do niches do to species dynamics in a stochastic model, and how does immigration play into it? In Chapter 4 we set out to understand what niche differentiation is actually doing to species lifetimes and in what sense it enhances them. We showed that a partially niche-differentiated assemblage, where the number of species regionally present exceeds the number of niches available at the local scale, may actually contain

fewer species and display higher species turnover than a completely neutral assemblage. This chapter also related the effect of niches on species lifetimes to species traits, and those results support ideas further explored in Chapter 5 that clusters are directly related to the varying pace of competitive sorting along the trait axis.

Chapter 5 performed the model exploration called for in Chapter 1, and found that persistent clustering is a general outcome in stochastic niche assembly under immigration, especially when competitive sorting is driven by trait similarity. On the other hand, we also showed that pattern may not form if immigration rates greatly exceed the speed of competitive sorting. Finally, we showed that functional overdispersion and even spacing between species does not occur under the types of niche differentiation mechanisms explored. This chapter provided the first abundance-weighted metric specifically designed to find and quantify clustering in ecological data. Combined, Chapters 4 and 5 paint a picture of clusters as groups of species vying for the same niche, their appearance being a consequence of the fact that niche differentiation enhances the lifetimes of those species centrally situated within their clusters but expedites exclusion of those located in gaps.

We also argued for refining inference methods with process-based null models. Whereas in Chapter 5 only immigration is considered, Chapter 6 showed that if an OTU/species can generate similar OTUs/species by random mutation/speciation, then clustering will appear even in the absence of niche mechanisms. Statistical null hypotheses based on data randomization ignore this, and may lead to false positive results. Together, Chapters 5 and 6 make the case for clustering in niche inference by use of a process-based null model that includes the relevant neutral processes.

8.2   Closing thoughts

Niche differentiation means higher species richness, lower extinction rates, longer lifetimes, functional overdispersion, and even spacing on trait axes. Without niche differentiation, traits and abundances are randomly distributed across species. All of the above are commonly held beliefs about the impact of niche differentiation on species assemblages, and in this dissertation each of them were shown not to hold generally in the context of stochastic assembly with immigration and/or speciation, and when pattern-forming neutral processes are considered in the null model. This dissertation emphasized the need to validate inference methods using niche models. This approach clarifies in which circumstances specific methods and metrics are expected to work, and when specific

expectations regarding the impact of niche assembly on coexistence and pattern are likely to be met. Using this approach, we presented a refined theoretical expectation, that under stochastic niche assembly with immigration from a diverse regional pool, clustering on trait axes is expected. We also provided a metric for measuring clustering pattern, which can be used to compare observations with neutral process-based nulls that factor in the potential for neutral clustering.

We stated in the Introduction that the main goal of this dissertation was to "advance theory connecting niche differentiation to coexistence and pattern, and contribute to improving niche inference methods." This was accomplished by showing the following: i. Avoidance of discontinuity in ecological strategies should lead to more realistic niche models with more realistic pattern predictions; ii. More widespread use of process-based null models should lead to inference tests with reduced risk of type I and type II errors; iii. Measuring extinction rates alone may be insufficient to reject neutrality, as species turnover is not necessarily lower under niche differentiation than under neutrality; iv. Clustering should be brought to mainstream niche inference, and we have provided a tool to assist in that purpose.

A theme has emerged of clustering as a potentially general outcome of niche differentiation. We have argued that the interplay of immigration and competitive sorting leads to clustering in stochastic niche models where competition is related to trait similarity. An inverse relation between intensity of competition and trait differences, at least between very similar species, is the quintessential property of niche differentiation, without which stabilization is impossible—intraspecific competition cannot exceed interspecific competition otherwise. By tying similarity-based competition with clustering, we have then opened the floor to speculation that clustering on niche axes could be, to be liberal with language, a new "law" of ecology (**??**). Whether this is the case will depend on if the results of Chapter 5 are borne out across a wider set of niche mechanisms, and if the connections between niche stabilization and similarity-based competition can be more firmly established mathematically and/or empirically.

Going forward, new metrics of trait pattern should be developed in lockstep with model outcomes, as was done here with clustering and the gap statistic. Existing metrics can be coopted to reflect new findings: for example, the abundance-weighted extension of the coefficient of variation offered in Chapter 5 can be further developed intro a metric, which would be particularly useful for systems with low regional diversity. The relationship between the competition matrix and the speed of competitive exclusion must be better

understood, as it predicts which niche mechanisms are more likely to lead to pattern in communities under propagule pressure and other stochastic processes likely to dilute pattern.

Another important future direction is to better understand how disturbance and dispersal limitation may affect pattern. This is likely to be critical for niche mechanisms with strongly asymmetric competition, such as the competition-colonization tradeoff. This mechanism requires that higher fecundity compensate for lower competitive ability. In our results, low-fecundity strong competitors dominated over high-fecundity poor competitors. This could be because in a spatially implicit model with a saturated environment where every site is permanently occupied, as implemented here, competitive ability is likely to matter much more than fecundity. The tradeoff may maximize coexistence in an environment under continual disturbance, which periodically resets the exclusion process by temporarily favoring species with high fecundity, and if species are dispersal-limited, which slows down the exclusion of high-fecundity poor competitors by low-fecundity strong competitors. With these phenomena included, the competition-colonization tradeoff may predict more coexistence, less asymmetry in abundance, and stronger patterning than we found here.

Appendices

Appendix A

Niche mechanisms and spatial scales: life history tradeoffs and example plant traits

Niche mechanisms requiring spatial variation in the environment can involve tradeoffs between different life history components like those between species' ability to outcompete other species for a limiting resource in a suboptimal environment and their ability to colonize or quickly make use of newly available resources. Models falling into that category include the forest-architecture hypothesis (Kohyama 1993), the successional niche (Pacala and Rees 1998), the competition-colonization tradeoff (Tilman 1994), the dominance-discovery tradeoff (Adler et al. 2007), and the tolerance-fecundity tradeoff (Muller-Landau 2010). At a small scale, there will be filtering for species best able to colonize or make use of resources given the age or stress level of the patches encompassed, i.e. for species at a particular positioning on the life history tradeoff axis. At a large scale, there should be trait patterning associated with niche differentiation, since there will be some limit to the similarity in species' positioning along the tradeoff axis required for robust stable coexistence (Fig. 2.4).

The scale at which a given trait should display signs of niche differentiation relates to the scale of the niche mechanism at play. For example, leaf economics traits such as specific leaf area (SLA) can plausibly contribute to stable coexistence through stand dynamics. Low SLA indicates investment in slow growth and long-lived leaves, suggesting a climax strategy where the plant endures in shaded areas, whereas high SLA reflects fast growth and short-lived leaves, indicative of a pioneer strategy where the plant grows fast when the opportunity arises such as in gaps in the canopy (Wright et al. 2004). Hence variation

in SLA values might lead to coexistence in a forest composed of a patchwork of stands in different stages of the successional cycle. One would thus expect to see niche related patterning in SLA across stands of different successional stages, and not within a stand or among neighbors, where filtering would predominate (though the existing evidence is ambivalent, see Cavender-Bares et al. 2004).

In contrast, a trait like root profile may show niche-related patterning at the smallest scales. For instance, root systems may segregate to obtain water and nutrients from different soil depths (de Kroon et al. 2003). Changes in rooting strategies in response to the presence of other species have been documented, although the evidence for overdispersion in root profiles at the neighborhood scales is limited and contradictory, indicating additional processes such as facilitation and species-specific interactions beyond resource depletion (Schenk 2006). Traits associated with Janzen-Connell effects can also exhibit niche-related patterning at small spatial scales (Sedio and Ostling 2013). Some authors measured strong negative effects of conspecific seedling and adult densities on growth and survival of seedlings at small scales, dropping to background levels within just a few meters' distance (Hubbell et al. 2001). Such short-scale ranges raise questions about the actual role to coexistence and the maintenance of diversity at larger scales, although it has been shown to have an influence on patterns of relative abundance (Comita et al. 2010).

Appendix B

Existing research on the role of intraspecific variation to coexistence

The impact of intraspecific variation on trait dispersion starts with its impact on coexistence. Bolnick et al. (2011) suggest that intraspecific variation should be considered in understanding the outcome of species interactions whenever the strength of those interactions depends non-linearly on traits, because in that case the average interaction strength will differ from the interaction strength estimated from the mean trait value. Fig 5 in the main text shows that intraspecific variation has the potential to either increase or decrease interspecific competition, which suggests intraspecific variation may sometimes enhance coexistence and sometimes make it less likely. Indeed, existing work suggests intraspecific variation affects coexistence both ways.

A few experiments have been made to find correlations between interspecific variation (either in the form of genotypic diversity or trait variation) and coexistence (measured through species richness or community invasibility), with results ranging from positive (Booth and Grime 2003; Lankau and Strauss 2007) to non-significant (Weltzin et al. 2003) to negative (Crutsinger et al. 2008). As an example of the latter, Crutsinger et al. (2008) found that a plant community comprised of phenotypically diverse species was more resistant to invasion, thus restricting the number of species a community can support.

A variety of theoretical hypotheses and models of the impact of intraspecific variation on coexistence have also been offered, mostly suggesting that intraspecific variation should promote coexistence. For example, Aarssen (1983) proposed that it can reduce extinction risk by allowing for tolerance to spatially or temporally varying local conditions. Vellend

(2006) verified this in a Lotka-Volterra model where within-species genotypic diversity promotes coexistence by allowing species to span multiple viable niches and giving them flexibility to respond to selection by converging or diverging from competitors. Yamauchi & Miki (2009) further explored Vellend's model with the introduction of genetic flow between phenotypes, sexual inheritance of traits, and environmental stochasticity, and reported mixed effects of intraspecific variation on species diversity in simulation outcomes depending on the scenario tested. (It should be noted that the case where the facilitation of coexistence is most strongly observed, when environmental fluctuations correlate with niche positions, in effect harbors additional niche dimensions along which the species are differentiated.)

From the standpoint of Chesson's framework, by bringing species together in trait space, intraspecific variation may in effect reduce fitness differences between species and hence decrease the amount of stabilization required for coexistence (Chesson and Rees 2007). This could also slow down competitive exclusion, although Lichstein et al. (2007) argue that this effect on the speed of the dynamics would have little impact on observed diversity in real systems. Lichstein et al. (2007) concluded for a model with intraspecific variation in seed quality that such variation enables coexistence only when there is a tradeoff between species means and variances. How common such a tradeoff might be is unknown.

Other processes have also been proposed. If an intransitive loop (i.e. a "rock-paper-scissors" competitive hierarchy) occurs between the different phenotypes of a species and individuals from another species, then both species can persist without one species completely dominating the other (Lankau and Strauss 2007). Finally, rather than being a direct agent of coexistence, intraspecific variation could be associated with high niche-dimensionality, which in turn is an agent of stable coexistence. Clark et al. (2007) provides an example of this.

Appendix C

Two-species coexistence under smooth and kinked kernels

Let us consider two competing species in equilibrium, placed along a trait axis at trait values x_1 and x_2. We assume $x_2 > x_1$ without loss of generality. The equations read

$$\frac{\mathrm{d}n(x_1)}{\mathrm{d}t} = n(x_1)\big(r_0(x_1) - a(x_1, x_1)n(x_1) - a(x_1, x_2)n(x_2)\big), \tag{C.1}$$

$$\frac{\mathrm{d}n(x_2)}{\mathrm{d}t} = n(x_2)\big(r_0(x_2) - a(x_2, x_2)n(x_2) - a(x_2, x_1)n(x_1)\big). \tag{C.2}$$

If the two species are closely packed then the difference $\Delta x = x_2 - x_1$ between the strategies of the two species will be small. When this is so, several expansions become possible. First,

$$r_0(x_2) = r_0(x_1 + \Delta x) \approx \underbrace{r_0(x_1)}_{r_0} + \underbrace{\frac{\mathrm{d}r_0}{\mathrm{d}x}(x_1)}_{c} \Delta x = r_0 + c\Delta x, \tag{C.3}$$

where we introduced the notations r_0 and c for the value and the slope of the function $r_0(x)$ at $x = x_1$, respectively (we assume $r_0(x)$ is differentiable). Second, by introducing the function $A(x) = a(x, x)$, we get

$$a(x_2, x_2) = A(x_2) = A(x_1 + \Delta x) \approx \underbrace{A(x_1)}_{a_x} + \underbrace{\frac{\mathrm{d}A}{\mathrm{d}x}(x_1)}_{w} \Delta x = a_x + w\Delta x, \tag{C.4}$$

where $a_x = a(x_1, x_1)$ and w is the slope measuring the difference between the two intraspecific competition coefficients $a(x_1, x_1)$ and $a(x_2, x_2)$. Third, the interspecific competition

coefficients are expanded as

$$a(x_1, x_2) = a(x_1, x_1 + \Delta x) \approx a(x_1, x_1) + \underbrace{\partial_2 a(x_1, x_1^+)}_{-k_x} \Delta x + \frac{1}{2} \underbrace{\partial_2^2 a(x_1, x_1^+)}_{-d_x} \Delta x^2$$

$$= a_x - k_x \Delta x - \frac{d_x}{2} \Delta x^2 \tag{C.5}$$

and

$$a(x_2, x_1) = a(x_2, x_2 - \Delta x) \approx \underbrace{a(x_2, x_2)}_{a_y} - \underbrace{\partial_2 a(x_2, x_2^-)}_{k_y} \Delta x + \frac{1}{2} \underbrace{\partial_2^2 a(x_2, x_2^-)}_{-d_y} \Delta x^2$$

$$= a_y - k_y \Delta x - \frac{d_y}{2} \Delta x^2, \tag{C.6}$$

where $\partial_k^n a(x, y)$ is the nth partial derivative of a with respect to the kth variable, evaluated at (x, y), and $\partial_k^n a(x, y^+)$ means the limit of the derivative as the second variable approaches y from values strictly higher than y itself. The derivatives in the expansions above are defined via the limiting procedure because in the kinked case the derivatives do not exist at zero trait difference. Moreover, even if the kernel is smooth, it might only be differentiable once and so its second derivative might only exist to the right and left of the maximum, not at the maximum itself. This procedure is justified since we assumed $x_2 > x_1$, therefore the competition coefficients $a(x_1, x_2)$ and $a(x_2, x_1)$ only need to be considered to the left and right of the kernel's maximum, respectively. Also, notice that the quantities r_0, a_x and a_y are positive due to the positivity of $r_0(x)$ and $a(x, y)$, and the positivity of k_x, k_y, d_x, and d_y is evident from the fact that the kernel is a decreasing function of $|x - y|$.

The dynamical equations may now be written as

$$\frac{\mathrm{d}n(x_1)}{\mathrm{d}t} = n(x_1) \left(r_0 - a_x n(x_1) - (a_x - k_x \Delta x - \frac{d_x}{2} \Delta x^2) n(x_2) \right), \tag{C.7}$$

$$\frac{\mathrm{d}n(x_2)}{\mathrm{d}t} = n(x_2) \left(r_0 + c \Delta x - a_y n(x_2) - (a_y - k_y \Delta x - \frac{d_y}{2} \Delta x^2) n(x_1) \right) \tag{C.8}$$

in this approximation.

The well-known inequalities expressing the necessary and sufficient conditions of stable

coexistence under two-species Lotka–Volterra competition read

$$\frac{a_{12}}{a_{22}} < \frac{r_{01}}{r_{02}} < \frac{a_{11}}{a_{21}} \tag{C.9}$$

(e.g., **?**). In our notation, $a_{12} = a(x_1, x_2)$, $a_{21} = a(x_2, x_1)$, $a_{11} = a_x$, $a_{22} = a_y$, $r_{01} = r_0$, and $r_{02} = r_0 + c\Delta x$. Applying the criterion to these parameters,

$$\frac{a_x - k_x \Delta x - (d_x/2)\Delta x^2}{a_y} < \frac{r_0}{r_0 + c\Delta x} < \frac{a_x}{a_y - k_y \Delta x - (d_y/2)\Delta x^2} \tag{C.10}$$

must be true for coexistence to happen. Let us take the inverse of these conditions:

$$\frac{a_y}{a_x - k_x \Delta x - (d_x/2)\Delta x^2} > 1 + \frac{c}{r_0}\Delta x > \frac{a_y}{a_x} - \frac{k_y}{a_x}\Delta x - \frac{d_y}{2a_x}\Delta x^2. \tag{C.11}$$

At this point, we will consider the smooth and the kinked case separately. We start with the smooth case. If the kernel is smooth, it is differentiable at its maximum and the value of the derivative is zero — therefore $k_x = k_y = 0$ and the quadratic terms are the first nontrivial orders of expansion for the kernel. Then the above condition reduces to

$$\frac{a_y}{a_x - (d_x/2)\Delta x^2} > 1 + \frac{c}{r_0}\Delta x > \frac{a_y}{a_x} - \frac{d_y}{2a_x}\Delta x^2. \tag{C.12}$$

Multiplying by $a_x - (d_x/2)\Delta x^2$ and neglecting terms that are higher order than quadratic, we get

$$a_y > a_x + \frac{ca_x}{r_0}\Delta x - \frac{d_x}{2}\Delta x^2 > a_y - \left(\frac{d_x a_y}{2a_x} + \frac{d_y}{2}\right)\Delta x^2. \tag{C.13}$$

We subtract a_y and use $a_y = a_x + w\Delta x$ to obtain

$$0 > \left(\frac{ca_x}{r_0} - w\right)\Delta x - \frac{d_x}{2}\Delta x^2 > -\left(\frac{d_x a_y}{2a_x} + \frac{d_y}{2}\right)\Delta x^2, \tag{C.14}$$

or, after adding $(d_x/2)\Delta x^2$ and dividing by Δx,

$$\frac{d_x}{2}\Delta x > \frac{ca_x}{r_0} - w > \left(\frac{d_x}{2} - \frac{d_x a_y}{2a_x} - \frac{d_y}{2}\right)\Delta x. \tag{C.15}$$

If $ca_x/r_0 - w$ is positive, there will exist a Δx so small that the first inequality cannot be satisfied. The same is true for the second inequality when $ca_x/r_0 - w$ is negative. This

puts a limit to the similarity of the two species: Δx must be large enough to satisfy both inequalities. Formally, the limit to the similarity of the species disappears when $ca_x/r_0 - w$ is zero, a nongeneric situation.

Having established the limits to the similarity of two competing species under smooth competition kernels, let us turn our attention to kinked ones. In this case the first-order expansion coefficients k_x and k_y are nonzero, rendering the second order negligible in comparison. Therefore in Eq. (C.11) we may neglect any terms that are quadratic or higher order. As a result, we get

$$\frac{a_y}{a_x - k_x \Delta x - (d_x/2)\Delta x^2} > 1 + \frac{c}{r_0}\Delta x > \frac{a_y}{a_x} - \frac{k_y}{a_x}\Delta x. \tag{C.16}$$

Multiplying by $a_x - k_x \Delta x - (d_x/2)\Delta x^2$ and neglecting all terms of quadratic or higher order leads to

$$a_y > a_x - k_x \Delta x + \frac{ca_x}{r_0}\Delta x > a_y - \left(\frac{k_x a_y}{a_x} + k_y\right)\Delta x. \tag{C.17}$$

Using $a_y = a_x + w\Delta x$, rearranging, and simplifying yields

$$0 > \frac{ca_x}{r_0} - k_x - w > -\frac{k_x a_y}{a_x} - k_y, \tag{C.18}$$

which is independent of Δx. The conclusion is that two species may be arbitrarily closely packed if the competition kernel is kinked, as long as these inequalities are satisfied.

Appendix D

Competition kernel as an overlap between sensitivities and impacts

Our purpose is to show that the competition kernel is always expressible as an overlap between two different functions called sensitivities and impacts (**?**). This expression does not depend on the assumptions that lead to the utilization overlap picture. The resource utilization overlap model turns out to be a special case of this general formalism where the sensitivity and impact functions are precisely proportional to one another.

As mentioned in the Introduction, species interactions are mediated through a number of regulating factors, i.e., variables that mediate the feedback loops between densities and growth rates. Familiar examples include resources, predators, pathogens, space, etc. We assume that there is a continuum of regulating entities in the system: $R(z)$ measures the quantity of the zth factor with $z \in [z_0, z_m] \subseteq \mathbb{R}$. Within this framework, the most general continuous time, continuous density model will read

$$\frac{\mathrm{d}n(x)}{\mathrm{d}t} = n(x)\, r\left(R(z, n), E\right), \tag{D.1}$$

where $n(x)$ is the density distribution along the trait axis, and E is the collection of all density-independent model parameters (they may depend on the trait values). Around a fixed point equilibrium with equilibrium distribution n^*, the linearization of the growth rates will read

$$\frac{\mathrm{d}n(x)}{\mathrm{d}t} \approx n(x)\Big(\underbrace{r(R(z, n^*), E)}_{0} + \delta r(R(z, n), E)\Big), \tag{D.2}$$

or

$$\frac{\mathrm{d}n(x)}{\mathrm{d}t} \approx n(x)\left(\int_{x_0}^{x_m} \frac{\delta r(x)}{\delta E(y)}\delta E(y)\,\mathrm{d}y + \int_{x_0}^{x_m}\int_{z_0}^{z_m} \frac{\delta r(x)}{\delta R(z)}\frac{\delta R(z)}{\delta n(y)}\delta n(y)\,\mathrm{d}z\,\mathrm{d}y\right), \qquad \text{(D.3)}$$

where we used the chain rule of differentiation (see Section 3.2 for the meaning of the functional derivative); $r(x)$ is shorthand for $r\left(R(x, n(x)), E(x)\right)$. The factor in the second term of the expansion multiplying the perturbed densities $\delta n(y)$ consists of two parts. The first part,

$$S(x, z) = \frac{\delta r(x)}{\delta R(z)}, \qquad \text{(D.4)}$$

is the *sensitivity* of the species with trait x to the zth regulating factor (**?**), since it measures how the growth rate of species x would change if the zth factor was slightly modified. The second part of the product,

$$I(y, z) = \frac{\delta R(z)}{\delta n(y)}, \qquad \text{(D.5)}$$

is the *impact* of species with trait y on the zth regulating factor. It tells us how the factors regulating the populations are themselves affected by a change in species abundances. As before in Section 3.2, the full factor multiplying the perturbed densities $\delta n(y)$ in Eq. (D.3) is the competition kernel, which in our case is the overlap of the sensitivities and impacts:

$$a(x, y) = \int_{z_0}^{z_m} \frac{\delta r(x)}{\delta R(z)}\frac{\delta R(z)}{\delta n(y)}\,\mathrm{d}z = \frac{\delta r(x)}{\delta n(y)} = \int_{z_0}^{z_m} S(x, z)I(y, z)\,\mathrm{d}z. \qquad \text{(D.6)}$$

Note that this formula applies to any ecological scenario near a fixed point, and as such, it is the proper generalization of the resource utilization overlap picture. The resource utilization function is a phenomenological construct that is intuitive and very useful, but not generalizable to arbitrary ecological situations. The sensitivities and impacts on the other hand are always well-defined, and the competition kernel is always obtained as their overlap integral. Indeed, the resource utilization model is simply the special case when the sensitivity and impact functions are strictly proportional to one another.

As an example, let us consider simple, linear resource competition, a continuous extension of **?**'s (**?**) model. The dynamics of the species densities is given by the equations

$$\frac{\mathrm{d}n(x)}{\mathrm{d}t} = n(x)\left(\int_{z_0}^{z_m} b(x, z)R(z)\,\mathrm{d}z - m(x)\right), \qquad \text{(D.7)}$$

where $R(z)$ is the zth resource, $b(x, z)$ is the potential growth the xth population is able to

achieve on a unit of the zth resource, and $m(x)$ is the density-independent mortality rate of species x. As we can see, the total birth rate is accumulated through the contribution of all the resources available to the species. The resources, in turn, have their own dynamics, which assumes logistic saturation in the absence of consumers and linear consumption in their presence:

$$\frac{dR(z)}{dt} = R(z)\left(R_0(z) - R(z) - \int_{x_0}^{x_m} f(y,z)n(y)\,dy\right), \tag{D.8}$$

where $R_0(z)$ is the maximum (saturation) quantity of resource z, and $f(y,z)$ is the rate at which species y depletes resource z. Assuming that the dynamics of the resources is fast compared to that of the densities, it is always in its equilibrium state:

$$R(z) = R_0(z) - \int_{x_0}^{x_m} f(y,z)n(y)\,dy. \tag{D.9}$$

Substituting Eq. (D.9) into Eq. (D.7) we obtain

$$\begin{aligned}
\frac{dn(x)}{dt} &= n(x)\left(\int_{z_0}^{z_m} b(x,z)\left(R_0(z) - \int_{x_0}^{x_m} f(y,z)n(y)\,dy\right)dz - m(x)\right) \\
&= n(x)\left(\underbrace{\int_{z_0}^{z_m} b(x,z)R_0(z)\,dz - m(x)}_{r_0(x)}\right) \\
&\quad - \int_{x_0}^{x_m}\left(\underbrace{\int_{z_0}^{z_m} b(x,z)f(y,z)\,dz}_{a(x,y)}\right)n(y)\,dy.
\end{aligned} \tag{D.10}$$

As we can see, the competition kernel is the overlap of the functions $b(x,y)$ and $-f(y,z)$. This suggests that these functions play the roles of sensitivities and impacts. Indeed, from their definitions we get

$$S(x,z) = \frac{\delta r(x)}{\delta R(z)} = \int_{z_0}^{z_m} b(x,z')\delta(z-z')\,dz' = b(x,z) \tag{D.11}$$

and

$$I(y,z) = \frac{\delta R(z)}{\delta n(y)} = -\int_{x_0}^{x_m} f(y',z)\delta(y-y')\,dy' = -f(y,z). \tag{D.12}$$

The original MacArthur resource utilization model is recovered when $b(x,z) = \alpha f(x,z)$

for some constant α. Since populations tend to influence those resources most that they depend upon the most, this assumption is reasonable — but it is neither ubiquitous nor necessary.

Appendix E

Generalization of the results of Section 3.5

Here we extend the results obtained in Section 3.5 from resource overlap to arbitrary models. The key to doing this is to write the competition kernel as the overlap of sensitivity and impact functions (see Appendix B); note that this is always possible and does not depend upon the specific assumptions of resource overlap models. The competition kernel is thus given by Eq. (D.6). Since we are interested in nondifferentiability at zero trait difference, we set $y = x$:

$$a(x,x) = \int_{z_0}^{z_m} S(x,z)I(x,z)\,dz, \tag{E.1}$$

where $S(x,z)$ and $I(x,z)$ are the sensitivity and impact functions, respectively, assumed to be bounded. We now show that discontinuities in the sensitivities and impacts occurring at corresponding points between the two functions is sufficient to lead to kinked kernels. Let us consider functions that contain a jump for every possible trait value x:

$$S(x,z) = \alpha\Theta(x - z_1(x)) + \eta(x,z), \tag{E.2}$$

$$I(x,z) = \beta\Theta(x - z_2(x)) + \zeta(x,z), \tag{E.3}$$

where α and β are constants, η and ζ are continuous functions, Θ is the Heaviside unit step function and $z_1(x)$, $z_2(x)$ are curves along which the sensitivity and impact functions possess a discontinuity (they depend on x because we allow for the possibility of each species having their discontinuity at different points).

Similarly to the procedure in Section 3.5, our strategy for determining whether $a(x,x)$ is nondifferentiable will be to take the second derivative of the kernel with respect to the

first variable and see whether the result obtained is infinitely large or not. The second derivative reads

$$\partial_1^2 a(x,x) = \int_{z_0}^{z_m} \partial_1^2 S(x,z) I(x,z) \, dz$$
$$= \alpha \int_{z_0}^{z_m} \delta'(x - z_1(x)) I(x,z) \, dz + \dots,$$

(E.4)

where $\partial_1^2 S(x,z)$ is the second partial derivative of S with respect to the first variable, evaluated at (x,z), δ' is the derivative of the Dirac delta function, and the ellipsis denotes all other terms the derivative produces that we have not written out, for the reason that those terms are necessarily finite and so they do not contribute to the nondifferentiability of the kernel. The integration can be performed with the help of the δ' function:

$$\partial_1^2 a(x,x) = -\alpha \partial_2 I(x,z_1(x)) + \dots,$$

(E.5)

which is infinitely large if I is discontinuous along $z_1(x)$. This of course happens when $z_1(x) = z_2(x)$. Therefore, if $S(x,z_1(x))$ and $I(x,z_1(x))$ are both discontinuous along some curve $z_1(x)$, then the resulting competition kernel is kinked.

Appendix F

Constructing the tolerance-fecundity tradeoff model

Here, we show that Muller-Landau's 2010 tolerance-fecundity tradeoff model arises from Eq. 6 of the main text when a simpler picture of seed tolerance is assumed.

First, we must convert the sum over sites to a sum over fecundities. This can be done for the following reason: sweeping over the stress axis is equivalent to sweeping over fecundity values, as each stress level is the threshold associated with a given fecundity. Mathematically, we can always write

$$\int (\text{integrand}) \, d\sigma = \int (\text{integrand}) \, ds = \int (\text{integrand}) \, \frac{ds}{dx} \, dx,$$

In our case, we will write

$$\int_{s_0}^{s_m} \frac{r(s)T(s,f)}{\int_{f_0}^{f_m} y \, p(y)T(s,y) \, dy} \, ds = -\int_{f_0}^{f_m} \frac{h'(x)g(x,f)}{\int_{f_0}^{f_m} y \, p(y)g(x,y) \, dy} \, dx. \tag{F.1}$$

where s_0 and s_m are the thresholds of the least and most tolerant species, respectively. The left-hand side of Eq. F.1 is simply the integral term in Eq. 6 (main text) with the stress threshold explicitly used as the integration variable, and the functions $h(x)$ and $g(x,f)$ are defined to have the properties

$$h'(x)dx = r(s)ds \qquad g(x,f) = T(s(x),f),$$

where $s(x)$ is the stress threshold of the species with fecundity x. Finally, the minus sign arises because of the change in the limits of the integration given the inverse relationship

146

between stress thresholds and fecundity: $\int_{\sigma_0}^{\sigma_m} ds \to \int_{f_m}^{f_0} dx = -\int_{f_0}^{f_m} dx$.

To see what $h(x)$ means, let us write $r(s)$ as $\frac{d\gamma}{ds}$, where $\gamma(s)$ is some function of s. Then, using the chain rule, we see that

$$\frac{dh}{dx} = \frac{d\gamma}{dx},$$

whereby we can say $h(x) = \gamma(s(x))$. The function γ is easily interpretable: $\gamma(s) = \int^s r(\sigma)\,d\sigma$, that is, $\gamma(s)$ is the total proportion of sites in the community whose stress level is equal to or less than s. The function $h(f)$, then, is the proportion of land whose stress level is equal to or less than the threshold of species f. Its functional form is determined by both the environment and the biology of the organisms, as it depends on the distribution of stress in the land and on how a given fecundity translates to its stress threshold. Notice that the tolerance-fecundity tradeoff, which comprises the essence of the model, requires that $h'(f) < 0$ for all values of f.

The function g is just a renaming of function T, and is not hard to interpret: just as $T(s(x), f)$ is the probability that species f can tolerate in a site with stress s (which is the threshold of species x), $g(x, f)$ is the probability that species f can tolerate a site whose stress level is the threshold of species x. Notice that, for fixed f, $g(x, f)$ is a forward sigmoid function of x, while $T(\sigma, f)$ is a backwards sigmoid function of σ (i.e., g starts low and transitions to 1 as x increases, while T shows the opposite behavior with increasing σ). We change notation from $T(\sigma, f)$ to $g(x, f)$ to highlight this shift of focus from sites to species, which flips the function.

Having justified all terms in Eq. F.1, we can now render our model in the form

$$\frac{dp}{dt}(f) = m\left(-f \int_{f_0}^{f_m} \frac{h'(x)\,g(x, f)}{\int_{f_0}^{f_m} y p(y) g(x, y)\,dy}\,dx - 1\right) p(f). \tag{F.2}$$

If we further assume that all species have a similar transition regime around their specific stress threshold, then the transition function becomes a function of a single variable, $g(x-f)$. Although this greatly simplifies things, such an assumption unfortunately has no a priori biological foundation. It is, however, implicitly used by Muller-Landau, and we apply it here to show the compatibility between the models. With that assumption, Eq. F.2 reads

$$\frac{dp}{dt}(f) = m\left(-f \int_{f_0}^{f_m} \frac{h'(x)\,g(x - f)}{\int_{f_0}^{f_m} y p(y) g(x - y)\,dy}\,dx - 1\right) p(f). \tag{F.3}$$

147

In **?**, the transition function is the unit step function $\Theta(x - f)$ (which is one for positive arguments and zero otherwise): $g(x - f) = \Theta(x - f)$. In biological terms, this means that species with fecundity values x equal to or higher than that of the focal species with fecundity f can tolerate sites whose stress level is that of the focal species with probability 1, and species whose fecundity x is less than f cannot tolerate that stress level at all. Thus Eq. F.3 finally becomes

$$\frac{dp}{dt}(f) = m\left(-f \int_{f}^{f_m} \frac{h'(x)}{\int_{f_0}^{x} y p(y)\, dy}\, dx - 1\right) p(f),$$
(F.4)

which is Muller-Landau's model put in the context of a continuum of species.

Appendix G

Tight packing in the tolerance fecundity model

Here, we analytically solve Muller-Landau's model for tight packing solutions. Simulations reported in the main text further demonstrate that tight packing is robust in this case. Additionally, we show that for a very general class of continuous transition functions g, namely analytic ones, tight packing is not a possible solution.

For tight packing to occur in any subset $[f_1, f_2]$ of the range $[f_0, f_m]$, $p(f)$ must satisfy the conditions $\frac{dp}{dt} = 0$, $p(f) > 0$ for all f in $[f_1, f_2]$. Applying these to Eq. F.4 – Muller-Landau's model written for a continuum of species – we obtain

$$\int_f^{f_m} \frac{h'(x)}{I(x)} dx + \frac{1}{f} = 0,$$

where $I(x) = \int_{f_0}^x dy\, yp(y)$. Differentiating twice the equation above with respect to f and performing appropriate substitutions gives

$$fh''(f) + 2h'(f) + p(f) = 0. \tag{G.1}$$

Eq. G.1 is our tight packing equation, and appears in the main text as Eq. 7. If it is satisfied in the subset $[f_1, f_2]$, then we get a continuum of coexisting species in that subset. Given that Eq. G.1 is a second-order differential equation, $h(f)$ must be further specified by two boundary conditions. Such boundary conditions are more easily obtained from the discrete version of the model, Eq. 1 of the main text. One such condition can be obtained by summing over all i's (recall that, in equilibrium, $\frac{dp_i}{dt} = 0$): $\sum_{i=1}^N p_i = h_1$ – which when extended for a continuum of species reads $\int_{f_0}^{f_m} dx\, p(x) = h(f_0)$. The other condition can also

be obtained from Eq. 1 by substituting $i = N$: $f_N^{-1} \sum_{i=1}^{N} f_i p_i = h_N$ – which in the continuum is $f_m^{-1} \int_{f_0}^{f_m} dx\, x\, p(x) = h(f_m)$. Although in principle Eq. G.1 is an equation for $p(f)$, the easiest way to arrive at a solution for it is to posit $p(f)$ and solve for h, subject to the boundary conditions discussed above. If no suitable h can be found, than the posited $p(f)$ is not a possible solution. As it turns out, there are infinitely many functions $h(f)$ for which there is a corresponding tight packing solution $p(f)$. We give one example in Fig. 2 in the main text.

Having argued that the tight packing equation has a wealth of solutions in the discontinuous model, we now turn to show that it has no solution at all in the continuous case, at least if the transition function is analytic.

We can see from Eq. F.3 that for the tight packing conditions to hold, we must have

$$f \int_{f_0}^{f_m} \frac{h'(x)}{J(x)} g(x - f)\, dx = -1 \qquad \forall\, f \in [f_0, f_m], \tag{G.2}$$

where $J(x) = \int_{f_0}^{f_m} dy\, y\, p(y) g(x - y)$. We now make the assumption that g is analytic. We will also require that the function $\xi(x) \equiv \frac{h'(x)}{J(x)}$ be differentiable at all points and decay to zero at $|x| \to \infty$. This is expected to hold for any reasonable set of functions $h(x)$ and $J(x)$ which obey the constraints of the model. The integral above is a function of f, and given our assumptions about g and ξ, it can be shown that the Taylor series of the integral above converges, and therefore the integral is an analytic function. Since f itself is of course an analytic function and the product of analytic functions is analytic, then the entire left-hand side of Eq. G.2 is analytic. Thus Eq. G.2 can only be satisfied if

$$f \int_{f_0}^{f_m} \xi(x)\, g(x - f)\, dx = -1 \qquad \forall\, f \in \mathbb{R}$$

That, however, cannot be, as the equality cannot hold at $f = 0$. Now, since two analytic functions differing at a point can only coincide at a finite number of points, we conclude that Eq. G.2 does not have a solution for any continuous range $[f_0, f_m]$ if g is analytic.

Although this proves that no analytic transition function would allow tight packing (let alone robust tight packing!), it does not prove that the same should hold for *any continuous* transition function. However, if we restrict ourselves to *differentiable* transition functions (a reasonable and fairly general assumption), then we can approximate any such function with arbitrary precision by an analytic function. Then, given that Eq. G.2 cannot be satisfied

if g is analytic, it stands to reason that no differentiable function g will work either. We believe a formal proof can be constructed, and leave it for future work.

We add that while here we can only ensure tight packing is not possible for analytic tolerance functions, we contend elsewhere (**?**) that, for any model with a differentiable competition kernel, tight packing, if present, is necessarily fragile.

Finally, we speculate that this no-tight packing bound could probably be extended to more general conditions. For instance, we were able to show (unpublished results) that the proof developed here holds when mortality, rather than a constant, is an analytic function of species fecundity.

Appendix H

The competition kernel

Here we show that the competition kernel in Muller-Landau's model, $a(f_1, f_2)$, possesses a discontinuous first derivative at $f_1 = f_2$ if and only if the tolerance function is discontinuous.

It follows from the definition of derivatives that the first derivative of a function is discontinuous at a point if and only if the second derivative diverges to infinity at that point. We will use that to show that the competition kernel has a divergent second derivative if and only if the tolerance function is discontinuous. We will use the transition function g defined in Online Appendix A, rather than the tolerance function T. Recall that $g(x, f)$ describes the probability that the species with fecundity f can tolerate stress level at the threshold of the species with fecundity x.

From Eq. F.2, if $g(x, f) = \Theta(x - f)$, the per capita growth rate is

$$R(f) = \frac{1}{p(f)} \frac{dp}{dt} = -m \left(f \int_{f_0}^{f_m} \frac{h'(x)\Theta(x - f)}{\int_{f_0}^{f_m} yp(y)\Theta(x - y)dy} dx - 1 \right)$$

From this we calculate the competition kernel:

$$a(f_1, f_2) = \frac{\delta R(f_1)}{\delta p(f_2)} = m f_1 f_2 \int_{f_0}^{f_m} \frac{h'(x)\Theta(x - f_1)\Theta(x - f_2)}{(I(x))^2} dx,$$

where $I(x) \equiv \int_{f_0}^{f_m} yp(y)\Theta(x - y)dy$, and take the first derivative with respect to f_2 (we could

152

equally well pick f_1 instead):

$$\frac{\partial}{\partial f_2} a(f_1, f_2) = m f_1 \left(\int_{f_0}^{f_m} \frac{h'(x)\Theta(x - f_1)\Theta(x - f_2)}{(I(x))^2} dx - f_2 \int_{f_0}^{f_m} \frac{h'(x)\Theta(x - f_1)\delta(x - f_2)}{(I(x))^2} dx \right).$$

The second derivative is, then,

$$
\begin{aligned}
\frac{\partial^2}{\partial f_2^2} a(f_1, f_2) &= m f_1 \left(-2 \int_{f_0}^{f_m} \frac{h'(x)\Theta(x - f_1)\delta(x - f_2)}{(I(x))^2} dx \right. \\
&\quad \left. - f_2 \frac{\partial}{\partial f_2} \left(\int_{f_0}^{f_m} \frac{h'(x)\Theta(x - f_1)\delta(x - f_2)}{(I(x))^2} dx \right) \right) \\
&= m f_1 \left(-2 \frac{h'(f_2)}{(I(f_2))^2} \Theta(f_2 - f_1) - f_2 \frac{\partial}{\partial f_2} \left(\frac{h'(f_2)}{(I(f_2))^2} \Theta(f_2 - f_1) \right) \right). \quad \text{(H.1)}
\end{aligned}
$$

Since we are only interested in whether or not the second derivative diverges, we can stop here and look at the behavior of the terms. Since $h(f)$ is well behaved and $I(f) \neq 0$ for all $f > f_0$, the first term never diverges, and is therefore of no consequence. The second term, however, contains a derivative of $\Theta(f_2 - f_1)$. The step function is discontinuous at $f_1 = f_2$, and the derivative of a discontinuous function diverges at the point of discontinuity (more specifically, the derivative of a function at a point of discontinuity is proportional to the Dirac delta function, which is uniformly zero at all points except for where its argument vanishes, at which point it is infinite). Precisely at $f_1 = f_2$, then, we have a singularity in Eq. H.1. This proves that the competition kernel has a discontinuous first derivative at the point where two species coincide if the transition function $g(x, f)$ is discontinuous when $x = f$.

In the same vein, the competition kernel when the transition function g is continuous can be shown to be smooth, i.e., it does not contain any point where the first derivative is discontinuous – no 'kinks'. To see that, we can follow the same steps as above and replace the discontinuous unit function with any continuous function g:

$$
\begin{aligned}
a(f_1, f_2) &= m f_1 f_2 \int_{f_0}^{f_m} \frac{h'(x)g(x - f_1)g(x - f_2)}{(J(x))^2} dx, \\
\frac{\partial}{\partial f_2} a(f_1, f_2) &= m f_1 \left(\int_{f_0}^{f_m} \frac{h'(x)g(x - f_1)g(x - f_2)}{(J(x))^2} dx - f_2 \int_{f_0}^{f_m} \frac{h'(x)g(x - f_1)g'(x - f_2)}{(J(x))^2} dx \right)
\end{aligned}
$$

153

where $J(x) \equiv \int_{f_0}^{f_m} y p(y) g(x - y) dy$. Proceeding to the second derivative, we have

$$\frac{\partial^2}{\partial f_2{}^2} a(f_1, f_2) = m f_1 \left(-2 \int_{f_0}^{f_m} \frac{h'(x) g(x - f_1) g'(x - f_2)}{(I(x))^2} dx \right.$$
$$\left. - f_2 \frac{\partial}{\partial f_2} \left(\int_{f_0}^{f_m} \frac{h'(x) g(x - f_1) g'(x - f_2)}{(I(x))^2} dx \right) \right).$$

Now it should be easy to see that none of the terms above diverge at any point, since they contain only smooth, well-behaved functions and their derivatives. This proves that the kernel is smooth.

Appendix I

The hierarchical competition-colonization model

To illustrate the generality of the impact of a kinked competition kernel on model predictions across different models, here we present a similar analysis for a different 'kinked model', the hierarchical competition-colonization model (**?**). We find that this model shares with Muller-Landau's the ability to support robust tight packing due to discontinuity in resource use. A continuous formulation of the model is presented, from which our simulations indicate robust tight packing is absent. Further tests verify that diversity levels and trait dispersion differ sharply between the original and the smooth formulations.

The original model is defined by the equations (**?**)

$$\dot{p}_i = f_i p_i \left(1 - \sum_{j=1}^{i} p_j\right) - p_i \sum_{j=i}^{i-1} f_j p_j - m_i p_i, \tag{I.1}$$

which determines the growth of species i, characterized by fecundity f_i and mortality m_i. The first term represents recruitment, which can occur in any patch happening to be empty or occupied by any species j with $j \geq i$. In fact, species i sees patches occupied by species j, $j > i$ as empty patches. The second term represents the converse: species i is promptly displaced by any species j, $j < i$. The last term is the mortality term. When we extend Eq. I.1 to a continuous pool of species, we obtain

$$\frac{dp}{dt}(f) = \left(f\left(1 - \int_{f_0}^{f} p(x)dx\right) - \int_{f_0}^{f} x\, p(x)dx - m(f) \right) p(f). \tag{I.2}$$

Let us first derive the conditions under which the solution to Eq. I.2 constitutes tight packing. For tight packing equilibrium, characterized by $dp/dt = 0$, $p(f) > 0$ for all f in $[f_0, f_m]$, we then must have

$$f\left(1 - \int_{f_0}^{f} p(x)dx\right) - \int_{f_0}^{f} x\,p(x)dx - m(f) = 0. \tag{I.3}$$

We can turn this integral equation into a differential equation by differentiating twice with respect to f. The resulting equation, as found in **?**, is

$$3p(f) + 2p'(f) + m''(f) = 0. \tag{I.4}$$

If $m(f)$ satisfies Eq. I.4, then our equilibrium will consist of arbitrarily close species coexisting in the community. Since it is a second-order differential equation, it requires two boundary conditions for the solution to be completely specified. Here, such conditions can be found from Eq. I.3:

$$m(f_0) = f_0,$$
$$m'(f_0) = 1 - 2f_0\,p(f_0).$$

The first is obtained by substituting $f = f_0$, and the second by taking the derivative with respect to f on both sides and then evaluating at $f = f_0$. Thus any solutions to Eq. I.4 subject to these initial conditions is a tight packing solution. We find one such solution and show with simulations that it is robust to small perturbations in $m(f)$ (Fig. I.1A).

Fixing the hierarchical competition-colonization model

As before, we take the model's equations to the species continuum, Eq. I.2, and introduce unit step functions to obtain

$$\frac{dp}{dt} = p(f)\left(f - \int_{f_0}^{f_m} dx\,(f + x)\,p(x)\,\Theta(f - x) - m(f)\right).$$

156

The discontinuity is again effectively removed by substituting the step function with a differentiable function:

$$\frac{dp}{dt} = p(f)\left(f - \int_{f_0}^{f_m} dx\,(f + x)\,p(x)\,g(f - x) - m(f)\right).$$

Whereas in the tolerance-fecundity tradeoff the function $g(x)$ introduces a continuous transition between full tolerance and absolute intolerance to stress, here it introduces a gradient of competitive effect across species. While in Kinzig et al.'s original model an individual could be displaced just as easily by any given seed from any outranking species, here a seed from a species higher up the hierarchy poses a stronger threat than another seed from a species closer down. The same goes for species being outranked by the focal species: the further down the hierarchy the other species is, the greater its competitive disadvantage per seed. In short, our reformulation of the model introduces degrees of competitive ability differences, thus replacing the binary {outranks, is outranked}.

Comparing tight packing, species diversity, and trait dispersion

The competition-colonization tredoff, in its discontinuous (i.e. with a kinked competition kernel) form as presented in **?**, supports robust tight packing (Fig. I.1A). As for its continuous, smooth kernel counterpart, though we are unable to produce analytical proof that tight packing is absent, we argue elsewhere (**?**) that models with smooth competition kernels cannot bear robust tight packing. Thus, if tight packing occurs, the smallest fluctuation in model parameters suffices to break the assemblage. Since such fluctuations could be easily brought about by noise and stochasticity in natural systems, a model which forecloses robust tight packing effectively precludes tight packing altogether.

Analogous tests for diversity levels and species distribution along the trait axis as those performed for the tolerance-fecundity tradeoff model in the main text produce similar results (Fig. I.1B to I.1D).

Figure I.1: **A**: tight packing solution for Kinzig et al.'s hierarchical competition-colonization tradeoff model with uniform mortality. As with Muller-Landau's model, the effect of a perturbation – represented here as a change in the death rate of one species – remains local and does not destroy the assemblage. **B** and **C**: comparison of diversity levels between the original and 'smooth' formulations of the competition-colonization tradeoff produce similar results as those performed with the tolerance-fecundity case, namely, the smooth version typically supports much lower diversity levels. **D**: trait dispersion differs markedly across the two formulations of the competition-colonization tradeoff. For perturbations around uniform mortality, the original model shows a consistent and strong tendency towards species clustering due to the formation of patches of continuous coexistence, whereas the smooth version shows no visible trend towards either overdispersion or clustering. As with the model analyzed in the main text, results depend quantitatively on model parameters (in this case, species mortality), but the two formulations consistently lead to quite disparate trait dispersion across parametrizations.
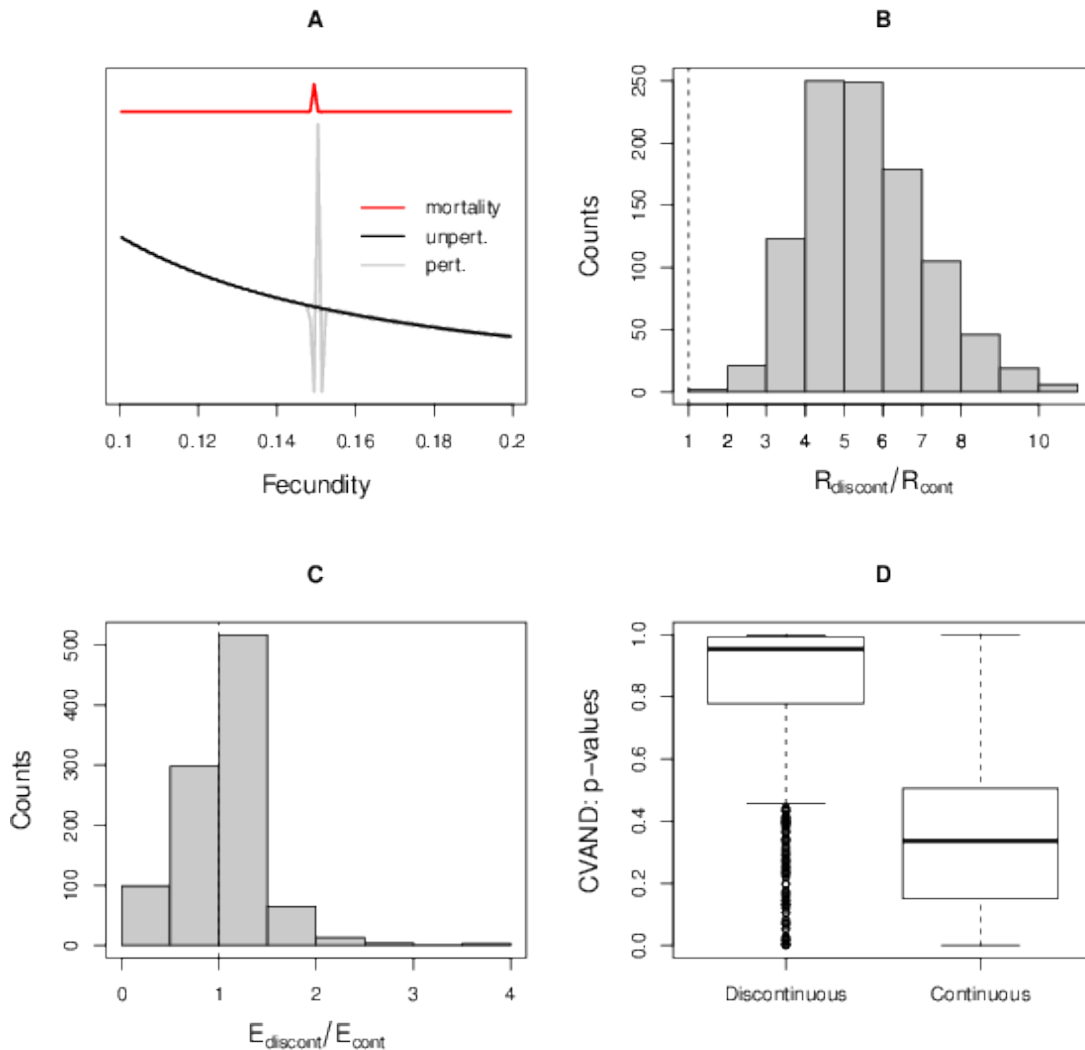
Appendix J

Supplementary information for Chapter 4

Figure J.1: Species mean local persistence times plotted against their abundance at the end of the simulation. No correlation is observed in the H_0 assemblage (left)(correlation index $\rho_{\log(N),\log(T)} = 0.008$), , in contrast with a clear positive correlation in the H_1 assemblage (right) ($\rho_{\log(N),\log(T)} = 0.86$). Abundances and local persistence times are logged for better visualization. Results shown for the baseline scenarios.
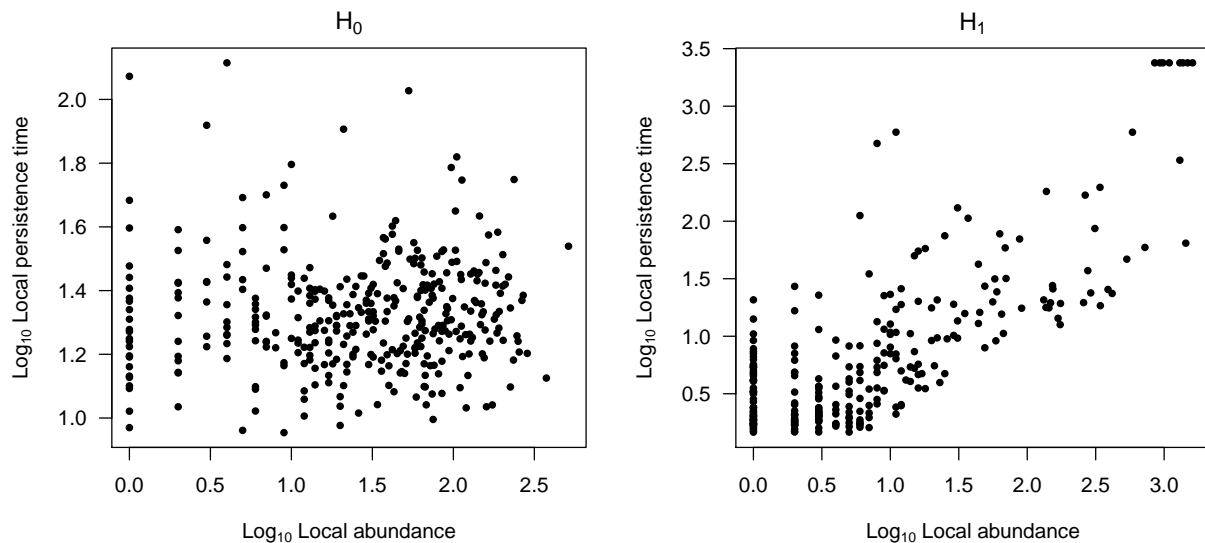
Figure J.2: Comparison between simulation results for a finite niche axis (left) and a circular axis (right). Competitive coefficients in the niche model used in the paper are a function of trait difference d_{ij}. In the finite axis case, $d_{ij} = |x_i - x_j|$, while in the circular case $d_{ij} = \min(|x_i - x_j|, 1 - |x_i - x_j|)$, where $x_i, x_j \in [0, 1]$ are the trait values of species i and j. The complete symmetry of the circular niche axis prevents it from spontaneously sorting into niches and gaps, whereas in the finite case the edge effects break that symmetry. A circular axis is typically used to avoid edge effects, but in our paper we use finite axis for realism. When the axis is finite the niche structure is more evident, and the local persistence times (red curves) confirm that the positions of the niches are set: species located at the niches are residents, while others come and go between competitive exclusion and recolonization events. Particularly in gaps, exclusion is very fast. By contrast, in the circular case no species is a resident, and all species have similar local persistence times. Simulation parameters identical to our baseline cases.
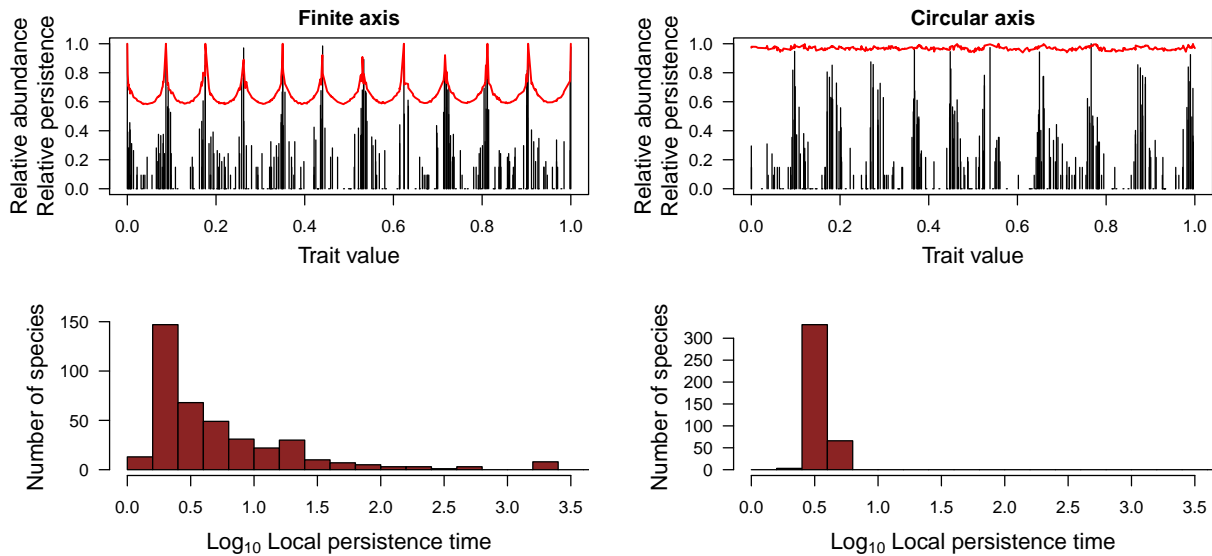
Figure J.3: Comparison between a Lotka-Volterra niche model with fast competitive exclusion ($\rho = 4$, used in the main text), and a model with slow competitive exclusion (the commonly used $\rho = 2$). In a deterministic formulation without immigration, both models have similar final states where only niche-diferentiated species survive, but the sorting is much faster for $\rho = 4$ (compare **A** and **B**, which are snapshots taken after the same number of simulation steps). As a result, stochastic dynamics under immigration differ considerably between these models. Comparing **C** and **D**, we notice that the abundance structure is much less marked in the $\rho = 2$ case. In fact, except for edge effects, there is hardly any abundance structure at all. This is confirmed by the local persistence time curve (in red): the edge species persist throughout the entire tracking period; those species in their immediate vicinity, being in strong competition with them, have particularly short local persistence times; and the remainder of the community have mostly undifferentiated local persistence times. This suggests that niche dynamics in the $\rho = 2$ model are so slow and weak that other forces such as immigration and edge effects dominate. The histograms (**E**) highlight the stark contrast between the local persistence time distribution of the $\rho = 2$ and $\rho = 4$ niche models, as well as $\rho = 2$ similarity to neutral dynamics (compare with Fig. 5.1C). We conclude that slow niche dynamics may be indistinguishable from neutrality in both abundance structure and local persistence times in the context of stochasticity with immigration, even though they would lead to the same final state in the corresponding deterministic closed community model.
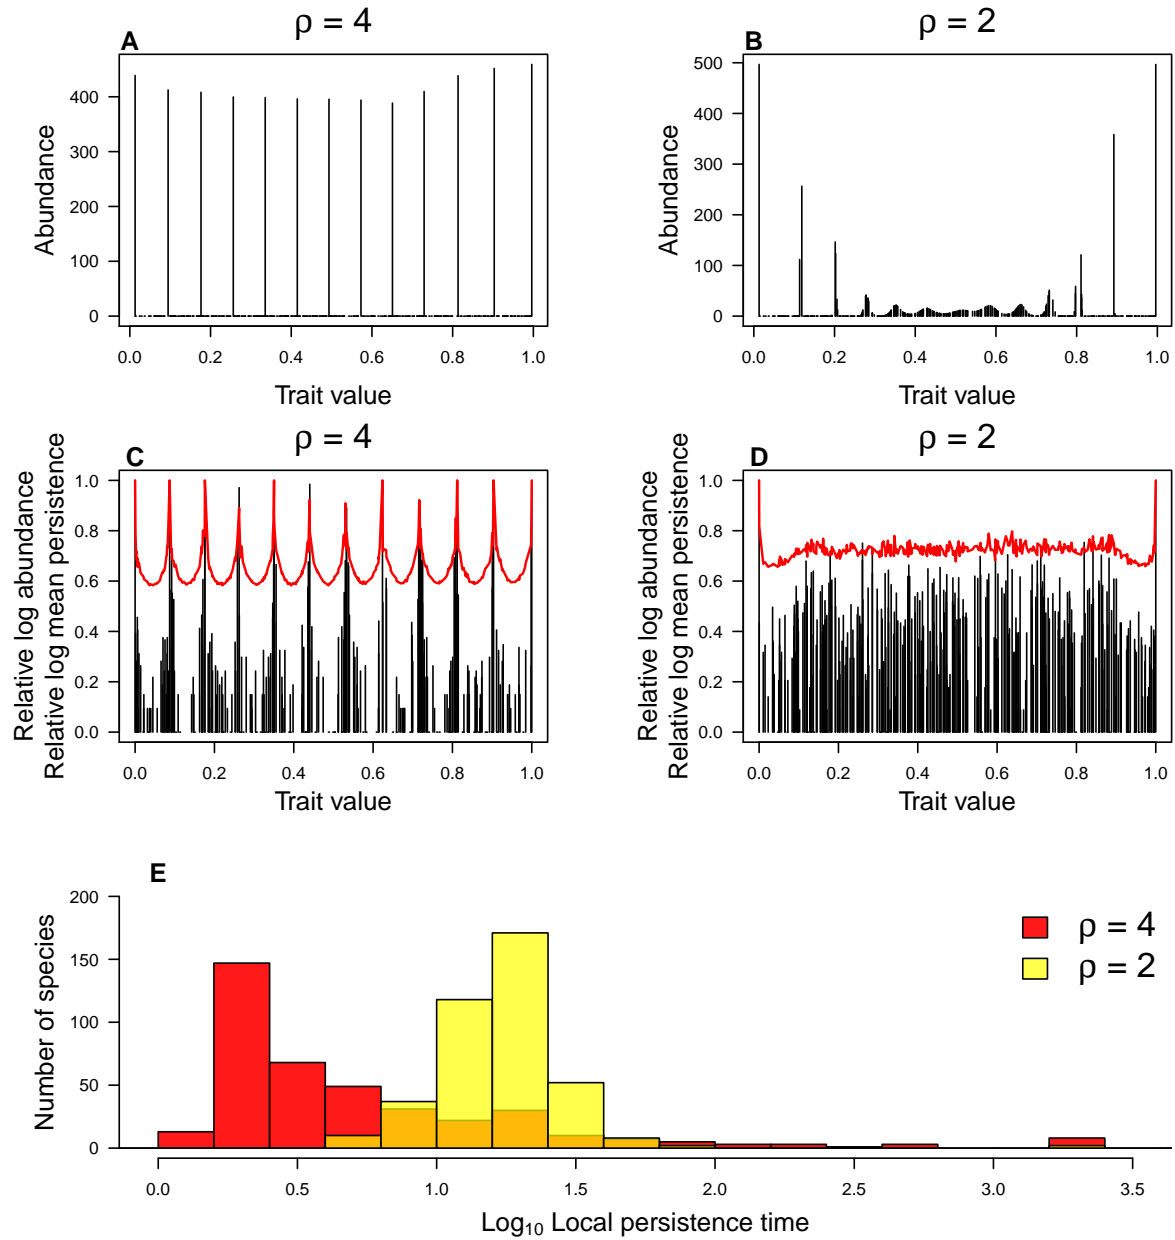
Figure J.4: When regional abundances differ between species, they heavily influence local persistence times. Species that are rare in the pool tend to have the lowest local persistence times, while the most regionally common species tend to be residents locally. The effect is seen both in H_0 and H_1 assemblages, but the correspondence is much tighter in the neutral case (left) than in the partially stabilized assemblage (right). Notice in particular that in the H_1 assemblage some residents are not particularly abundant in the regional pool. Simulation parameters are identical to the baseline cases except for the regional pool.
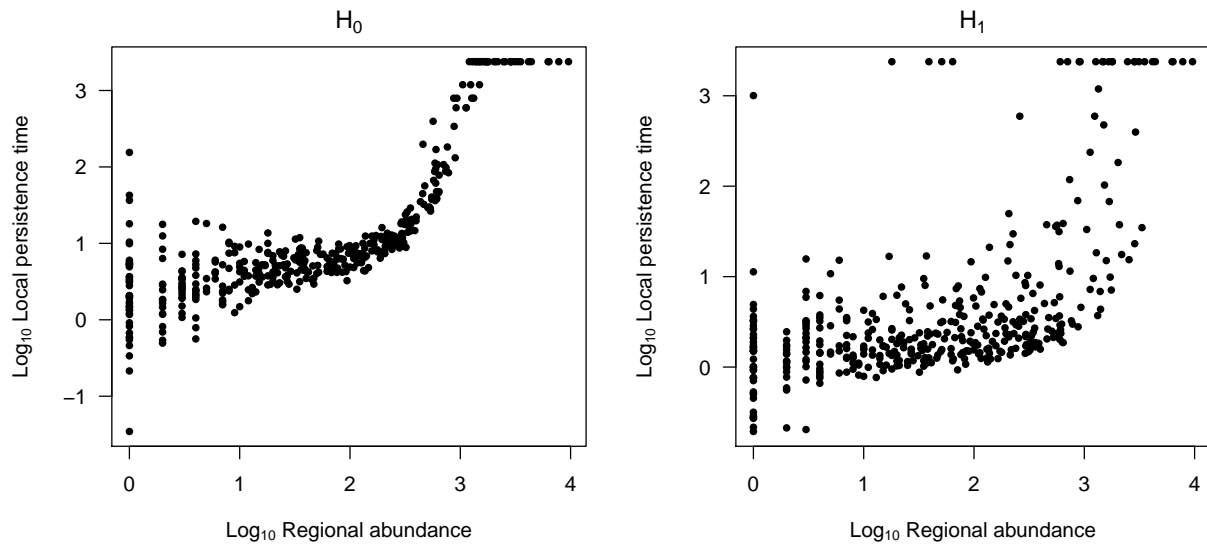
Figure J.5: Comparison between species local persistence times when immigration was low ($m = 0.01$) and high ($m = 0.1$). Points represent the same species across the two immigration scenarios, and their position indicates the base-10 logarithm of their local persistence times, in community turnovers. In both the H_0 and H_1 assemblages, points mostly fall above the 1-to-1 line (red), indicating that most species had longer local persistence times in the high immigration case. In particular, notice that many species are promoted to resident status when $m = 0.1$. Having a local persistence time above a certain cutoff in the low immigration case seems to be a better determiner of resident status in the high immigration case in the H_0 assemblage than in the H_1 assemblage, indicating that stabilization may counter the effects of immigration on local persistence times. Simulation parameters are identical to the baseline cases except for the regional pool and immigration rates.
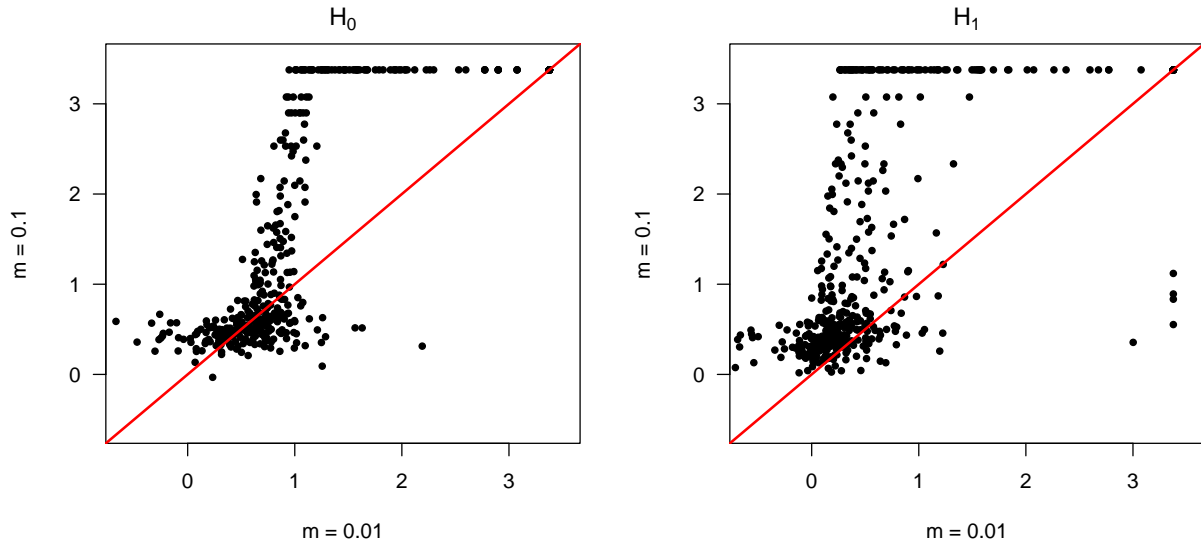
Figure J.6: Distribution of abundances (stems) and local persistence times (red curves) as a function of species trait value in H_0 (top) and H_1 (bottom) assemblages where intrinsic growth rates differ between species (other parameters are set as in our baseline case). Notice the stark contrast between the impact of the intrinsic growth rates r_i on the neutral and stabilized assemblages. While in the neutral assemblage both local persistence times and abundances are dominated by r_i both in the parabolic and random r_i treatments, in the stabilized assemblage the r_i distribution has a much subtler effect. In the H_1, neither r_i forming a parabolic trend along the niche axis (left) nor randomly assigned r_i (right) interfere with the resident status of niche differentiated species. The r_i do have an impact, however. In the parabolic case, the local persistence times increase towards the edges of the axis, where the r_i are actually smaller. This is possibly due to stronger competitive interactions among higher-abundance species near the center of the axis. Abundances and local persistence times logged to facilitate visualization.
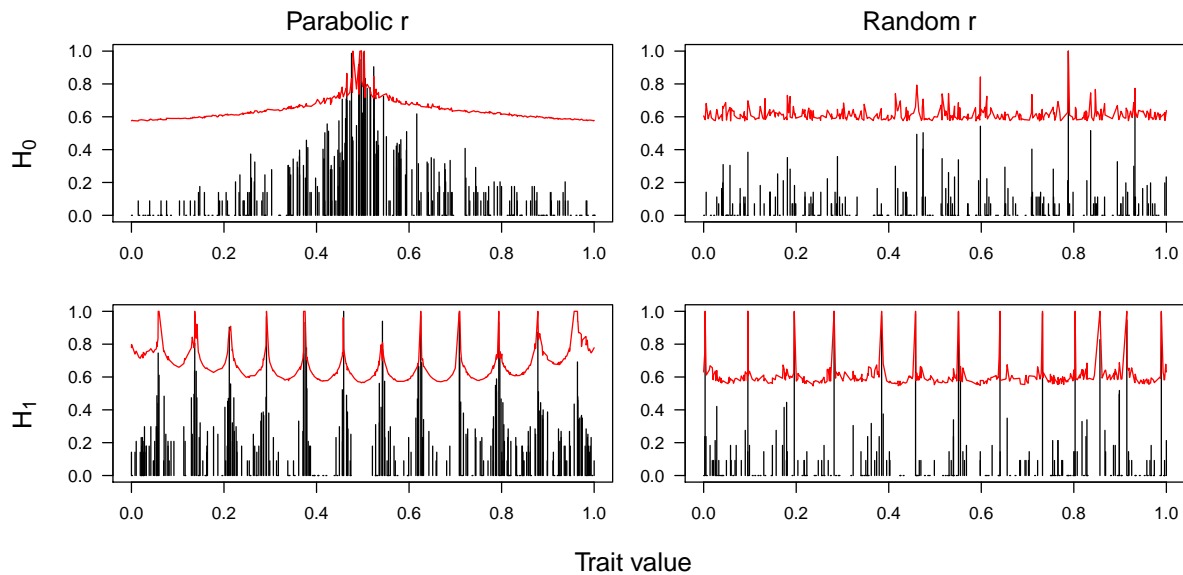

165

Figure J.7: Comparison between assemblages with different niches-to-species ratios. As the ratio increases, more species are able to acquire high abundance and become residents (stem plots and red dots). As a result, the community average local persistence time increases from lower than neutral to higher than neutral (vertical red and blue lines). This change in the mean is driven almost entirely by the increase in residents, as they have a local persistence time that is orders of magnitude longer than most of the non-residents. The local persistence time distribution among the latter remains similar or shorter than neutral (red bars correspond to the niche community, blue bars to the neutral assemblage). See caption of Figure 3 in the main text for simulation parameters.
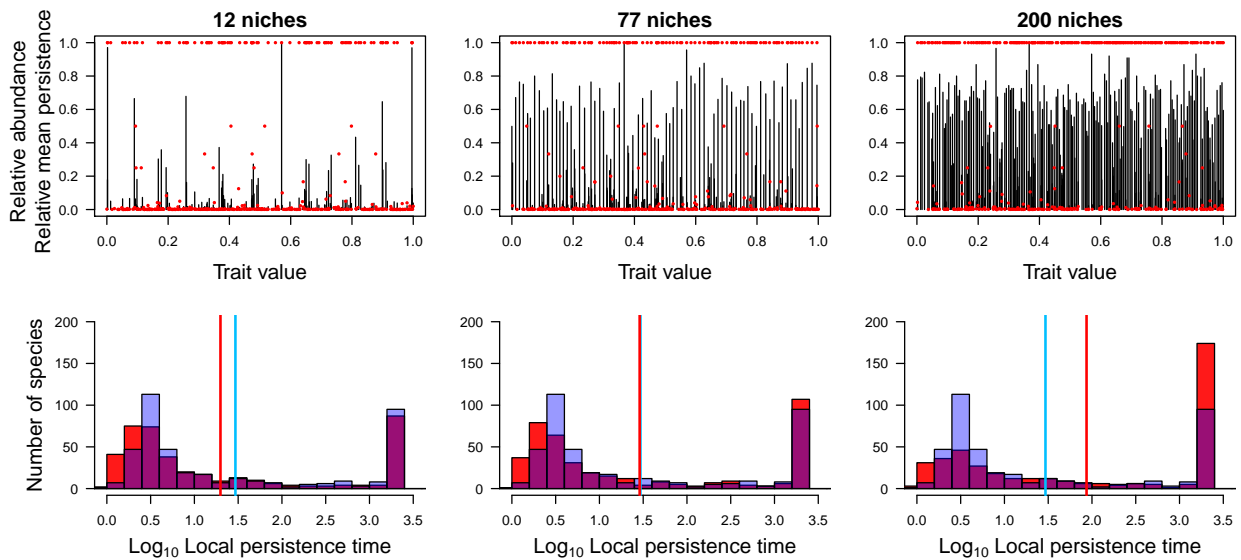
Appendix K

Detailed description of metrics

Clustering metric: the gap statistic

The gap statistic method estimates the number of clusters that best fits the data relative to reference data sets (**?**). For each candidate number of clusters k within a provided range $[k_{\min}, k_{\max}]$, it calculates the goodness of fit of k clusters to the data, as well as the goodness of fit of k clusters to the reference sets. The estimated number of clusters $\hat{k}$ is the k that maximizes the difference in goodness of fit—the "gap"—between the data and the average of the reference sets. The method is quite general, and both the reference sets and the measure of goodness of fit must be specified. Here, following **?**, we use the k-means clustering procedure to provide the measure of goodness of fit.

k-means clustering is a well established method for assigning each data point to one of a previously specified number k of clusters (**?**). It does so by finding the k points in data space that minimize within-cluster distances W_k, i.e. the average distance between members of the same cluster.[1] In our case, the data points are individuals and the data space is the one-dimensional trait axis. The number of clusters, k, must be provided. We apply the gap statistic to estimate k, using $-\log(W_k)$ as the goodness of fit, and reference sets formed by randomizing observed abundances across the species while keeping their trait values.

---

[1]Strictly speaking, W_k, known as the dispersion, is the average sum of squares of within-cluster distances. In more formal terms, the k-means clustering procedure partitions data space of any dimensionality into the set of Voronoi cells that minimizes W_k.

Notice that it does not suffice to compare $-\log(W_k)$ directly across different k's because this number necessarily increases with k, as the average within-cluster distance always decreases for increasing numbers of clusters. This is the reason for calculating the gap statistic between the data and the reference sets: comparing against the references is a way to find which increase in k produces the biggest increase in $-\log(W_k)$, above and beyond what is expected from the increase in k. The reason we use $-\log(W_k)$ rather than simply W_k is that the expected decrements in W_k with increasing k are multiplicative rather than additive. See **?** for more details.

Our metric is then defined as

$$Gap = \begin{cases} -\log(W_{\hat{k}}) + \overline{\log(W_{\hat{k}}^{\text{ref}})}, & \text{if } \hat{k} > 1 \\ 0, & \text{if } \hat{k} = 1 \end{cases} \tag{K.1}$$

where $\hat{k}$ is the estimated number of clusters, $-\log(W_{\hat{k}})$ is the goodness of fit to the data when $k = \hat{k}$, and $-\overline{\log(W_{\hat{k}}^{\text{ref}})}$ is the average goodness of fit to the reference distributions when $k = \hat{k}$. The distinction between $\hat{k} > 1$ and $\hat{k} = 1$ is due to our interest in finding evidence for more than one cluster, because under the interpretation that each cluster represents a niche (see main text), an outcome of a single cluster does not qualify as coexistence based on niche differentiation, and is thus treated as a negative result.

In order to quantify the degree to which a given community is clustered relative to what would be expected from neutral dynamics, we then compare the Gap observed for that community against a set of neutral communities, and calculate the standard score,

$$z = \frac{Gap - E[Gap_{\text{neutral}}]}{SD[Gap_{\text{neutral}}]}.$$

To further verify that identified clusters are meaningful units related to the competitive dynamics, we compare the lifetimes of species identified as the medoids of their clusters to expectations given their regional abundance, and with lifetimes of other species in the community. A medoid is the species whose average dissimilarity to all others within its cluster is minimal. We find medoids using the "cluster" package in R (**?**). Lifetimes are measured for each species by taking the average time between introduction and extirpation events in our simulations.

Other metrics

In addition to our new metric, we also test our niche scenarios for pattern with three other metrics that have been used in the literature.

Rao's quadratic index (**?**): the expected trait distance between two individuals randomly sampled with replacement from the community. Defined as $Q = \sum d_{ij} p_i p_j$, where p_i and p_j are the relative abundances of species i and j, and d_{ij} is the trait distance (absolute value of trait difference) between them, and the sum is over all species pairs.

Functional dispersion, introduced in (**?**), is the abundance-weighted mean distance between species and the community centroid—the point on the trait axis representing the trait value of the average individual. Defined as FDis $= \sum z_i p_i$, where z_i is the distance between species i and the centroid.

Coefficient of variation between adjacent species: tests for even spacing (**?**). Given a set of species sorted by trait value, it is defined as CV $= \sqrt{\mathrm{Var}[d_i]}/E[d_i]$, where d_i is the trait distance between species i and species $i+1$. Often the trait range is used in the denominator instead of the mean (**??**).

A niche-differentiated community is expected to have high Rao, high FDis, and low CV. The CV has also been used to look for clustering, where it is expected to be high (**?**).

Appendix L

Detailed description of niche models and simulation design

Simulation design

Communities start as a random draw from the pool, and in each subsequent step one individual is randomly chosen to die and is replaced by a new recruit. A proportion of recruits fulfilling vacancies left by deaths are immigrants from the pool, and the remainder are local birth events. The particular model used dictates the probabilities that each species is selected for each death and each recruitment event. We simulate dynamics for 5.0×10^7 steps to allow for a stationary state to be reached (i.e. to allow for species abundance distributions averaged across a set of runs to settle to relatively stationary values). We then run for another 5.0×10^7 steps while keeping track of all introduction and extirpation events, for a total of 1.00×10^8 steps, corresponding to roughly 5,000 community turnovers. At the end of a run, trait pattern is calculated using the trait values and abundances of extant species, and each species' average lifetime is computed by calculating the average number of steps that species persisted between being introduced through immigration and being extirpated due to competitive exclusion or drift. Unless otherwise noted for specific niche scenarios, our parameter choices were as follows: the local community has $J = 21,000$ individuals, approximately the number of trees larger than 10 cm diameter at breast height in the 50-hectare plot on Barro Colorado Island (**?**). The regional pool, generated with R package untb (**?**), is a neutral metacommunity with diversity parameter $\theta = 50$ matching fits to BCI's species abundance distribution (**?**). The size of the pool was set to $J_M = 150,000$ individuals, which leads to approximately $S = 410$ species. We checked that bigger sizes

do not change neutral predictions for the SAD. Species trait values are randomly drawn from a uniform distribution between 0 and 1 (standard uniform distribution). Following empirical estimates for Barro Colorado Island, we set the immigration rate m to 0.08 (**?**).

MacArthur-Levins scenarios

The deterministic immigration-free version of the Lotka-Volterra model is

$$\frac{\mathrm{d}N_i}{\mathrm{d}t} = r_i N_i - \sum_j A_{ij} N_j N_i , \tag{L.1}$$

where the competition coefficient A_{ij} quantifies the impact of species j on species i, and r_i is species i's intrinsic fitness, which determines its growth when not limited by competition. In the set of niche scenarios which we will refer to as MacArthur-Levins scenarios, niche differentiation is achieved by setting $A_{ij} = \exp[-((x_i - x_j)/w)^\rho]$, where x_i and x_j are the trait values of species i and j, w sets the scale for the decline of competition between species with increasingly different trait values, and ρ sets the speed of exclusion dynamics (niche sorting). In contrast, neutral competition is achieved by setting $A_{ij} = 1$ between all species.

In our stochastic implementation of Lotka-Volterra competition, competition affects the probability of death, while births are density independent. The probability that species i is selected for a death event is $\sum_j A_{ij} N_i N_j / \sum_i \sum_j A_{ij} N_i N_j$, where the sums are over all species currently present in the local community. The denominator normalizes the probabilities to 1. The new recruit can either be an immigrant with probability m or a local birth event with probability $1 - m$. The probability that species i is chosen at this stage is $m q_i + (1 - m) r_i N_i / \sum_j r_j N_j$, where q_i is species i's relative abundance in the regional pool.

In our "baseline" MacArthur-Levins scenario, we set $\rho = 4$, which causes fast niche sorting, $w = 0.063$, which leads to approximately 13 niches[1], and $r_i = 1$ for all species. To test the effect of immigration on clustering we included variations of the baseline with m set to 0.005, 0.01, 0.05, 0.08, and 0.15. Regional diversity is also likely to affect clustering pattern if the number of niches available to species is fixed. We tested this hypothesis by including variations of the baseline with different regional diversity while holding constant the number of niches. This was done by setting the regional diversity parameter θ to 5, 10, 20, 30, and 50.

---

[1]i.e. in the determinstic formulation this choice leads to an equilibrium with 13 stably coexisting species.

In a community with a finite trait axis, species near the edge of the axis have an automatic advantage because they face fewer competitors from one side. We test the impact of these edge effects by including a scenario with a circular axis. This is done by defining the trait difference between species i and j as $d_{ij} = \min(|x_i - x_j|, 1 - |x_i - x_j|)$.

We consider abiotic filtering in two of our MacArhur-Levins scenarios, modeled as affecting the intrinsic rate of growth r_i of each species. In the purifying filtering scenario, named in analogy with purifying selection in population genetics, environmental conditions favor species with intermediate trait values over those with extreme values. This is implemented by defining $r_i = x_i(1 - x_i)$, where $x_i \in [0, 1]$ is the trait value of species i. In the random filtering scenario, the species characteristic that is filtered by the environment is unrelated to the trait being examined. This is implemented by drawing r_i from a standard uniform distribution. In all other scenarios, $r_i = 1$ for all species. Note that a scaling constant in r_i does not affect the dynamics and outcome of our simulations. We include one final scenario, called pure filtering, where species are under purifying selection but compete neutrally without niche differences (done by setting $A_{ij} = 1$ for all i, j). This is of course not a niche scenario, but we include it to compare and contrast clustering in this scenario against complete neutrality and the niche scenarios (see Fig. 6.4 in main text).

The choice of $\rho = 4$ for our baseline scenario is a deliberate departure from the common choice of Gaussian competition function (made by setting $\rho = 2$) used by MacArthur and Levins (**?**), because the latter has been shown to allow stable coexistence of species with arbitrarily similar traits and is sensitive to edge effects and other factors (**?**). We hypothesize that the Gaussian prescription will lead to weak pattern due to slow niche sorting, and include a Gaussian scenario to test this hypothesis. A slow-sorting scenario should also be sensitive to edge effects caused by the finite trait axis. We therefore also include a Gaussian scenario with a circular axis.

Finally, a defining characteristic of the MacArthur-Levins model, as mentioned in the main text, is that competition is strictly and completely determined by trait similarity, therefore not allowing for any particularities in interspecific competitive interactions caused by the actual trait values of the species involved or any other factor. While capturing the essence of the idea that species can mitigate competition by differing in strategies, this assumption of complete symmetry along the trait axis is likely to be often violated in nature. We test the importance of that property by including two scenarios where the parameter w is not constant across the matrix A_{ij}. In the first, $A_{ij} = \exp[-((x_i - x_j)/w_i)^4]$, where w_i increases linearly between 0.015 for species 1 and 0.15 for species S. In this case, a

pair of species with large trait values competes more intensely than a pair with equal trait difference but small trait values. A hypothetical example of this situation is when the trait has a metabolic cost that must be compensated by higher resource consumption. In our second scenario with non-constant w, we set $A_{ij} = \exp[-((x_i - x_j)/w_{ij})^4]$, where w_{ij} is drawn randomly from a uniform distribution between 0.015 and 0.15 for each pair of species i and j. This adds noise to the dependence of competitive interactions on trait similarity, which is still present but only on average across pairs of species.

Niche differentiation by preferred resource

In its deterministic form, the Rosenzweig-MacArthur model (**?**) is defined by the dynamic equations

$$
\begin{aligned}
\frac{dX_a}{dt} &= r\left(1 - \frac{X_a}{K}\right)X_a - \sum_i C_{ai}X_aN_i \\
\frac{dN_i}{dt} &= \epsilon \sum_a C_{ai}X_aN_i - \mu N_i \,,
\end{aligned}
\tag{L.2}
$$

where X_a is the abundance of resource a, N_i is the abundance of consumer species i, μ is consumer mortality, identical for all species, r and K are the resources intrinsic growth rate and carrying capacity, also identical for all resources, C_{ai} quantifies the amount of resource a consumed by consumer species i, and ϵ is a conversion (efficiency) factor. In our stochastic implementation we use the same zero-sum Moran process described for our MacArthur-Levins scenarios. We assume density-independent mortality, and resource consumption affects consumer birth rates. As death rate is identical across species, the probability that species i with abundance N_i is chosen for a death event is simply $N_i/\sum_j N_j$. The probability that species i is selected for a recruitment event is $m q_i + (1 - m) \sum_a C_{ai}X_aN_i/\sum_j \sum_a C_{aj}X_aN_j$. Resources follow the same stochastic birth-death process as consumers, except that there is no immigration. Resource a is selected for a death event with probability $\sum_i C_{ai}X_aN_i/\sum_b \sum_i C_{bi}X_bN_i$, and for a birth event with probability $(1 - X_a/K)X_a/\sum_b(1 - X_b/K)X_b$. If a resource population is depleted to zero, the resource is extirpated and cannot be replenished.

Notice that the only parameters that affect the stochastic dynamics are the resource carrying capacity and the consumption matrix. We define $C_{ai} = C_0 \exp[-(d_{ai}/\sigma)^2]$, where d_{ai} is the difference between resource a and the preferred resource of consumer i, σ controls

how fast the consumption declines as d_{ai} increases, and C_0 is the consumption level of the preferred resource. We assume C_0 to be identical across all consumers, and therefore it cancels out of the probabilities defined above. We set $\sigma = 0.03$.

The Rosenzweig-MacArthur model simplifies to the MacArthur-Levins model if the resource dynamic equations are solved for equilibrium and the resource equilibrium abundances substituted in the consumer dynamic equations (**?**). This assumes faster resource dynamics relative to consumer dynamics, and ignores the possibility of resource extinction. Abrams showed that resource depletion or exclusion has important effects on the competitive interactions between consumers (**?**), as well as on coexistence outcomes (**?**). We therefore consider two scenarios of this model, one with $K = 100$, and the other with $K = 400$. In the first, most resources are never severely depleted by consumption, whereas in the second resource depletion is severe and many resources are extirpated. Both scenarios start with 500 resources with trait value drawn from a standard uniform distribution.

Niche differentiation by preferred abiotic environment

The third type of niche model we examined was introduced by **?**, and is based on Tilman's stochastic niche model (**?**) and was also studied by **?**. It assumes that different species are optimally adapted to different abiotic environmental conditions, and their ability to colonize different local environments declines with the difference to their optimal environment (their niche optimum). As with the stochastic models described above, Schwilk and Ackerly assumed a saturated community where every site is occupied, and each death, which occurs randomly in one of the sites, is followed by a lottery between all propagules arriving at the vacated site, some of which can be from the regional pool. Which propagule wins the lottery is probabilistically determined by their suitability to the local environment.

A species' "trait" in this model is its niche optimum, which is drawn from a standard uniform distribution. The suitability of site a to species i is $F_{ai} = \exp[-0.5(\Delta_{ai}/0.05)^2]$, where Δ_{ai} is the difference between the local environment a and the niche optimum of species i. There are 1,000 linearly arranged sites, forming a gradient of fixed abiotic conditions between 0 and 1.

Dispersal limitation is implemented via a Gaussian dispersal kernel, $D_{ab} = \exp[-0.5((a-b)/\sigma)^2]$, which sets the probability that a propagule from an individual located at site

b arrives at the vacated site *a*. We explore two dispersal scenarios of this model: local dispersal ($\sigma = 50$), where the probability of arrival of propagules to a site quickly declines with the distance to the parent, and global dispersal ($\sigma = 10^5$), where the probability of arrival is independent of distance to the parent. This corresponds to a case where the physical distances are small relative to the dispersal ability of the individuals.

The probability that site *a* is occupied by species *i* in each time step is then $mp_i + (1 - m) \sum_{b \neq a} \phi_{i,a,b} / \sum_{j=1}^{S} \sum_{c \neq a} \phi_{j,a,c}$, where $\phi_{i,a,b} = F_{ai} D_{ab}$ and p_i is the regional abundance of species *i*.

Niche differentiation by life history strategy

In the competition-colonization tradeoff model (**??**), species trade off fecundity with competitive ability: the number of propagules produced trades off with the ability of propagules to displace individuals of other species. Species are ordered by fecundity value, with species 1 being the least fecund and most competitive, and species *S* being the most fecund and least competitive. This model as originally formulated is described by the dynamic equations (**?**)

$$\frac{dp_i}{dt} = f_i p_i (1 - \sum_{j=1}^{i} p_j) - p_i \sum_{j=1}^{i-1} f_j p_j - \mu p_i , \tag{L.3}$$

where p_i is the relative abundance of species *i* in the local community, f_i is its fecundity (i.e. the number of propagules produced by each individual during time interval dt), and μ is the mortality rate, here assumed identical across all species. After some algebra, Equation (L.3) can be rewritten (see **?**)

$$\frac{dp_i}{dt} = (f_i - \mu)p_i - \sum_{j=1}^{S} \Theta_{ij}(f_i + f_j)p_j p_i , \tag{L.4}$$

where $\Theta_{ij} = 1$ if $i > j$, 1/2 if $i = j$, and 0 if $i < j$. It has been argued that for realistic purposes the strict hierarchy must be replaced with a probabilistic approach where higher-ranking species are more likely to displace lower-ranking ones (**?**). We follow this approach by replacing the abrupt function Θ_{ij} with the continuous function $G_{ij}(f_i, f_j) = 0.5 (1 - \tanh[\nu(f_j - f_i)])$, which is equal to 0.5 when $f_i = f_j$ and asymptotes to 1 and 0 when $f_i \gg f_j$ and $f_i \ll f_j$, respectively. In the limit of $\nu \to \infty$ one recovers the step function.

Notice that Equation (L.4) has the same format as the Lotka-Volterra model, Equation (L.1). We use that correspondence in our stochastic formulation of this model. Again, we implement a saturated community under a stochastic birth-death process with immigration from a regional pool, and place the density-dependent effects on the probability of being selected for death. The probability that death befalls species i is thus $\sum_j A_{ij}N_iN_j / \sum_i \sum_j A_{ij}N_iN_j$, where $A_{ij} = G_{ij}(f_i + f_j)$ is the competition coefficient, and the probability that a recruit is from species i is $m f_i q_i / \sum_j f_j q_j + (1 - m) r_i N_i / \sum_j r_j N_j$, where $r_i = f_i - \mu$. Notice that immigrants are more likely to be from species with high fecundity. The logarithm of the fecundities were drawn from a uniform distribution between 0 and $\log 500$. We set mortality rate $\mu = 1$, and steepness parameter $\nu = 0.15$ in function G.

The other lifetime tradeoff model we examine is Muller-Landau's tolerance-fecundity tradeoff model (**?**). This model assumes that species trade off fecundity with stress tolerance, and the environment varies in stress levels. The model can be written (**?**)

$$\frac{dN_i}{dt} = \mu \sum_a h_a \frac{T_{ai} f_i N_i}{\sum_j T_{aj} f_j N_j} - \mu N_i \tag{L.5}$$

where f_i is species i's fecundity, T_{ai} is its tolerance to stress level a, h_a is the number of sites with stress level a, and μ is the mortality rate, common to all species. In this model, mortality is density-independent, and competition occurs at the recruitment stage.

Recruitment occurs as follows: for each recently vacated site of stress level a, all species j contribute propagules according to their fecundity f_j and abundance N_j, which are then weighted by their tolerance to stress a. A proportion $T_{ai} f_i N_i$ of those viable propagules belong to species i. The recruitment at all sites with stress a is equal to this proportion multiplied by the number of sites with that stress level, h_a. Total recruitment is then calculated by summing over all stress levels.

In our stochastic implementation of this model, which follows the same zero-sum recipe used in the other models, species are selected to die in proportion to their abundances, and the probability that species i is selected for a recruitment event is $m f_i q_i / \sum_j f_j q_j + (1 - m) I_i / \sum_j I_j$, where $I_i = \sum_a h_a (T_{ai} f_i N_i / \sum_j T_{aj} f_j N_j)$. There are 21,000 sites of 400 different stress levels in the local community. The number of sites of each stress level is uniformly distributed. The stress levels are drawn from a standard uniform distribution, as are the fecundity values (note that multiplicative factors in the fecundity values cancel out of the stochastic dynamics). The tolerance function we used was $T_{ai}(s_a, f_i) = 0.5(1 - \tanh[100(s_a + f_i - 1)])$, where s_a is the stress level with index a. It has a sigmoidal shape, ranges from 0 to

1, and monotonically decreases with stress level and fecundity value[2].

Deterministic immigration-free simulations

For each niche scenario described above[3], we also ran deterministic immigration-free simulations in order to compare and contrast the outcomes with our simulations of stochastic niche assembly under immigration.

In these simulations, the initial community corresponds to the regional pool used in the stochastic simulations. That is, all of the approximately 410 species are initially present, with abundances equal to their regional abundances, totaling 150,000 individuals. Those abundances are then simultaneously updated at identical time intervals following the respective niche model's dynamic equations, without immigration, until stable equilibrium is reached—i.e. until only those species that are able to stably coexist remain and have reached their equilibrium abundances. The time interval is set so that on average each update corresponds to one death and/or one birth at the beginning of the simulation. We observed that 2 to 10 million updates sufficed for equilibrium to be reached, depending on the model. Equilibrium was tested by checking that all abundances were changing by less than one individual across the most recent 100,000 updates.

---

[2]In the original formulation by Muller-Landau, T_{ai} was a step function, but **?** argued that this is an intrinsically unbiological assumption because it assumes that differences in tolerance strategy can remain large between arbitrarily similar species, hence our use of a smooth function instead

[3]Except for the Schwilk-Ackerly model of competition for suitable habitat, which is mathematically similar to the MacArthur-Levins model with Gaussian competition function (the rationale is the same behind the demonstration that the Rosenzweig-MacArthur model, in the limit of fast resource dynamics, reduces to the MacArthur-Levins model due to the fact that the overlap between two Gaussian resource use curves yields a Gaussian competition curve. See **?**).

Appendix M

Supplementary figures for Chapter 5

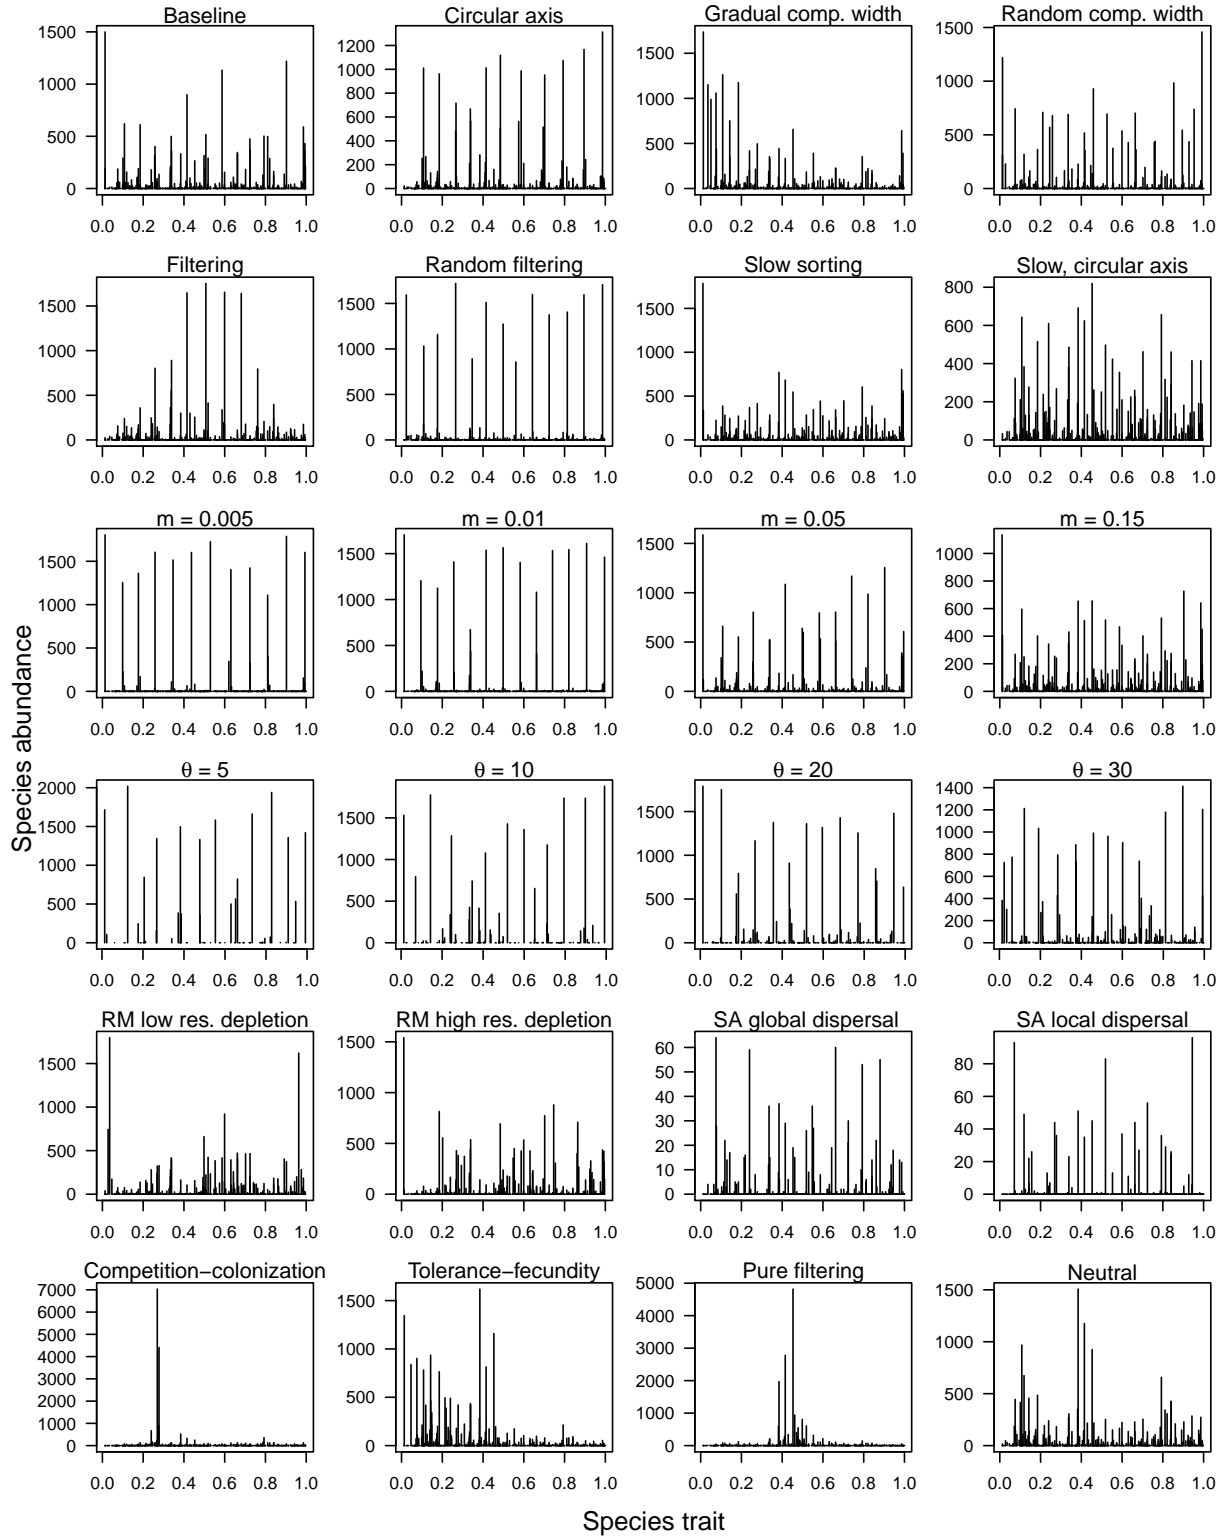Figure M.1: Example simulation outcome of each scenario.

179

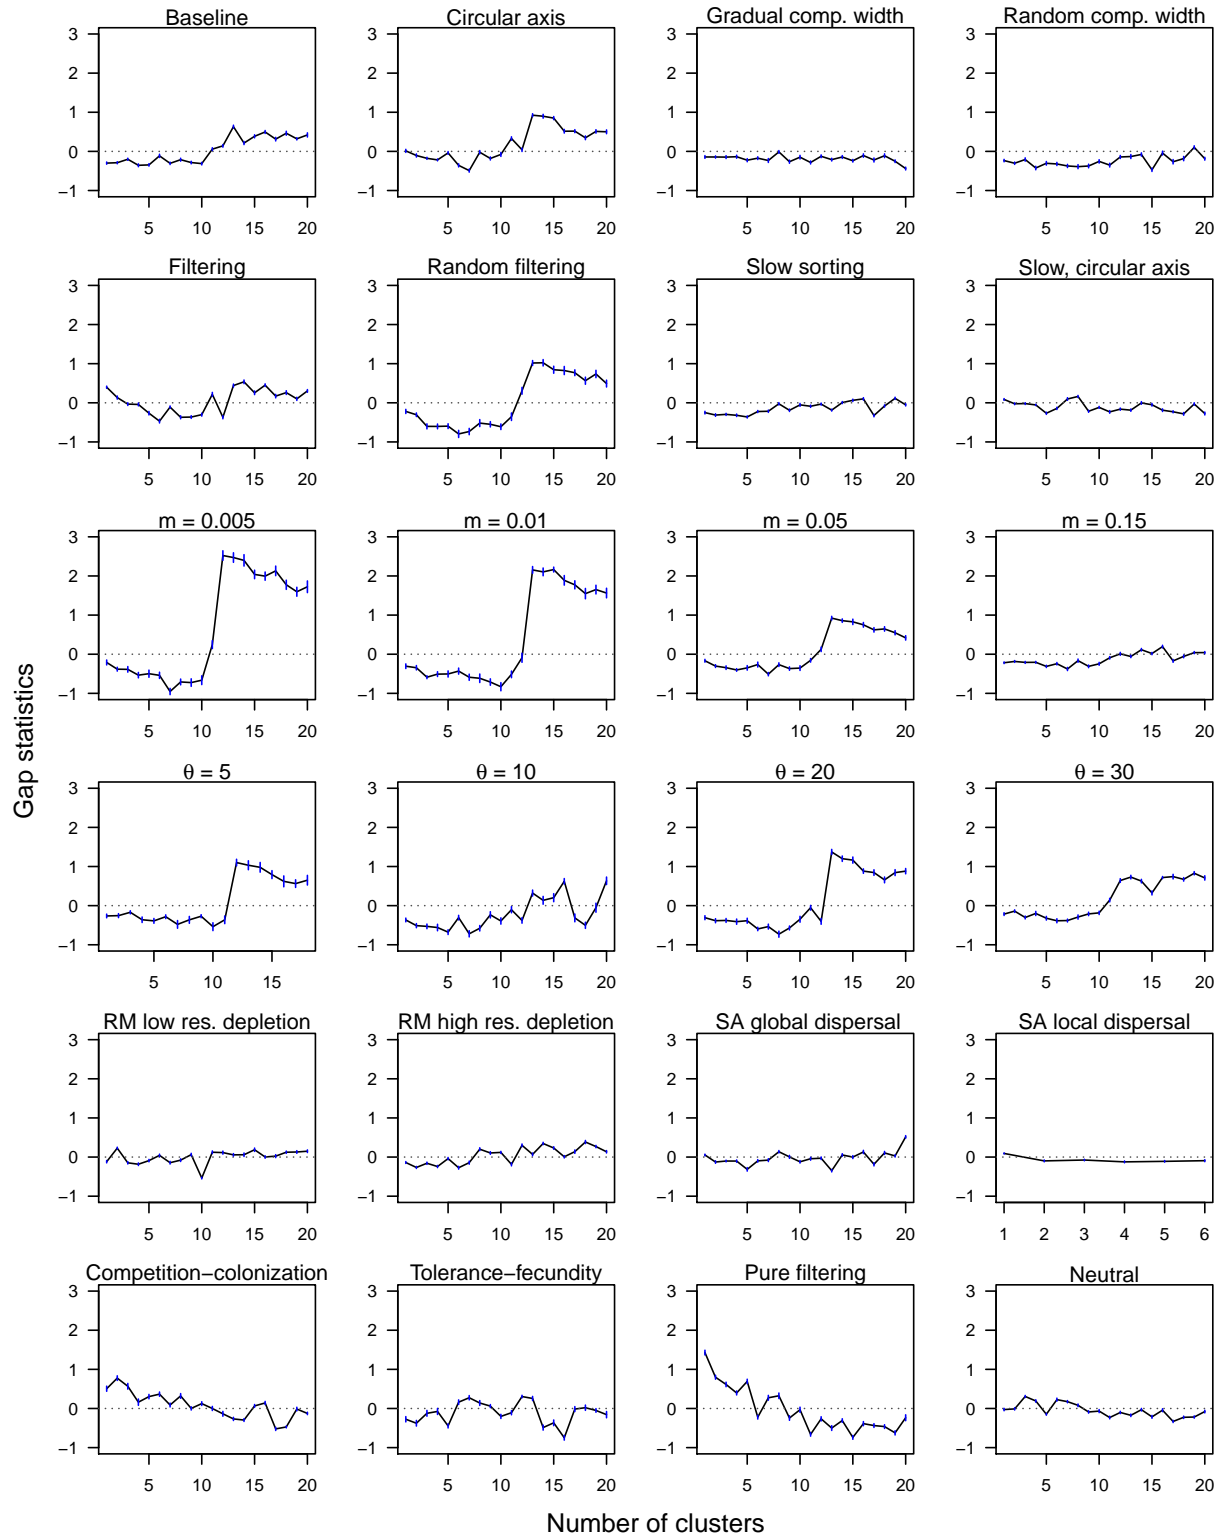Figure M.2: Gap curves corresponding to each scenario shown in Fig. M.1.

Figure M.3: Species lifetimes versus regional abundance for each scenario shown in Fig. M.1. See Fig. 6.5 in main text for details.
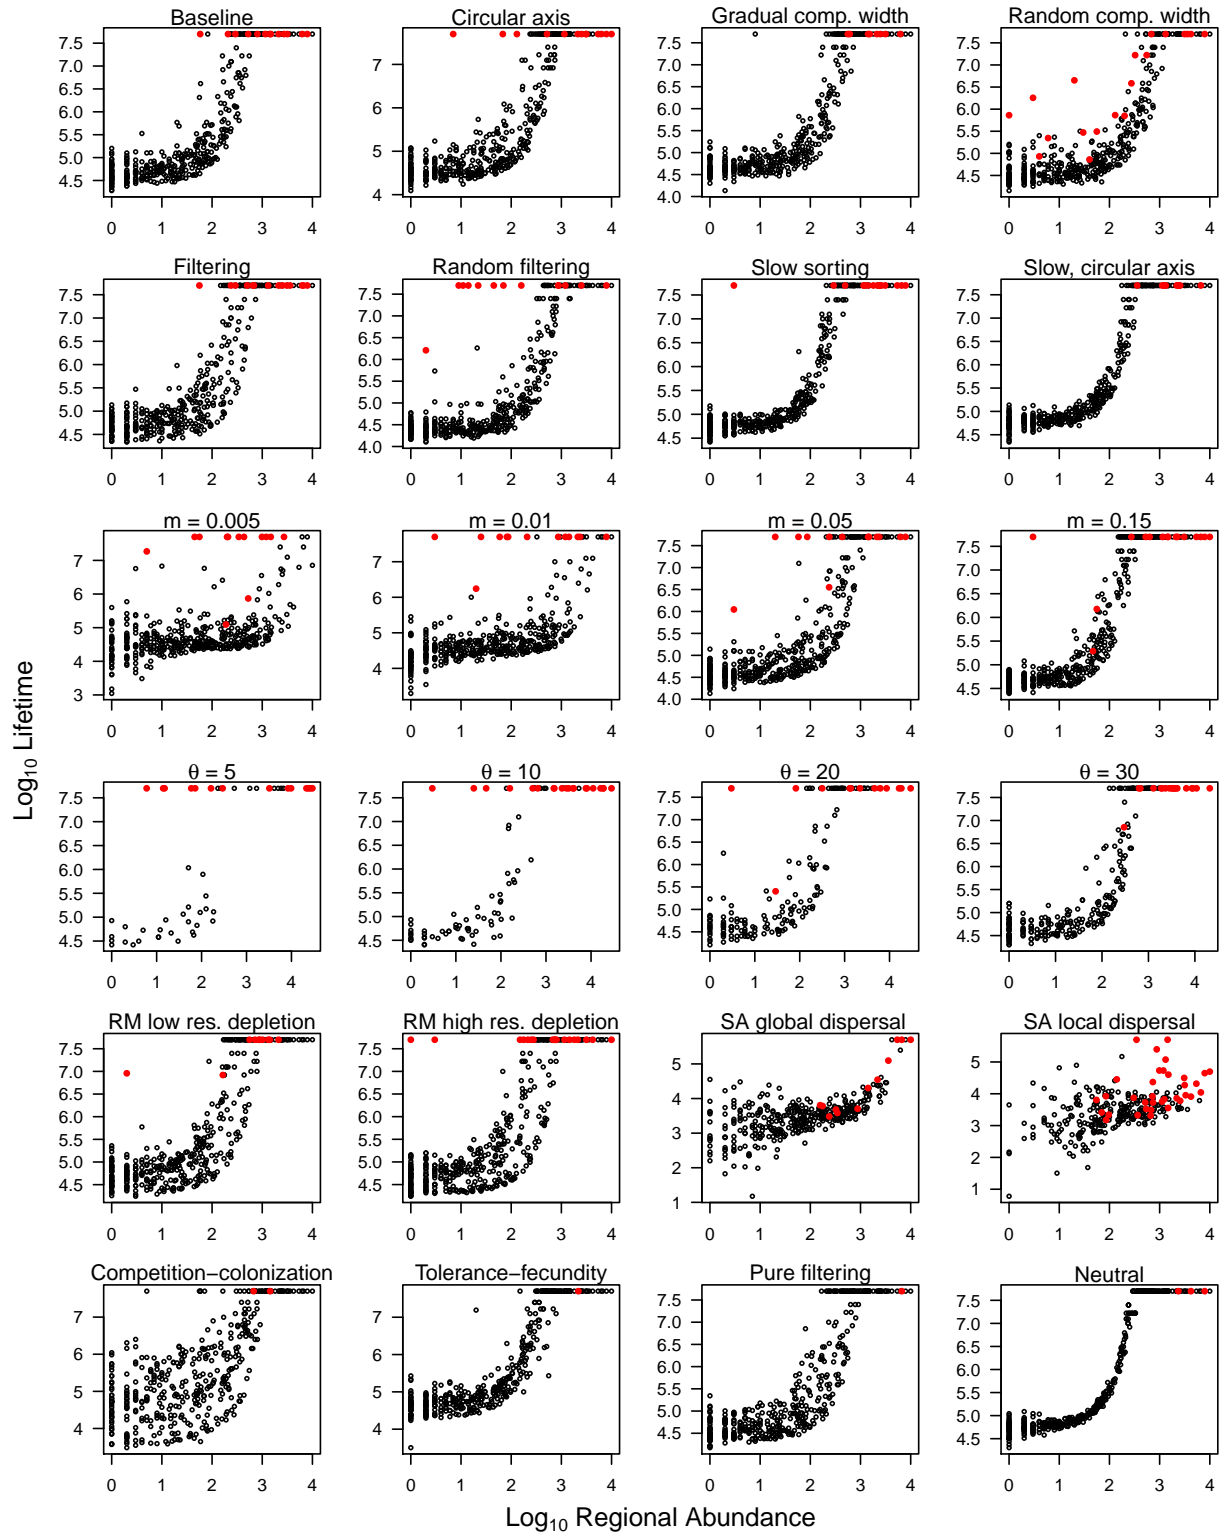
Figure M.4: Trend in coefficient of variation for each of the scenarios shown in Fig. M.1. See Fig. 6.6 in main text for details.
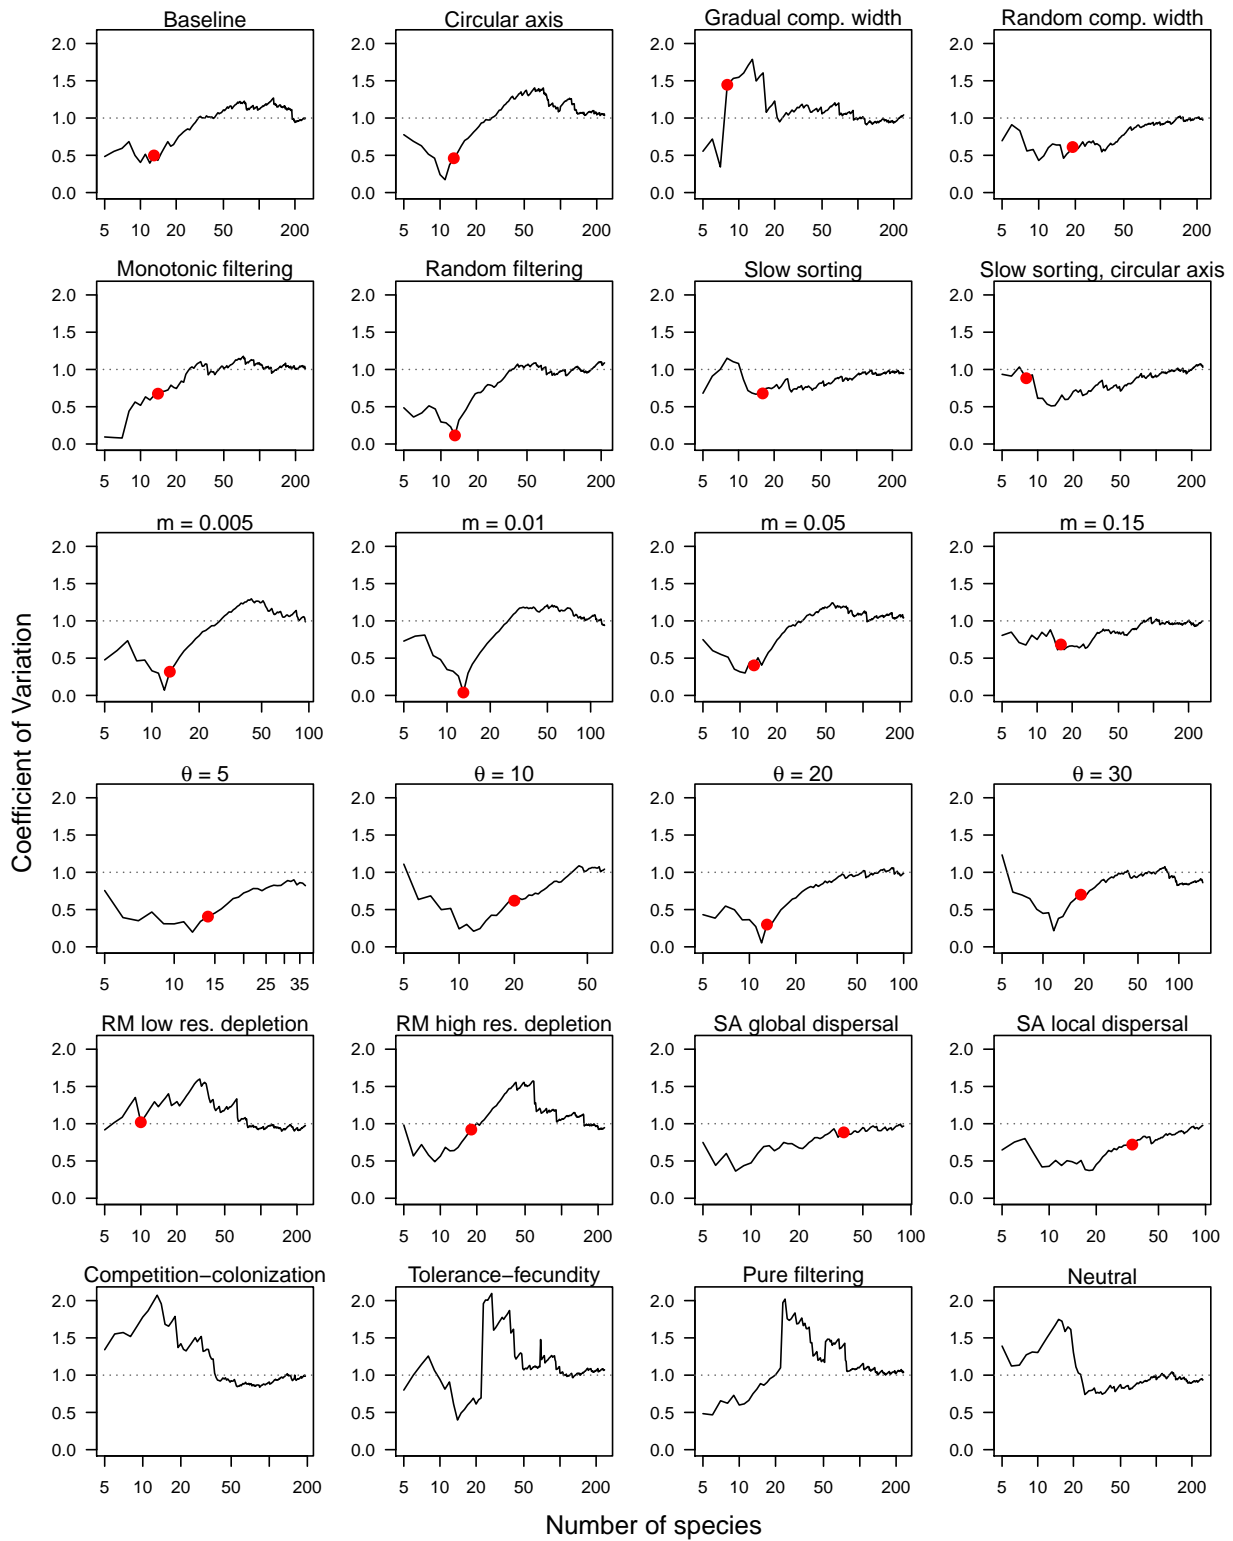
Figure M.5: Deterministic immigration-free outcomes of all scenarios.
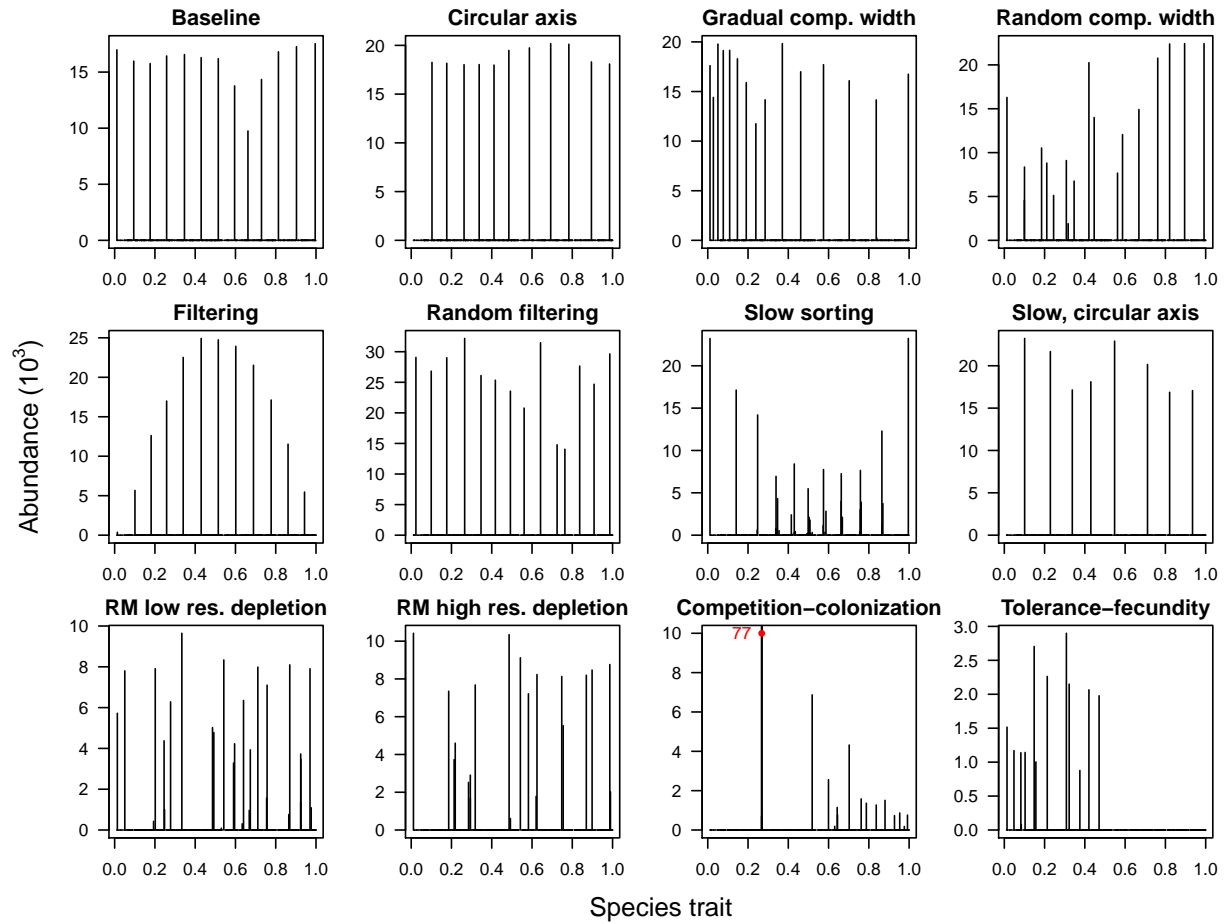
Figure M.6: Example simulation outcomes of the resource-consumer model for low and high resource depletion (top and bottom rows, respectively). Resources are shown on the left, consumers on the right. In the high resource depletion scenario, gaps left by resource extirpation cause analogous gaps among consumer species. These gaps possibly explain why clustering was stronger in this scenario than the low depletion case.

Figure M.7: Average standard score of the gap statistic across ten runs of the finite-axis slow-sorting MacArthur-Levins scenario, for five different immigration rates. As with the fast-sorting scenario (Fig. 6.2 in main text), there is a negative trend in clustering with immigration. Notice that clustering is strong for sufficiently low immigration rates.

Appendix N

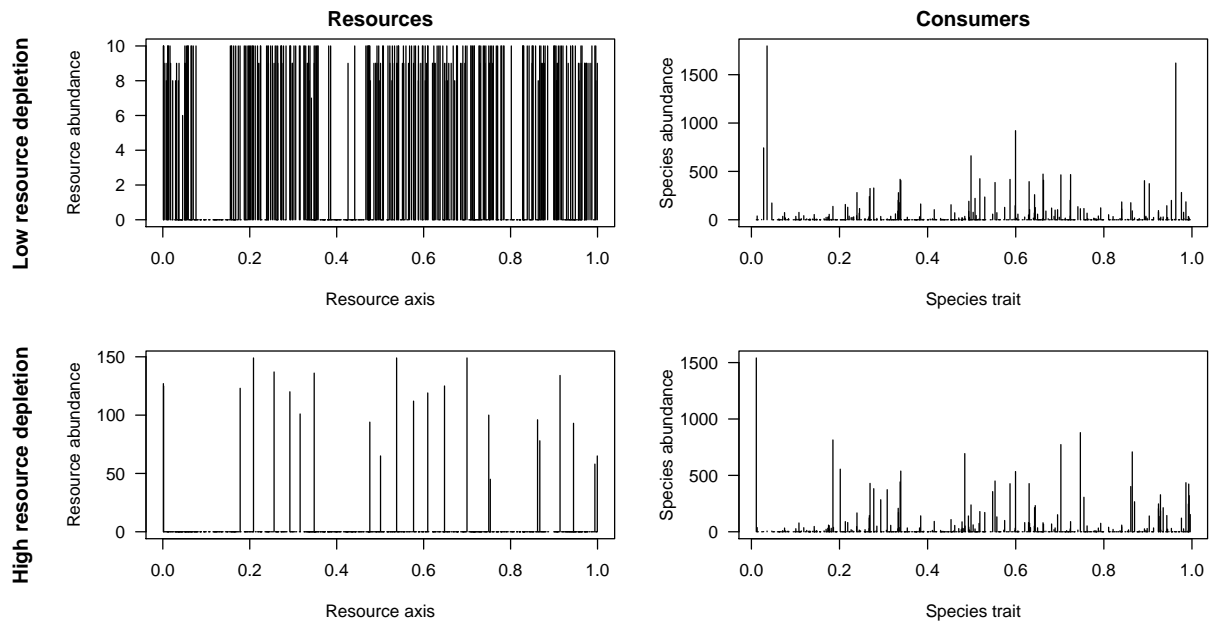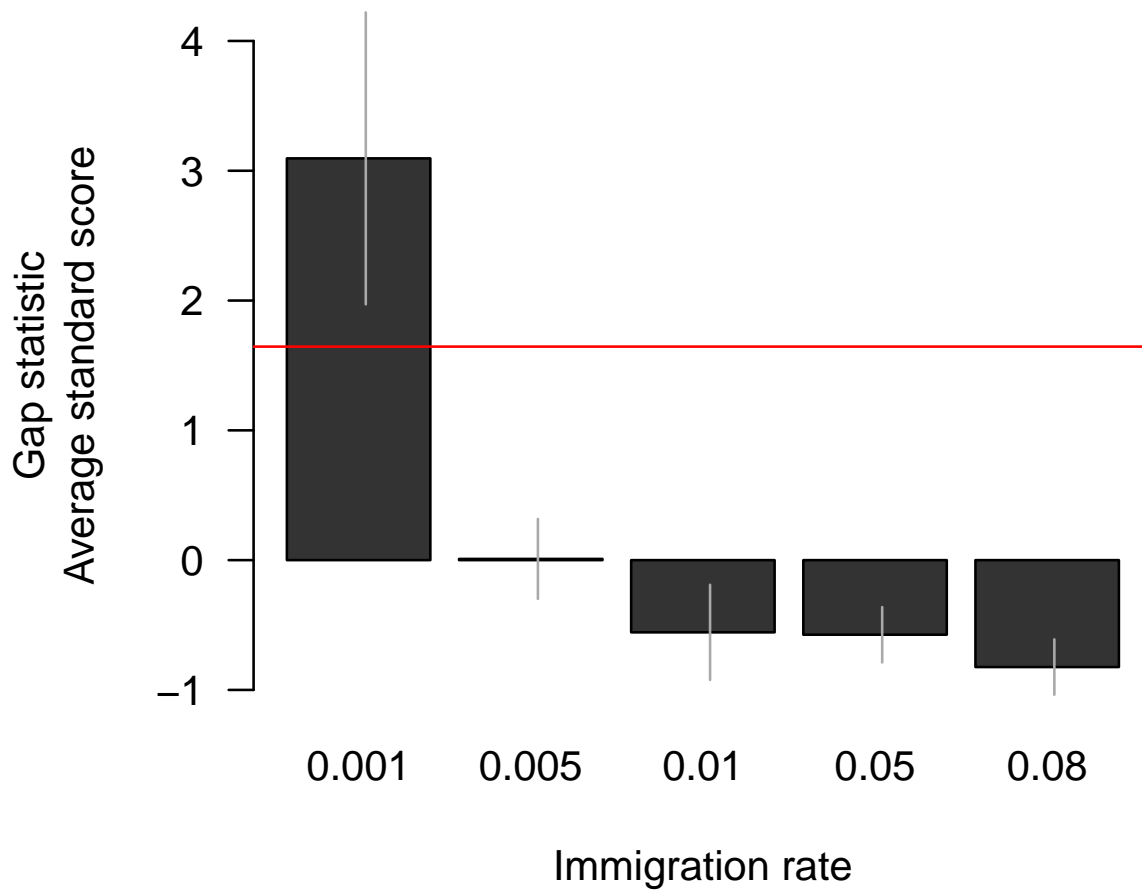# Demonstration that neutral communities can be clustered in genotype space

Using a master equation approach, Houchmandzadeh and Vallade (2003) demonstrated the formation of spatial clusters in a neutral model of birth, death, and migration, where the latter could only occur between neighboring geographic locations. In the long term, individual organisms tend to aggregate spatially due to the local nature of geographic dispersal. Here we employ the same technique to show that clusters also arise in a neutral evolution scenario where haploid individuals die and reproduce at random, and mutations, when they occur, are limited to replacement of a single nucleotide (point mutation).

Consider a genetic segment L bases long. The genotype is defined by the set of bases A, C, G, T, at each site of this segment. There are 4^L possible genotypes. The state of the community is the set of abundances of each genotype. Suppose that individuals die and replicate with a chance of mutation under a purely neutral regime. Further, suppose there can only be point mutations, that is, mutations that replace a single base nucleotide with another nucleotide (no indels and no large-scale mutations are allowed).

Our goal is to show that in the long term there will be a positive correlation between the abundance of a genotype and its neighbors in genotype space, irrespective of the initial composition of the community. Furthermore, this correlation should decline for progressively dissimilar genotypes.

Point mutations link the abundances of "first neighbors" – pairs of genotypes that differ from one another in only one site. In our neutral scenario, birth and death rates of a genotype are otherwise unaffected by the abundances of the other genotypes. We will

use the master equation to derive the covariances between abundances of genotypes. The master equation is a differential equation for the probability P that the system is observed at a given state $\mathbf{n}$. It contains transition rates W describing the probability of change in the abundance of a given genotype. It reads

$$\frac{\partial P(\mathbf{n})}{\partial t} = \sum_{i=1}^{4^L} P(n_i - 1)W^+(n_i - 1) - P(n_i)W^+(n_i)$$

$$+ P(n_i + 1)W^-(n_i + 1) - P(n_i)W^-(n_i) \tag{N.1}$$

where boldface $\mathbf{n}$ stands for the set $n_1, \ldots, n_k, \ldots, n_{4^L}$, and $P(n_i)$ is shorthand for $P(\mathbf{n}) = P(n_1, \ldots, n_i, \ldots, n_{4^L})$. In words, this equation says that the change in probability of finding the community at state $\mathbf{n}$ is the balance between the probability that the community was at a state with one individual too many or one too few and transitions into state $\mathbf{n}$, and the probability that it was at state $\mathbf{n}$ and leaves by gain or loss of an individual.

The model is defined by the transition rates $W^+(n_k)$ and $W^-(n_k)$, which express the rates at which a genotype k with abundance n_k gains and loses one individual, respectively. The interaction-free model is defined by the following transition rates:

$$W^+(n_k) = b(1 - \mu)n_k + b\mu \frac{1}{3L} \sum_{k'} n_{k'} \tag{N.2}$$

$$W^-(n_k) = dn_k \tag{N.3}$$

The second expression is the rate at which genotype k loses one individual, which occurs via death of an individual of that genotype. The first represents the rate at which genotype k gains an individual, which occurs either via the birth of a clone or via the birth of a mutant sired by another genotype. Notice that only the first neighbors of genotype k, written k', can mutate into k.

The birth and death rates of a genotype are unaffected by the community size J, which is free to vary. Consistent with our assumption of neutrality, we assume that the per capita rates b and d are uniform across genotypes. We can further assume that the community size should not consistently increase or decrease through time, which implies $b = d$.

Let u_{kl} be the covariance between genotypes k and l at time t, that is,

$$u_{kl}(t) = \langle n_k(t)n_l(t) \rangle - \langle n_k(t) \rangle \langle n_l(t) \rangle,$$

where the angle brackets denote an average over $\mathbf{n}$ (in effect averaging over a large number of replicates of the community). The correlation is then $\rho_{kl}(t) = u_{kl}(t)/(u_{kk}(t)u_{ll}(t))^{1/2}$.

In Appendix O we show, in close analogy with Houchmandzadeh and Vallade 2003, that u_{kl} satisfies the following differential equation:

$$\frac{du_{kl}}{d\tau} = -2u_{kl} + \frac{1}{3L}\sum_{k'} u_{lk'} + \frac{1}{3L}\sum_{l'} u_{kl'} \qquad \text{for } l \neq k, \tag{N.4}$$

$$\frac{du_{kk}}{d\tau} = -2u_{kk} + \frac{2}{3L}\sum_{k'} u_{kk'} + \frac{2}{\mu}\langle n_k \rangle + \frac{d\langle n_k \rangle}{d\tau} \tag{N.5}$$

where μ is the probability of mutation per birth event, $\tau = \mu b t$ is dimensionless time rescaled by the mutation and birth rates, and k' and l' represent the first neighbors of genotypes k and l, respectively. The time-dependence is omitted for brevity of notation.

Now, we are ultimately interested in the correlation of genotypes based on their genetic distance, i.e. the number of nucleotide differences between them. Genotype k has distance $z = 0$ from itself, its first neighbors k' stand a distance $z = 1$ from it, and so on. In the limit of an infinitely large ensemble, the covariance u_{kl} between the abundances of all pairs of genotypes (k, l) separated by distance z converges to the same number, u_z. Thus u_{kk} converges to u_0 for every genotype k, $u_{kk'}$ converges to u_1, and so forth. Also, if k and l differ in z loci, then among the first neighbors of k, z differ from l in $z - 1$ loci, $2z$ differ in z loci, and $3(L - z)$ in $z + 1$ loci. The same is true, of course, of the first neighbors of l with respect to k.

We can then rewrite Eqns. (N.4) and (N.5) as

$$\frac{du_0}{d\tau} = -2u_0 + 2u_1 + \frac{2}{\mu}\langle n_k \rangle + \frac{d\langle n_k \rangle}{d\tau} \tag{N.6}$$

$$\frac{du_z}{d\tau} = -2u_z + \frac{2}{3L}(zu_{z-1} + 2zu_z + 3(L - z)u_{z+1}) \tag{N.7}$$

for $z = 1, 2, \ldots, L$. Solving for $\rho_z = u_z/u_0$ in equilibrium, we get

$$\rho_0 = 1 \tag{N.8}$$

$$\rho_1 = 1 - \frac{1}{\mu}\frac{\langle n_k \rangle}{u_0} \tag{N.9}$$

$$\rho_{z+1} = \rho_z - \frac{z}{3(L - z)}(\rho_{z-1} - \rho_z) \tag{N.10}$$

[1]Expressions N.8 to N.10 guarantee that the set $\{\rho_z\}$ forms a monotonically decreasing series in z. This indicates higher correlations between pairs with shorter genetic distance. Thus fortuitous fluctuations in the abundance of a genotype will tend to boost the abundance of its near neighbors in genotype space, and vice-versa. This leads to clustering in genotype space: some regions are rich with abundant genotypes, while others are more sparsely occupied because of negative fluctuations that augment each other. Note that the abundances of all genotypes have the same ensemble mean (Eqn. O.1 in Appendix O): it is the correlations of fluctuations around the mean in specific communities that reveal the structure caused by small mutations under neutral dynamics.

The result suggests that the distribution of distances to nearest modal genotype caused by point mutations would be monotonic. We verify this with simulations in the main test. It is in principle possible to use the mathematical framework of neutral theory to derive $\langle n_k \rangle$ and u_0 as a function of J and μ. For example, by approximating the expected genotype richness in the community using equation (7) in (Vallade and Houchmandzadeh 2003), one can calculate $\langle n_k \rangle \simeq (-\mu \log \mu)^{-1}$. If both the mean and the variance can be so expressed, we will have analytical neutral predictions for the correlations between all pairs of genotypes, and this would be complemented with simulations to determine variability in order to carry out a statistical analysis.

---

[1]Expressions N.8 and N.9 show that $\rho_1 < \rho_0$ and ρ_1 is positive if $\langle n_k \rangle < \mu u_0$. We conjecture that ρ_1 is always positive, and tends to zero for increasingly lower mutation rates. Notice that although it may seem from equation N.9 that ρ_1 is negative for low μ, this is not necessarily the case because $\langle n_k \rangle$ and u_0 are also dependent on μ. Such may not be the case of a zero-sum formulation where the community size is fixed, since higher abundance for certain genotypes necessarily means lower abundance for others. But in either case the set $\{rho_z\}$ will be monotonically decreasing, indicating more positive correlations between pairs with shorter genetic distance.

Appendix O

Calculating correlations with the master equation

The master equation can be used to calculate dynamic equations for the moments $\langle n_k(t) \rangle$, $\langle n_k^2(t) \rangle$, $\langle n_k n_l(t) \rangle$, which can then be used to find a dynamic equation for the covariances $u_{kl}(t) = \langle n_k n_l(t) \rangle - \langle n_k(t) \rangle \langle n_l(t) \rangle$. (We'll subsequently omit the time-dependence for brevity of notation.)

To illustrate the derivation method, let's start with the equation for the expected abundance of genotype k, $\langle n_k \rangle$. Letting boldface $\mathbf{n}$ stand for the set $n_1, \ldots, n_k, \ldots, n_{4^L}$, we write

$$\langle n_k \rangle = \sum_{\mathbf{n}} n_k P(\mathbf{n})$$

$$\frac{d \langle n_k \rangle}{dt} = \frac{d}{dt} \sum_{\mathbf{n}=0}^{\infty} n_k P(\mathbf{n})$$

$$= \sum_{n_k=0}^{\infty} n_k \sum_{n_{l \neq k}=0}^{\infty} \frac{\partial P(\mathbf{n})}{\partial t}$$

$$= \sum_{n_k=0}^{\infty} n_k \sum_{n_{l \neq k}=0}^{\infty} \sum_{i=1}^{4^L} (P(n_i - 1)W^+(n_i - 1) - P(n_i)W^+(n_i)$$

$$+ P(n_i + 1)W^-(n_i + 1) - P(n_i)W^-(n_i)).$$

In the last line, we substituted the master equation (N.1) from Appendix N. Also, notice that $P(n_i)$ is shorthand notation for $P(n_1, \ldots, n_i, \ldots, n_{4^L})$ emphasizing the probability that the abundance of genotype i is equal to n_i. Next, we recognize that all but the k-term in the

sum $\sum\limits_i$ vanish, because once we sum over $\sum\limits_{n_l}$, the l-term in $\sum\limits_i$ cancels out (Houchmandzadeh and Vallade, 2003). This leaves us with

$$
\begin{aligned}
\frac{d\langle n_k \rangle}{dt} &= \sum_{\mathbf{n}=0}^{\infty} n_k P(n_k - 1) W^+(n_k - 1) - n_k P(n_k) W^+(n_k) \\
&\quad + n_k P(n_k + 1) W^-(n_k + 1) - n_k P(n_k) W^-(n_k) \\
&= \sum_{\mathbf{n}=0}^{\infty} W^+(n_k) P(n_k) - W^-(n_k) P(n_k) \\
&= \langle W^+(n_k) - W^-(n_k) \rangle \\
&= -\frac{\mu}{J} \left(\langle n_k \rangle - \frac{1}{3L} \sum_{k'} \langle n_{k'} \rangle \right).
\end{aligned}
$$

In the last line, we plugged in the transition rates that define the neutral model, Eqns. (N.2) and (N.3). In words, the changes in the expected value of genotype k in a fixed-size community are due to the balance between mutations in and out of genotype k. The equilibrium solution is

$$
\langle n_k \rangle = \frac{1}{3L} \sum_{k'} \langle n_{k'} \rangle. \tag{O.1}
$$

That is, the expected abundance of genotype k is the mean of the expected abundance of its first neighbors (a genotype has $3L$ first neighbors). Extending this result to all genotypes, we conclude that all genotypes have the same mean. This reflects the neutrality of the model.

Now we can move on to $\langle n_k n_l \rangle$, with $l \neq k$. Again, the master equation can help us:

$$
\begin{aligned}
\frac{d\langle n_k n_l \rangle}{dt} &= \frac{d}{dt} \sum_{\mathbf{n}=0}^{\infty} n_k n_l P(\mathbf{n}) \\
&= \sum_{\mathbf{n}=0}^{\infty} n_k n_l \frac{\partial P(\mathbf{n})}{\partial t} \\
&= \sum_{\mathbf{n}=0}^{\infty} n_k n_l \sum_{i=1}^{4^L} \Big(P(n_i - 1) W^+(n_i - 1) - P(n_i) W^+(n_i) \\
&\quad + P(n_i + 1) W^-(n_i + 1) - P(n_i) W^-(n_i) \Big).
\end{aligned}
$$

Using the same logic as above, the only surviving terms in the sum $\sum\limits_i$ are the k- and

l-terms, which can then be simplified upon summing over n_k and n_l.

$$\frac{d\langle n_k n_l \rangle}{dt} = \sum_{\mathbf{n}=0}^{\infty} n_k n_l \Big[P(n_k - 1)W^+(n_k - 1) - P(n_k)W^+(n_k)$$
$$+ P(n_k + 1)W^-(n_k + 1) - P(n_k)W^-(n_k)$$
$$+ P(n_l - 1)W^+(n_l - 1) - P(n_l)W^+(n_l)$$
$$+ P(n_l + 1)W^-(n_l + 1) - P(n_l)W^-(n_l) \Big]$$
$$= \sum_{\mathbf{n}=0}^{\infty} n_l \left(W^+(n_k)P(n_k) - W^-(n_k)P(n_k) \right)$$
$$+ n_k \left(W^+(n_l)P(n_l) - W^-(n_l)P(n_l) \right)$$
$$= \langle n_l \left(W^+(n_k) - W^-(n_k) \right) \rangle + \langle n_k \left(W^+(n_l) - W^-(n_l) \right) \rangle$$
$$= \frac{\mu}{J} \left(-2\langle n_l n_k \rangle + \frac{1}{3L} \sum_{k'} \langle n_l n_{k'} \rangle + \frac{1}{3L} \sum_{l'} \langle n_k n_{l'} \rangle \right),$$

where k' and l' are the first neighbors of genotypes k and l, respectively.
Now we can piece those together to get the dynamic equation for u_{kl}:

$$\frac{du_{kl}}{d\tau} = \frac{d\langle n_k n_l \rangle}{d\tau} - \langle n_k \rangle \frac{d\langle n_l \rangle}{d\tau} - \langle n_l \rangle \frac{d\langle n_k \rangle}{d\tau}$$
$$= -2u_{kl} + \frac{1}{3L} \sum_{k'} u_{lk'} + \frac{1}{3L} \sum_{l'} u_{kl'}, \qquad (O.2)$$

where $\tau = \mu b t$ is dimensionless time rescaled by the mutation and birth rates.
The equation for u_{kk} is similar but slightly different:

$$\frac{d\langle n_k n_k \rangle}{dt} = \sum_{\mathbf{n}=0}^{\infty} n_k^2 \Big[P(n_k - 1)W^+(n_k - 1) - P(n_k)W^+(n_k)$$
$$+ P(n_k + 1)W^-(n_k + 1) - P(n_k)W^-(n_k) \Big]$$
$$= 2\langle n_k(W^+(n_k) - W^-(n_k)) \rangle + \langle W^+(n_k) + W^-(n_k) \rangle$$
$$= 2\mu b \left(-\langle n_k^2 \rangle + \frac{1}{3L} \sum_{k'} \langle n_k n_{k'} \rangle + \frac{1}{\mu} \langle n_k \rangle \right) + \frac{d\langle n_k \rangle}{dt}$$

$$\frac{du_{kk}}{d\tau} = 2 \left(-u_{kk} + \frac{1}{3L} \sum_{k'} u_{kk'} + \frac{1}{\mu} \langle n_k \rangle \right) + \frac{d\langle n_k \rangle}{d\tau} \qquad (O.3)$$

192

Appendix P

Summary of our example power analysis

Building on our neutral simulations, we implemented selection by setting up death rates d as a function of the Hamming distance z to the dominant genotype – $d(z) = 1 + 0.05z$. This makes for a threefold difference between the death rates of the fittest and least fit genotypes. In every other respect, our selective regime is identical to the neutral regime. Results are shown in figure C1. The ribbons show the 95% confidence intervals for the metric applied on our set of neutral (grey) and selective (blue) regimes. Both show a monotonic decrease with genetic distance, but the selective regime is steeper. To quantify that difference, we ran a χ^2-based power analysis.

For each of the neutral communities and each of the selective regime communities, we calculate its χ^2-statistic as $\chi^2 = \sum_{z=1}^{L} (O_z - E_z)^2 / E_z$, where E_z is the mean of the metric at distance z across the set of simulated neutral communities, and O_z is the observed value for that particular community. We then compare the χ^2 of each selective-regime community to the set of χ^2 for all neutral communities, and calculate its p-value by determining the fraction of neutral χ^2 that are larger than the χ^2 of the selective community. The power of the test is the percentage of selective regime cases that were determined to be significantly different from the set of neutral communities, i.e., the percentage of selective regime cases whose p-value was lower than 0.05. The metric was able to distinguish 200 out of 228 selective regime communities from the set of 127 neutral communities at the $p = 0.05$ level, thus at 88% power. Lessening the selection gradient would lower power.

Figure P.1: Distribution of distances to nearest modal genotype observed in neutral communities driven by drift and mutation. (Normalized Hamming distance is defined as z/L, where z is the number of discordant loci and L is the total number of loci.) Numbers on the y axis reflect the frequency of rare-modal pairs. Grey band corresponds to neutral scenario with restriction to point mutations, and blue band corresponds to neutral scenario with mutations of unconstrained size. Bands span 95% confidence intervals across 127 (neutral) or 228 (selective) runs. Community size for all runs is 1,000, and mutation rate is 0.1 per genome per generation, for a genome with 40 loci. Selection implemented on death rates: rate $d(z)$ for a genotype with Hamming distance z to dominant genotype is $d(z) = 1 + 0.05z$.
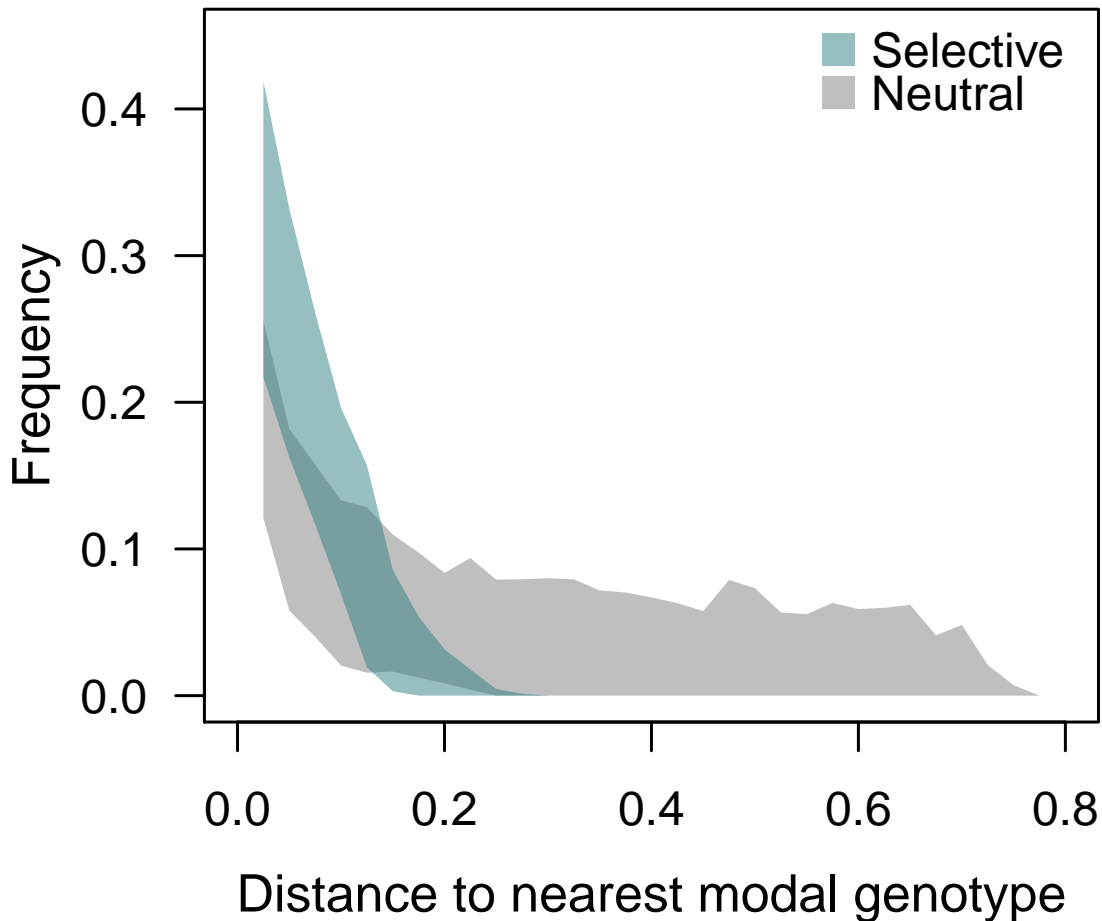
Bibliography

Abrams, P. (1975). Limiting Similarity and the Form of the Competition Coefficient. *Theoretical population biology*, 8:356–375.

Abrams, P. (1983). The Theory of Limiting Similarity. *Annual Review of Ecology and Systematics*, 14(1):359–376.

Abrams, P. (1990). Ecological vs Evolutionary Consequencs of Competition. *Oikos*, 57(1):147–151.

Abrams, P. a. and Rueffler, C. (2009). Coexistence and limiting similarity of consumer species competing for a linear array of resources. *Ecology*, 90(3):812–22.

Abrams, P. a., Rueffler, C., and Dinnage, R. (2008). Competition-similarity relationships and the nonlinearity of competitive effects in consumer-resource systems. *The American naturalist*, 172(4):463–74.

Adler, F. R. and Mosquera, J. (2000). Is Space Necessary? Interference Competition and Limits to Biodiversity. *Ecology*, 81(11):3226.

Adler, F. R. and Muller-Landau, H. C. (2005). When do localized natural enemies increase species richness? *Ecology Letters*, 8(4):438–447.

Adler, P. B., Ellner, S. P., and Levine, J. M. (2010). Coexistence of perennial plants: an embarrassment of niches. *Ecology letters*, 13(8):1019–29.

Adler, P. B., Fajardo, A., Kleinhesselink, A. R., and Kraft, N. J. B. (2013). Trait-based tests of coexistence mechanisms. *Ecology letters*.

Adler, P. B., HilleRisLambers, J., and Levine, J. M. (2007). A niche for neutrality. *Ecology letters*, 10(2):95–104.

Adler, P. B., Salguero-Gómez, R., Compagnoni, A., Hsu, J. S., Ray-Mukherjee, J., Mbeau-Ache, C., and Franco, M. (2014). Functional traits explain variation in plant life history strategies. *Proceedings of the National Academy of Sciences of the United States of America*, 111(2):740–5.

Akey, J. M., Eberle, M. a., Rieder, M. J., Carlson, C. S., Shriver, M. D., Nickerson, D. a., and Kruglyak, L. (2004). Population history and natural selection shape patterns of genetic variation in 132 genes. *PLoS biology*, 2(10):e286.

Albert, C. H., de Bello, F., Boulangeat, I., Pellet, G., Lavorel, S., and Thuiller, W. (2012). On the importance of intraspecific variability for the quantification of functional diversity. *Oikos*, 121(1):116–126.

Albert, C. H., Grassein, F., Schurr, F. M., Vieilledent, G., and Violle, C. (2011). When and how should intraspecific variability be considered in trait-based plant ecology? *Perspectives in Plant Ecology, Evolution and Systematics*, 13(3):217–225.

Alonso, D., Etienne, R. S., and McKane, A. J. (2006). The merits of neutral theory. *Trends in ecology & evolution*, 21(8):451–7.

Araújo, M. S., Bolnick, D. I., and Layman, C. a. (2011). The ecological causes of individual specialisation. *Ecology letters*, 14(9):948–58.

Armstrong, R. A. and McGehee, R. (1976). Coexistence of Species Competing for Shared Resources. *Theoretical Population Biology*, 9:317–328.

Barabás, G. and D'Andrea, R. (2016). The effect of intraspecific variation and heritability on community pattern and robustness. *Ecology Letters*.

Barabás, G., D'Andrea, R., Rael, R., Meszéna, G., and Ostling, A. M. (2013a). Emergent Neutrality or Hidden Niches? *Oikos*, 122:1565–1572.

Barabás, G., D'Andrea, R., Rael, R., Meszéna, G., and Ostling, A. M. (2013b). Species packing in nonsmooth competition models. *Theoretical Ecology*, 6(1):1–19.

Barabás, G. and Meszéna, G. (2009). When the exception becomes the rule: the disappearance of limiting similarity in the Lotka-Volterra model. *Journal of Theoretical Biology*, 258(1):89–94.

Barabás, G., Meszéna, G., and Ostling, A. (2012). Community robustness and limiting similarity in periodic environments. *Theoretical Ecology*, 5(2):265–282.

Barabás, G., Pásztor, L., Meszéna, G., and Ostling, A. (2014). Sensitivity analysis of coexistence in ecological communities: theory and application. *Ecology Letters*, 17:1479–1494.

Beaumont, M. A. (2010). Approximate Bayesian Computation in Evolution and Ecology. *Annual Review of Ecology, Evolution, and Systematics*, 41(1):379–406.

Bell, G. (2001). Neutral macroecology. *Science (New York, N.Y.)*, 293(5539):2413–8.

Berlin, T. H. and Kac, M. (1952). The spherical model of a ferromagnet. *Physical Review*, 86(6):821–835.

Bolnick, D. I., Amarasekare, P., Araújo, M. S., Bürger, R., Levine, J. M., Novak, M., Rudolf, V. H. W., Schreiber, S. J., Urban, M. C., and Vasseur, D. a. (2011). Why intraspecific trait variation matters in community ecology. *Trends in ecology & evolution*, 26(4):183–92.

Bond, E. M. and Chase, J. M. (2002). Biodiversity and ecosystem functioning at local and regional spatial scales. *Ecology Letters*, 5(4):467–470.

Bonsall, M. B., Jansen, V. a. a., and Hassell, M. P. (2004). Life history trade-offs assemble ecological guilds. *Science (New York, N.Y.)*, 306(5693):111–4.

Botta-Dukát, Z. (2005). Rao's quadratic entropy as a measure of functional diversity based on multiple traits. *Journal of Vegetation Science*, 16(5):533–540.

Breitschwerdt, E., Jandt, U., and Bruelheide, H. (2015). Do newcomers stick to the rules of the residents? Designing trait-based community assembly tests. *Journal of Vegetation Science*, 26(2):219–232.

Brychkov, Y. A. and Shirokov, Y. M. (1970). Asymptotic behavior of fourier transforms. *Theoretical and mathematical physics*, 4(3):847–853.

Burns, J. and Strauss, S. (2011). More closely related species are more ecologically similar in an experimental test. *Proceedings of the National Academy of Sciences*, 108(13):5302–5307.

Buttel, L. a., Durrett, R., and Levin, S. a. (2002). Competition and species packing in patchy environments. *Theoretical population biology*, 61(3):265–76.

Carboni, M., de Bello, F., Janeček, v., Doležal, J., Horník, J., Lepš, J., Reitalu, T., and Klimešová, J. (2014). Changes in trait divergence and convergence along a productivity gradient in wet meadows. *Agriculture, Ecosystems and Environment*, 182:96–105.

Carlucci, M. B., Debastiani, V. J., Pillar, V. D., and Duarte, L. D. S. (2015). Between- and within-species trait variability and the assembly of sapling communities in forest patches. *Journal of Vegetation Science*, 26(1):21–31.

Carroll, I. T. and Nisbet, R. M. (2015). Departures from neutrality induced by niche and relative fitness differences. *Theoretical Ecology*, 8(4):449–465.

Case, T. (2000). An illustrated guide to theoretical ecology.

Cavalli-Sforza, L. (1966). Population structure and human evolution. *Proceedings of the Royal Society B*, 164:362–379.

Cavender-Bares, J., Keen, A., and Miles, B. (2006). Phylogenetic Structure of Floridian Plant Communities Depends on Taxonomic and Spatial Scale. *Ecology*, 87(7):S109–S122.

Cavender-Bares, J., Kozak, K. H., Fine, P. V. a., and Kembel, S. W. (2009). The merging of community ecology and phylogenetic biology. *Ecology letters*, 12(7):693–715.

Chase, J. M. and Leibold, M. (2003). *Ecological Niches: Linking Classical and Contemporary Approaches*. The University of Chicago Press.

Chave, J., Muller-Landau, H. C., and Levin, S. A. (2002). Comparing classical community models: theoretical consequences for patterns of diversity. *The American Naturalist*, 159(1):1–23.

Chesson, P. (1990). Macarthur's consumer-resource model. *Theoretical Population Biology*, 37(1):26–38.

Chesson, P. (1991). A Need for Niches ? *Trends in Ecology & Evolution*, 6(1):26–28.

Chesson, P. (2000). Mechanisms of maitenance of species diversity. *Annual Review of Ecology and Systematics*, 31(2000):343–358.

Chesson, P. and Ellner, S. (1989). Invasibility and stochastic boundedness in monotonic competition models. *Journal of Mathematical Biology*, 27(2):117–138.

Chesson, P. and Huntly, N. (1997). The Roles of Harsh and Fluctuating Conditions in the Dynamics of Ecological Communities. *The American Naturalist*, 150(5):519–553.

Chisholm, R. A. and Lichstein, J. W. (2009). Linking dispersal, immigration and scale in the neutral theory of biodiversity. *Ecology Letters*, 12(12):1385–93.

Chisholm, R. A. and Pacala, S. W. (2010). Niche and neutral models predict asymptotically equivalent species abundance distributions in high-diversity ecological communities. *Proceedings of the National Academy of Sciences of the United States of America*, 107(36):15821–5.

Clark, J. S. (2012). The coherence problem with the Unified Neutral Theory of Biodiversity. *Trends in Ecology & Evolution*, 27(4).

Cody, M. L. (1999). Assembly rules at different scales in plant and bird communities. In Weiher, E. and Keddy, P., editors, *Ecological assembly rules: perspecitves, advances, retreats*, pages 165–204. Cambridge University Press.

Condit, R., Chisholm, R. A., and Hubbell, S. P. (2012). Thirty Years of Forest Census at Barro Colorado and the Importance of Immigration in Maintaining Diversity. *PLoS ONE*, 7(11):1–6.

Coomes, D. A. and Grubb, P. J. (2003). Colonization, tolerance, competition and seed-size variation within functional groups. *Trends in Ecology & Evolution*, 18(6):283–291.

Cornwell, W. K. and Ackerly, D. D. (2009). Community assembly and shifts in plant trait distributions across an environmental gradient in coastal California. *Ecological Monographs*, 79(1):109–126.

Cornwell, W. K., Schwilk, D. W., and Ackerly, D. D. (2006). A trait-based test for habitat filtering: Convex hull volume. *Ecology*, 87(6):1465–1471.

Courbaud, B., Vieilledent, G., and Kunstler, G. (2010). Intra-specific variability and the competition–colonisation trade-off: coexistence, abundance and stability patterns. *Theoretical Ecology*, 5(1):61–71.

Czernilofsky, A., Kurland, C., and Stöffler, G. (1975). 30s ribosomal proteins associated with the 3'-terminus of 16s rna. *Febs Letters*, 58(1-2):281–284.

D'Andrea, R., Barabás, G., and Ostling, A. (2013). Revising the tolerance-fecundity trade-off; or, on the consequences of discontinuous resource use for limiting similarity, species diversity, and trait dispersion. *The American Naturalist*, 181(4):E91–101.

D'Andrea, R. and Ostling, A. (2015). Can Clustering in Genotype Space Reveal "Niches"? *The American Naturalist*, 187(1):130–135.

D'Andrea, R. and Ostling, A. (2016). Challenges in linking trait patterns to niche differentiation. *Oikos*.

Dayan, T. and Simberloff, D. (2005). Ecological and community-wide character displacement: the next generation. *Ecology Letters*, 8(8):875–894.

de Bello, F. (2011). The quest for trait convergence and divergence in community assembly: are null-models the magic wand? *Global Ecology and Biogeography*, pages no–no.

de Bello, F., Price, J., Munkemuller, T., Liira, J., Zobel, M., Thuiller, W., Gerhold, P., Gotzenberger, L., Lavergne, S., Leps, J., Zobel, K., and Partel, M. (2012). Functional species pool framework to test for biotic effects on community assembly. *Ecology*, 93(10):2263–2273.

Deza, M. M. and Deza, E. (2006). *Dictionary of Distances*. Elsevier.

Doebeli, M., Blok, H. J., Leimar, O., and Dieckmann, U. (2007). Multimodal pattern formation in phenotype distributions of sexual populations. *Proceedings. Biological sciences / The Royal Society*, 274(1608):347–357.

Eriksson, O. (2005). Game theory provides no explanation for seed size variation in grasslands. *Oecologia*, 144(1):98–105.

Ernebjerg, M. and Kishony, R. (2011). Dynamic phenotypic clustering in noisy ecosystems. *PLoS computational biology*, 7(3):e1002017.

Falconer, D. S. (1981). *Introduction to quantitative genetics*. Pearson Education India.

Fort, H., Scheffer, M., and van Nes, E. (2010). The clumping transition in niche competition: a robust critical phenomenon. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(05):P05005.

Fort, H., Scheffer, M., and van Nes, E. H. (2009). The paradox of the clumps mathematically explained. *Theoretical Ecology*, 2(3):171–176.

Frank, S. a. (2009). The common patterns of nature. *Journal of evolutionary biology*, 22(8):1563–85.

Franzén, D. (2004). Plant species coexistence and dispersion of seed traits in a grassland. *Ecography*, 27(2):218–224.

Funkhouser, L. J. and Bordenstein, S. R. (2013). Mom Knows Best: The Universality of Maternal Microbial Transmission. *PLoS Biology*, 11(8):1–9.

Fussmann, G. F., Loreau, M., and Abrams, P. A. (2007). Eco-evolutionary dynamics of communities and ecosystems. *Functional Ecology*, 21(3):465–477.

Futuyma, D. J. and Moreno, G. (1988). The evolution of ecological specialization. *Annual Review of Ecology and Systematics*, pages 207–233.

Gause, G. F. (1934). Experimental Analysis of Vito Volterra's Mathematical Theory of the Struggle for Existence. *Science*, 79(2036):16–17.

Geritz, S. A., van der Meijden, E., and Metz, J. A. (1999). Evolutionary dynamics of seed size and seedling competitive ability. *Theoretical population biology*, 55(3):324–343.

Gillespie, D. T. (1977). Exact Stochastic Simulation of couple chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361.

Google Ngram Viewer (2015). Limiting Similarity. *Google books Ngram Viewer*.

Gotelli, N. J. and Graves, G. R. (1996). *Null Models in Ecology*. Smithsonian Institution.

Gotelli, N. J. and McGill, B. J. (2006). Null versus neutral models: what's the difference? *Ecography*, 29(5):793–800.

Götzenberger, L., de Bello, F., Bråthen, K. A., Davison, J., Dubuis, A., Guisan, A., Lepš, J., Lindborg, R., Moora, M., Pärtel, M., Pellissier, L., Pottier, J., Vittoz, P., Zobel, K., and Zobel, M. (2012). Ecological assembly rules in plant communities-approaches, patterns and prospects. *Biological Reviews*, 87(1):111–127.

Grant, P. R. (1969). Colonization of islands by ecologically dissimilar species of birds. *Canadian Journal of Zoology*, 47(1):41–43.

Grant, P. R. and Grant, B. R. (2006). Evolution of character displacement in darwin's finches. *science*, 313(5784):224–226.

Gravel, D., Canham, C. D., Beaudet, M., and Messier, C. (2006). Reconciling niche and neutrality: the continuum hypothesis. *Ecology letters*, 9(4):399–409.

Gross, K., Cardinale, B. J., Fox, J. W., Gonzalez, A., Loreau, M., Polley, H. W., Reich, P. B., and van Ruijven, J. (2014). Species richness and the temporal stability of biomass production: a new analysis of recent biodiversity experiments. *The American naturalist*, 183:1–12.

Gyllenberg, M. and Meszéna, G. (2005). On the impossibility of coexistence of infinitely many strategies. *Journal of Mathematical Biology*, 50:133–160.

Haegeman, B. and Loreau, M. (2011). A mathematical synthesis of niche and neutral theories in community ecology. *Journal of Theoretical Biology*, 269(1):150–65.

Haegeman, B., Sari, T., and Etienne, R. S. (2014). Predicting coexistence of plants subject to a tolerance-competition trade-off. *Journal of Mathematical Biology*, 68(7):1815–47.

Hankin, R. K. S. (2007). Introducing untb, an r package for simulating ecological drift under the unified nuetral theory of biodiversity. *Journal of Statistical Software*, 22.

Hardin, G. (1960). The competitive exclusion principle.

Harte, J. (2003). Ecology: Tail of death and resurrection. *Nature*, 424(6952):1006–1007.

Harte, J. (2011). *Maximum entropy and ecology: a theory of abundance, distribution, and energetics*. Oxford University Press.

Hedrick, P. W. (2011). *Genetics of populations*. Jones & Bartlett Learning.

Herben, T. and Goldberg, D. E. (2014). Community assembly by limiting similarity vs. competitive hierarchies: Testing the consequences of dispersion of individual traits. *Journal of Ecology*, 102(1):156–166.

Hernández-García, E., López, C., Pigolotti, S., and Andersen, K. H. (2009). Species competition: coexistence, exclusion and clustering. *Philosophical Transactions of the Royal Society A*, 367(1901):3183–95.

HilleRisLambers, J., Adler, P., Harpole, W., Levine, J., and Mayfield, M. (2012). Rethinking Community Assembly through the Lens of Coexistence Theory. *Annual Review of Ecology, Evolution, and Systematics*, 43(1):227–248.

Holt, R. D. (2006). Emergent neutrality. *Trends in Ecology & Evolution*, 21(10):531–3.

Houchmandzadeh, B. and Vallade, M. (2003). Clustering in neutral ecology. *Physical Review E*, 68(6):061912.

Hubbell, S. P. (2001). *The Unified Neutral Theory of Biodiversity and Biogeography*. Princeton University Press, Princeton, NJ.

Hubbell, S. P. (2006). Neutral theory and the evolution of ecological equivalence. *Ecology*, 87(6):1387–98.

Hubbell, S. P. and Foster, R. B. (1986). Biology, chance, and history and the structure of tropical rain forest tree communities. *Community Ecology*, 19:314–329.

Hulshof, C. M. and Swenson, N. G. (2010). Variation in leaf functional trait values within and across individuals and species: an example from a Costa Rican dry forest. *Functional Ecology*, 24(1):217–223.

Hutchinson, G. E. (1959). Homage to Santa Rosalia or Why Are There So Many Kinds of Animals? *The American Naturalist*, 93(870):145–159.

Ingram, T. and Shurin, J. B. (2009). Trait-based assembly and phylogenetic structure in northeast Pacific rockfish assemblages. *Ecology*, 90(9):2444–53.

Jeraldo, P., Sipos, M., Chia, N., Brulc, J. M., Dhillon, A. S., Konkel, M. E., Larson, C. L., Nelson, K. E., Qu, A., Schook, L. B., Yang, F., White, B. A., and Goldenfeld, N. (2012). Quantification of the relative roles of niche and neutral processes in structuring gastrointestinal microbiomes. *PNAS*, 109(25):9692–9698.

Jiang, L., Tan, J., and Pu, Z. (2010). An experimental test of Darwin's naturalization hypothesis. *The American Naturalist*, 175(4):415–23.

Johansson, M. E. and Keddy, P. A. (1991). Intensity and asymmetry of competition between plant pairs of different degrees of similarity: an experimental study on two guilds of wetland plants. *Oikos*, 60(1):27–34.

Kinzig, A., Levin, S., Dushoff, J., and Pacala, S. (1999). Limiting similarity, species packing, and system stability for hierarchical competition-colonization models. *The American Naturalist*, 153(4):371–383.

Kneitel, J. M. and Chase, J. M. (2004). Trade-offs in community ecology: linking spatial scales and species coexistence. *Ecology Letters*, 7(1):69–80.

Kohl, K. D., Stengel, A., and Dearing, M. D. (2015). Inoculation of tannin-degrading bacteria into novel hosts increases performance on tannin-rich diets. *Environmental microbiology*.

Kohyama, T. (1993). Size-Structured Tree Populations in Gap-Dynamic Forest–The Forest Architecture Hypothesis for the Stable Coexistence of Species. *Journal of Ecology*, 81(1):131–143.

Kohyama, T. (2006). The effect of patch demography on the community structure of forest trees. *Ecological Research*, 21(3):346–355.

Kraft, N. J. B. and Ackerly, D. D. (2010). Functional trait and phylogenetic tests of community assembly across spatial scales in an Amazonian forest. *Ecological Monographs*, 80(3):401–422.

Kraft, N. J. B., Godoy, O., and Levine, J. M. (2015). Plant functional traits and the multidimensional nature of species coexistence - Supporting Information. *Proceedings of the National Academy of Sciences*, 112(3):797–802.

Kraft, N. J. B., Valencia, R., and Ackerly, D. D. (2008). Functional traits and niche-based tree community assembly in an Amazonian forest. *Science*, 322:580–582.

Krebs, C. J. (2001). *Ecology: The experimental analysis of distribution and abundance*. Pearson Benjamin Cummings.

Kunstler, G., Lavergne, S., Courbaud, B., Thuiller, W., Vieilledent, G., Zimmermann, N. E., Kattge, J., and Coomes, D. a. (2012). Competitive interactions between forest trees are driven by species' trait hierarchy, not phylogenetic or functional similarity: implications for forest community assembly. *Ecology letters*, 15(8):831–40.

Lake, J. K. and Ostling, A. (2009). Comment on "Functional traits and niche-based tree community assembly in an Amazonian forest". *Science (New York, N.Y.)*, 324(5930):1015; author reply 1015.

Laliberté, E. and Legendre, P. (2010). A distance-based framework for measuring functional diversity from multiple traits. *Ecology*, 91(1):299–305.

Lasky, J. R., Bachelot, B., Muscarella, R., Schwartz, N., Forero-Montaña, J., Nytch, C. J., Swenson, N. G., Thompson, J., Zimmerman, J. K., and Uriarte, M. (2015). Ontogenetic shifts in trait-mediated mechanisms of plant community assembly. *Ecology*, 96(8):2157–2169.

Lasky, J. R., Uriarte, M., Boukili, V. K., and Chazdon, R. L. (2014a). Trait-mediated assembly processes predict successional changes in community diversity of tropical forests. *Proceedings of the National Academy of Sciences of the United States of America*, 111(15):5616–21.

Lasky, J. R., Yang, J., Zhang, G., Cao, M., Tang, Y., and Keitt, T. H. (2014b). The role of functional traits and individual variation in the co-occurrence of Ficus species. *Ecology*, 95(4):978–990.

Lauring, A. S., Frydman, J., and Andino, R. (2013). The role of mutational robustness in RNA virus evolution. *Nature reviews. Microbiology*, 11(5):327–36.

Lawson, D. J. and Jensen, H. J. (2007). Neutral evolution in a biological population as diffusion in phenotype space: Reproduction with local mutation but without selection. *Physical Review Letters*, 98(March):2–5.

Lawton, J. H. (1999). Are there general laws in ecology? *Oikos*, pages 177–192.

Leigh, E. G., Wright, S. J., Herre, E. A., and Putz, F. E. (1993). The decline of tree diversity on newly isolated tropical islands: A test of a null hypothesis and some implications. *Evolutionary Ecology*, 7(1):76–102.

Leimar, O., Doebeli, M., and Dieckmann, U. (2008). Evolution of Phenotypic Clusters Through Competition and Local Adaptation Along an Environmental Gradient. *Evolution*, 62(4):807–822.

Letten, A. D., Keith, D. A., and Tozer, M. G. (2014). Phylogenetic and functional dissimilarity does not increase during temporal heathland succession. *Proceedings of the Royal Society B*, 281(20142102).

Levin, S. (1970). Community equilibria and stability, and an extension of the competitive exclusion principle. *The American Naturalist*, 104(939):413–423.

Levin, S. and Pimentel, D. (1981). Selection of intermediate rates of increase in parasite-host systems. *American Naturalist*, pages 308–315.

Levins, R. (1969). Some demographic and genetic consequences of environmental heterogeneity for biological control. *Bulletin of the Entomological society of America*, 15(3):237–240.

Levins, R. and Culver, D. (1971). Regional Coexistence of Species and Competition between Rare Species. *Proceedings of the National Academy of Sciences of the United States of America*, 68(6):1246–1248.

Levy, R. and Borenstein, E. (2013). Metabolic modeling of species interaction in the human microbiome elucidates community-level assembly rules. *PNAS*, 110(31):12804–12809.

MacArthur, R. (1970). Species packing and competitive equilibrium for many species. *Theoretical population biology*, 1(1):1–11.

MacArthur, R. and Levins, R. (1967). The Limiting Similarity , Convergence , and Divergence of Coexisting Species. *The American Naturalist*, 101(921):377–385.

MacArthur, R. H. (1962). Some generalized theorems of natural selection. *Proceedings of the National Academy of Sciences*, 48(11):1893–1897.

MacArthur, R. H. and Wilson, E. O. (1967). *Theory of Island Biogeography.(MPB-1)*, volume 1. Princeton University Press.

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297, Oakland, CA, USA. University of California, Berkeley, University of California Press.

Maechler, M., Rousseeuw, P., Struyf, A., Hubert, M., and Hornik, K. (2016). *cluster: Cluster Analysis Basics and Extensions*. R package version 2.0.4 — For new features, see the 'Changelog' file (in the package source).

Magurran, A. E. and Henderson, P. a. (2003). Explaining the excess of rare species in natural species abundance distributions. *Nature*, 422(April):714–716.

Maire, V., Gross, N., Börger, L., Proulx, R., Wirth, C., Pontes, L. D. S., Soussana, J.-F., and Louault, F. (2012). Habitat filtering and niche differentiation jointly explain species relative abundance within grassland communities along fertility and disturbance gradients. *The New phytologist*, 196(2):497–509.

Manly, B. (1991). *Randomization and Monte-Carlo Mehtods in Biology*. Chapman & Hall, London.

Mason, N. and Wilson, J. (2006). Mechanisms of species coexistence in a lawn community: mutual corroboration between two independent assembly rules. *Community Ecology*, 7(1):109–116.

Mason, N. W. H., Mouillot, D., Lee, W. G., Wilson, J. B., and Functional, J. B. (2005). Functional richness, functional evenness and functional divergence: the primary components of functional diversity. *Oikos*, 111(February):112–118.

May, R. M. (1973). *Stability and Complexity in Model Ecosystems*. Princeton University Press, Princeton, NJ.

May, R. M. and MacArthur, R. H. (1972). Niche overlap as a function of environmental variability. *Proceedings of the National Academy of Sciences of the United States of America*, 69(5):1109–1113.

Mayfield, M. M. and Levine, J. M. (2010). Opposing effects of competitive exclusion on the phylogenetic structure of communities. *Ecology letters*, 13(9):1085–93.

McGill, B. J. (2003). Strong and weak tests of macroecological theory. *Oikos*, 102(3):679–685.

McGill, B. J., Enquist, B. J., Weiher, E., and Westoby, M. (2006). Rebuilding community ecology from functional traits. *Trends in ecology & evolution*, 21(4):178–85.

McGill, B. J., Hadly, E. a., and Maurer, B. a. (2005). Community inertia of Quaternary small mammal assemblages in North America. *Proceedings of the National Academy of Sciences of the United States of America*, 102(46):16701–6.

Meszéna, G. (2005). Adaptive dynamics: the continuity argument. *Journal of evolutionary biology*, 18(5):1182–1185.

Meszéna, G., Gyllenberg, M., Pasztor, L., and Metz, J. (2006). Competitive exclusion and limiting similarity: A unified theory. *Theoretical population biology*, 69:68–87.

Metz, J., Geritz, S., Meszéna, G., Jacobs, F., and van Heerwaarden, J. (1996). Adaptive dynamics: a geometrical study of the consequences of nearly faithful reproduction. *Stochastic and Spatial Structures of Dynamical Systems. North-Holland, Amsterdam*, pages 183–231.

Metz, J., Mylius, S., and Diekmann, O. (2008). When does evolution optimize? *Evolutionary Ecology Research*, 10(5):629–654.

Moles, A. T. and Westoby, M. (2004). Seedling survival and seed size: a synthesis of the literature. *Journal of Ecology*, 92(3):372–383.

Mouillot, D., Mason, W. H. N., Dumay, O., and Wilson, J. B. (2005). Functional regularity: a neglected aspect of functional diversity. *Oecologia*, 142(3):353–359.

Muller-Landau, H. C. (2010). The tolerance – fecundity trade-off and the maintenance of diversity in seed size. *Proceedings of the National Academy of Sciences*, 107(9):4242–4247.

Nee, S. (2005). The neutral theory of biodiversity: Do the numbers add up? *Functional Ecology*, 19(1):173–176.

Ochman, H., Elwyn, S., and Moran, N. a. (1999). Calibrating bacterial evolution. *Proceedings of the National Academy of Sciences of the United States of America*, 96(22):12638–12643.

Ostling, A. M. (2012). Large-scale spatial synchrony and the stability of forest biodiversity revisited. *Journal of Plant Ecology*, 5(1):52–63.

Pacala, S. W., Canham, C. D., and Silander Jr, J. A. (1993). Forest models defined by field measurements: I. The design of a northeastern forest simulator. *Canadian Journal of Forest Research*, 23:1980–88.

Pacala, S. W. and Rees, M. (1998). Models suggesting field experiments to test two hypotheses explaining successional diversity. *The American Naturalist*, 152(5):729–37.

Pacala, S. W. and Tilman, D. (1994). Limiting Similarity in Mechanistic and Spatial Models of Plant Competition in Heterogeneous Environments. *The American Naturalist*, 143(2):222–257.

Paine, C. E. T., Baraloto, C., Chave, J., and Hérault, B. (2011). Functional traits of individual trees reveal ecological constraints on community assembly in tropical rain forests. *Oikos*, 120(5):720–727.

Pigolotti, S. and Cencini, M. (2013). Species abundances and lifetimes: From neutral to niche-stabilized communities. *Journal of Theoretical Biology*, 338:1–8.

Pigolotti, S., Lopez, C., Hernández-García, E., and H., K. A. (2010). How Gaussian competition leads to lumpy or uniform species distributions. *Theoretical Ecology*, 3(2):89–96.

Pigot, A. L. and Etienne, R. S. (2015). A new dynamic null model for phylogenetic community structure. *Ecology Letters*, 18:153–163.

Pillar, V. D., Duarte, L. D. S., Sosinski, E. E., and Joner, F. (2009). Discriminating trait-convergence and trait-divergence assembly patterns in ecological community gradients. *Journal of Vegetation Science*, 20(2):334–348.

Poorter, A. L., Wright, S. J., Paz, H., Ackerly, D. D., Condit, R., Harms, E., Licona, J. C., Mazer, S. J., Webb, C. O., and Wright, I. J. (2008). Are Functional Traits Good Predictors of Demographic Rates ? Evidence from Five Neotropical Forests. *Ecology*, 89(7):1908–1920.

R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rosenzweig, M. L. (1978). Competitive speciation. *Biological Journal of the Linnean Society*, 10(3):275–289.

Rosenzweig, M. L. and MacArthur, R. H. (1963). Graphical representation and stability conditions of predator-prey interactions. *American Naturalist*, pages 209–223.

Rosindell, J., Hubbell, S. P., and Etienne, R. S. (2011). The unified neutral theory of biodiversity and biogeography at age ten. *Trends in ecology & evolution*, 26(7):340–8.

Roughgarden, J. (1974). Species packing and the competition function with illustrations from coral reef fish. *Theoretical population biology*, 5:163–186.

Roughgarden, J. (1979). Theory of population genetics and evoltionary ecology: an introduction. In *Theory of population genetics and evoltionary ecology: an introduction*, pages 534–536. Macmillan, New York.

Roughgarden, J. (1983). Competition and Theory in Community Ecology. *The American Naturalist*, 122(5):583–601.

Rudin, W. (1973). Functional analysis. 1973.

Rueffler, C., Van Dooren, T. J. M., and Metz, J. a. J. (2007). The interplay between behavior and morphology in the evolutionary dynamics of resource specialization. *The American naturalist*, 169(2):E34–52.

Schamp, B. S., Chau, J., and Aarssen, L. W. (2008). Dispersion of traits related to competitive ability in an old-field plant community. *Journal of Ecology*, 96:204–212.

Scheffer, M. and van Nes, E. H. (2006). Self-organized similarity, the evolutionary emergence of groups of similar species. *Proceedings of the National Academy of Sciences of the United States of America*, 103(16):6230–5.

Schleuter, D., Daufresne, M., Massol, F., and Argillier, C. (2010). A user's guide to functional diversity indices. *Ecological Monographs*, 80(3):469–484.

Schoener, T. W. (1974). Resource partitioning in ecological communities. *Science (New York, N.Y.)*, 185(4145):27–39.

Schwilk, D. W. and Ackerly, D. D. (2005). Limiting similarity and functional diversity along environmental gradients. *Ecology Letters*, 8(3):272–281.

Schwinning, S. and Weiner, J. (1998). Mechanisms determining the degree of size asymmetry in competition among plants. *Oecologia*, 113:447–455.

Seedorf, H., Griffin, N. W., Ridaura, V. K., Reyes, A., Cheng, J., Rey, F. E., Smith, M. I., Simon, G. M., Scheffrahn, R. H., Woebken, D., Spormann, A. M., Van Treuren, W., Ursell, L. K., Pirrung, M., Robbins-Pianka, A., Cantarel, B. L., Lombard, V., Henrissat, B., Knight, R., and Gordon, J. I. (2014). Bacteria from Diverse Habitats Colonize and Compete in the Mouse Gut. *Cell*, 159(2):253–266.

Segura, A. M., Calliari, D., Kruk, C., Conde, D., Bonilla, S., and Fort, H. (2011). Emergent neutrality drives phytoplankton species coexistence. *Proceedings. Biological sciences / The Royal Society*, 278(1716):2355–61.

Shipley, B. (2010). *From plant traits to vegetation structure: chance and selection in the assembly of ecological communities*. Cambridge University Press.

Shipley, B., Vile, D., and Garnier, E. (2006). From Plant Traits to Plant Communities: A Statistical Mechanistic Approach to Biodiversity. *Science*, 314(5800):812–814.

Siefert, A. (2012). Incorporating intraspecific variation in tests of trait-based community assembly. *Oecologia*, 170(3):767–75.

Siepielski, A. M., Hung, K.-L., Bein, E. E. B., and McPeek, M. a. (2010). Experimental evidence for neutral community dynamics governing an insect assemblage. *Ecology*, 91(3):847–57.

Siepielski, A. M. and McPeek, M. A. (2010). On the evidence for species coexistence: a critique of the coexistence program. *Ecology*, 91(11):3153–3164.

Silva, R. R. and Brandão, C. R. F. (2014). Ecosystem-wide morphological structure of leaf-litter ant communities along a tropical latitudinal gradient. *PLoS ONE*, 9(3).

Simberloff, D. and Boecklen, W. (1981). Santa Rosalia Reconsidered : Size Ratios and Competition. *Evolution*, 35(6):1206–1228.

Sterck, F., Markesteijn, L., Schieving, F., and Poorter, L. (2011). Functional traits determine trade-offs and niches in a tropical forest community. *Proceedings of the National Academy of Sciences of the United States of America*, 108(51):20627–32.

Strong Jr, D. R., Szyska, L. A., and Simberloff, D. (1979). Test of Community-Wide Character Displacement Against Null Hypotheses. *Evolution*, 33(3):897–913.

Stubbs, W. J. and Wilson, J. (2004). Evidence for limiting similarity in a sand dune community. *Journal of Ecology*, 92(4):557–567.

Swenson, N. G. and Enquist, B. J. (2009). Opposing Assembly Mechanisms in a Neotropical Dry Forest : Implications for Phylogenetic and Functional Community Ecology. *Ecology*, 90(8):2161–2170.

Swenson, N. G., Enquist, B. J., Thompson, J., and Zimmerman, J. K. (2007). The influence of spatial and size scale on phylogenetic relatedness in tropical forest communities. *Ecology*, 88(7):1770–80.

Szabo, P. and Meszéna, G. (2006). Limiting similarity revisited. *Oikos*, 112:612–619.

TerHorst, C. P., Miller, T. E., and Powell, E. (2010). When can competition for resources lead to ecological equivalence ? *Evolutionary Ecology Research*, 2010(12):843–854.

Thompson, K., Petchey, O. L., Askew, A. P., Dunnett, N. P., Beckerman, A. P., and Willis, A. J. (2010). Little evidence for limiting similarity in a long-term study of a roadside plant community. *Journal of Ecology*, 98(2):480–487.

Tibshirani, R., Walther, G., and Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63:411–423.

Tilman, D. (1982). *Resource competition and community structure*. Princeton University Press, monographs edition.

Tilman, D. (1994). Competition and Biodiversity in Spatially Structured Habitats. *Ecology*, 75(1):2–16.

Tilman, D. (2004). Niche tradeoffs, neutrality, and community structure: a stochastic theory of resource competition, invasion, and community assembly. *Proceedings of the National Academy of Sciences of the United States of America*, 101(30):10854–61.

Turelli, M. (1981). Niche Overlap and Invasion of Competitors in Random Environments I . Models without Demographic Stochasticity. *Theoretical Population Biology*, 56(198 1):1–56.

Ulrich, W. (2004). Species co-occurrences and neutral models: reassessing J. M. Diamond's assembly rules. *Oikos*, 107(3):603–609.

Uriarte, M., Swenson, N. G., Chazdon, R. L., Comita, L. S., John Kress, W., Erickson, D., Forero-Montaña, J., Zimmerman, J. K., and Thompson, J. (2010). Trait similarity, shared ancestry and the structure of neighbourhood interactions in a subtropical wet forest: implications for community assembly. *Ecology letters*, 13(12):1503–14.

van der Plas, F., Janzen, T., Ordonez, A., Fokkema, W., Reinders, J., Etienne, R. S., and Olff, H. (2015). A new modeling approach estimates the relative importance of different community assembly processes. *Ecology*, 96(6):1502–1515.

Vandermeer, J. H. (1975). Interspecific Competition: A New Approach to the Classical Theory. *Science*, 188(4185):253–255.

Vergnon, R., Dulvy, N. K., and Freckleton, R. P. (2009). Niches versus neutrality: uncovering the drivers of diversity in a species-rich community. *Ecology letters*, 12(10):1079–90.

Vergnon, R., van Nes, E. H., and Scheffer, M. (2012). Emergent neutrality leads to multimodal species abundance distributions. *Nature communications*, 3(may 2011):663.

Villéger, S., Mason, N. W. H., and Mouillot, D. (2008). New multidimensional functional diversity indices for a multifaceted framework in functional ecology. *Ecology*, 89(8):2290–2301.

Violle, C., Enquist, B. J., McGill, B. J., Jiang, L., Albert, C. H., Hulshof, C., Jung, V., and Messier, J. (2012). The return of the variance: intraspecific variability in community ecology. *Trends in ecology & evolution*, 27(4):244–52.

Violle, C., Navas, M.-L., Vile, D., Kazakou, E., Fortunel, C., Hummel, I., and Garnier, E. (2007). Let the concept of trait be functional! *Oikos*, 116(5):882–892.

Volkov, I., Banavar, J. R., Hubbell, S. P., and Maritan, A. (2003). Neutral theory and relative species abundance in ecology. *Nature*, 424(6952):1035–7.

Volterra, V. (1927). *Variazioni e fluttuazioni del numero d'individui in specie animali conviventi*. C. Ferrari.

Waples, R. S., Do, C., and Chopelet, J. (2011). Calculating Ne and Ne/N in age-structured populations: a hybrid Felsenstein-Hill approach. *Ecology*, 92(7):1513–22.

Weiher, E., Clarke, G. D. P., and Keddy, P. a. (1998). Community Assembly Rules, Morphological Dispersion, and the Coexistence of Plant Species. *Oikos*, 81(2):309.

Weiher, E. and Keddy, P. A. (1999). Relative abundance and evenness patterns along diversity biomass gradients. *Oikos*, 87(2):355–361.

Wilke, C. O. (2005). Quasispecies theory in the context of population genetics. *BMC Evolutionary Biology*, 5:44.

Wills, C., Condit, R., Foster, R. B., and Hubbell, S. P. (1997). Strong density- and diversity-related effects help to maintain tree species diversity in a neotropical forest. *Proceedings of the National Academy of Sciences of the United States of America*, 94(4):1252–1257.

Wilson, J. B. (1999). Assembly rules in plant communities. In Weiher, E. and Keddy, P., editors, *Ecological assembly rules: perspecitves, advances, retreats*, pages 130–164. Cambridge University Press.

Wilson, J. B. and Stubbs, W. J. (2012). Evidence for assembly rules: limiting similarity within a saltmarsh. *Journal of Ecology*, 100(1):210–221.

Wilson, M. K. and Briggs, C. (1955). The normal flora of the bovine rumen ii. quantitative bacteriological studies. *Journal of Applied Bacteriology*, 18(2):294–306.

Wright, I. J., Reich, P. B., Westoby, M., Ackerly, D. D., Baruch, Z., Bongers, F., Cavender-Bares, J., Chapin, T., Cornelissen, J. H. C., Diemer, M., Flexas, J., Garnier, E., Groom, P. K., Gulias, J., Hikosaka, K., Lamont, B. B., Lee, T., Lee, W., Lusk, C., Midgley, J. J., Navas, M.-L., Niinemets, U., Oleksyn, J., Osada, N., Poorter, H., Poot, P., Prior, L., Pyankov, V. I., Roumet, C., Thomas, S. C., Tjoelker, M. G., Veneklaas, E. J., and Villar, R. (2004). The worldwide leaf economics spectrum. *Nature*, 428(6985):821–7.

Wright, S. J., Kitajima, K., Kraft, N. J. B., Reich, P. B., Wright, I. J., Bunker, D. E., Condit, R., Dalling, J. W., Davies, S. J., Díaz, S., Engelbrecht, B. M. J., Harms, K. E., Hubbell, S. P., Marks, C. O., Ruiz-Jaen, M. C., Salvador, C. M., and Zanne, A. E. (2010). Functional traits and the growth-mortality trade-off in tropical trees. *Ecology*, 91(12):3664–74.

Xiao, X., O'Dwyer, J. P., and White, E. P. (2016). Comparing process-based and constraint-based approaches for modeling macroecological patterns. *Ecology*, 97(5):1228–1238.

Yachi, S. and Loreau, M. (1999). Biodiversity and ecosystem productivity in a fluctuating environment: the insurance hypothesis. *Proceedings of the National Academy of Sciences of the United States of America*, 96(4):1463–1468.

Yan, B., Zhang, J., Liu, Y., Li, Z., Huang, X., Yang, W., and Prinzing, A. (2012). Trait assembly of woody plants in communities across sub-alpine gradients: Identifying the role of limiting similarity. *Journal of Vegetation Science*, 23(4):698–708.

Yoshiyama, R. M. and Roughgarden, J. (1977). Species Packing in Two Dimensions. *The American Naturalist*, 111(977):107–121.

Yul, D. W. and Wilson, H. B. (2001). The Competition-Colonization Is Dead ; Long Live the Competition-Colonization. *The American Naturalist*, 158(1):49–63.

Zhou, S.-R. and Zhang, D.-Y. (2008). A Nearly Neutral Model of Biodiversity. *Ecology*, 89(1):248–258.