

CHAPTER TEN

Reconciling Dispersal-Assembly and Niche-Assembly Theories

The principal aim of this book has been to show that a powerful, formal neutral theory of biodiversity and biogeography can be constructed on the foundation of the original theory of island biogeography. This theory describes the dynamics of species richness and relative species abundance in local communities and metacommunities that are undergoing ecological drift, random dispersal, and random speciation. The theory's key assumptions are, first, that the dynamics of ecological communities are a stochastic, zero-sum game, and second, that this game is neutral and played by individuals that are identical in their probabilities of giving birth, dying, and dispersing—and even of speciating. In chapter 1, I explored the philosophical divide between the dispersal assembly and niche assembly perspectives of ecological communities. In this concluding chapter, I anticipate a truly unified theory that at a more fundamental level reconciles these two apparently conflicting perspectives.

My prediction is that this reconciliation will resemble the unification of population genetics, which successfully combined the theories of genetic drift and of selection. Such a reconciliation in ecology will depend on wider acceptance of the fact that ecological drift is a universally present process that must inevitably operate continuously in real ecological communities, just as genetic drift is always operating on gene frequencies in populations. The appropriate question therefore is not whether ecological drift exists, but under what

circumstances is it quantitatively important. I would argue that we really have no idea because we have not looked adequately nor in the right ways. Our half-century preoccupation in ecology with equilibrium theories of coexistence in niche-assembled communities (chapter 1), has led us away from engaging in a thorough examination of the importance of ecological drift, both empirically and theoretically.

I believe that community ecology will have to rethink completely the classical niche-assembly paradigm from first principles. Such a rethinking is important because there are an embarrassingly large number of empirically well-documented and widespread patterns in community ecology that remain almost totally unexplained by contemporary theoretical ecology. These include regularities in the patterns of relative species abundance, species-area relationships, and even phylogeny. The neutral theory outlined in this book does a better job at explaining these patterns than current niche assembly theory does. In fact, many nonobvious patterns predicted by the neutral theory have been substantiated empirically, so the neutral theory is unexpectedly rich in new, testable predictions. This is not to argue that all of these patterns are due to neutrality, but it does shift the burden of proof to those who would argue that niche assembly is necessary to explain these patterns. I firmly believe that a new synthetic theory of biodiversity and biogeography will be developed which integrates ecological drift and demographic and environmental stochasticity into niche assembly theory, but this will require a deep reevaluation of the current niche-assembly theoretical paradigm. There is abundant evidence that such a fundamental rethinking has already begun. A number of ecologists have expressed ideas that are similar but not identical to those presented here (e.g., Levins 1968, Tilman 1988, Charnov 1993).

One of the great surprises to me in developing this theory has been just how well, in fact, it *does* work. It works astonishingly well in spite of making what might appear to be a

RECONCILING THE THEORIES

false and crippling assumption, namely, that all individuals are identical. As I will argue below, however, this assumption is far closer to the truth than might at first be appreciated. Why does the neutral theory do so well? In this chapter, I present my thoughts on the answer to this question, and its implications for the reconciliation of niche-assembly and dispersal-assembly perspectives.

The success of this neutral theory may be distressing to some people, but it may be of some comfort to note that the neutrality assumptions in this neutral theory are far more biologically interesting than simply asserting that "nothing is going on." Indeed, there are many possible neutral models that could be studied that would allow far more complexity in the ecological interactions among individuals than those studied in this book (chapter 1). The essential point to bear in mind is that assuming ecological equivalence among individuals is a very permissive definition of neutrality that potentially allows a great deal of interesting biology to be incorporated in neutral theories. In what follows I discuss the reasons why I think assuming ecological equivalence among individuals in trophically defined communities may not be such an unreasonable approximation of nature after all.

I believe that the success of the neutral theory is powerful testimony to the fact that, in some regards, ecological nature is far simpler than we might at first suppose. If one always had to specify the unique dynamical behavior of each and every species and all their interactions with resources and each other in a community, then the neutral theory simply would not work at all. The very fact that the neutral theory does so well implies that the dynamics of ecological communities are governed by rules that have much lower dimensionality than their species richness might suggest.

The key to reconciling niche assembly and dispersal assembly, I believe, lies in recognizing an initially counter-intuitive truth. It is not immediately obvious why, but niche

differentiation actually *greatly reduces* the dimensionality of ecological communities. In principle, there is no necessary reason to expect this. However, the niche differentiation of virtually all species in trophically well defined communities is of a special type, confined to a very restricted region of possible life history space. Virtually all niche differentiation is constrained to obey a small set of scaling laws and fitness invariance rules that embody a series of life history trade-offs shared by all species in the community. These trade-offs and fitness invariance rules can be fully specified by many fewer parameters than would be required to uniquely characterize the niche of each and every species.

By *fitness invariance* I mean that there are different trade-off combinations of life-history traits that confer equivalent per capita relative fitnesses on the species exhibiting them. This must be true by an almost self-evident proof. All species that manage to persist in a community for long periods with other species must exhibit net long-term population growth rates of nearly zero. A given species may win a little here, and lose a little there, but over the long term and over large landscapes, its net growth rate must be very close to zero. If this were not the case, i.e., if some species should manage to achieve a positive growth rate for a considerable length of time, then from our first principle of the biotic saturation of landscapes, it must eventually drive other species from the community. But if all species have the same net population growth rate of zero on local to regional scales, then ipso facto they must have identical or nearly identical per capita relative fitnesses.

Consider the tropical forest on Barro Colorado Island (BCI), Panama. In the 50 ha plot there are more than three hundred coexisting tree and shrub species with stems greater than 1 cm in diameter at breast height (dbh). Although the BCI forest is species rich, closer examination shows that virtually all BCI tree species are niche differentiated along a few major life history trade-offs. The best

RECONCILING THE THEORIES

known of these is the trade-off between shade tolerance and growth rate. Shade-tolerant trees have high tolerance (high survival) to shade stress but are unable to grow rapidly in full sun. Conversely, shade-intolerant pioneers are able to grow very rapidly in the high light environment of a large treefall gap, but they have very poor tolerance of shade (low survival). This life history trade-off can be visualized graphically by plotting mean annual survival rate of a tree species in the shade against median annual growth rate in full sun. The trade-off is manifest as a narrow locus of points in niche space along which all BCI tree species are arrayed (fig. 10.1). Most of the tree species in the BCI forest are shade tolerant, and so there are more species at the shade-adapted end of the manifold than at the sun-adapted end. The physiological basis for this trade-off is well known (e.g., Bazzaz and Pickett 1980, Coley et al. 1985, Reich et al. 1999). This is an example of what has been called a *life history manifold* because it is an evolutionary attractor in niche space. It represents a series of unavoidable or nearly unavoidable life history design or constraint functions that every species in the trophically defined community must obey in order to stay in the game. Clearly, this manifold greatly reduces the dimensionality problem of characterizing life-history variation in the BCI tree community. The manifold can be described by many fewer parameters than could each and every tree species individually. With two parameters, one can characterize the principal linear axis, and with a few more parameters, one can describe even the distribution of species densities along the manifold.

Moreover, niche differentiation in the BCI tree community is further simplified because a majority of the other major life history traits of BCI tree species also covary predictably along the shade tolerance-intolerance manifold. For example, shade-tolerant species in general produce fewer, larger seeds that disperse relatively locally and germinate quickly in understory shade. Shade-intolerant pioneer species tend to produce large numbers of small seeds that




FIG. 10.1. Niche differentiation along the shade tolerance–shade intolerance life history manifold for BCI tree species. Each dot represents a species mean for stems 1–4 cm diameter at breast height (dbh). The x -axis is the average relative growth rate (percent diameter increase per year) of a species in full sun (in large light gaps). The y -axis is the average annual percentage survival of the same species in understory shade (average light level 1–2% of full sun). Most BCI tree species are shade adapted.

disperse widely but remain dormant in the soil for several years, waiting for a light cue to germinate (Dalling et al. 1997, Vasques-Yanes and Smith 1982). Shade-tolerant species have longer-lived leaves that are more invested in chemical defenses against pathogen and herbivore attack than are the shorter-lived leaves of pioneers (Coley et al. 1985). Indeed, Peter Reich and his colleagues have described a whole suite of tightly covarying physiological and life-historical traits connected to leaf longevity, which they have been developing into a general theory of life-history evolution in plants (Reich et al. 1997, 1998, 1999).

RECONCILING THE THEORIES

The irony of this argument in the context of the present neutral theory is that it turns out that *niche differentiation along life history trade-offs is the very mechanism by which per capita relative fitnesses are equalized among the coexisting species in a community*. In light of this, the question that we started with, namely, "Which is more important, niche assembly or dispersal assembly?" now appears somewhat simplistic because these two theoretical perspectives are fundamentally and intrinsically intertwined and interdependent mechanistically. This logic leads to a set of perhaps deeper and subtler questions.

If per capita relative fitnesses are equalized by life-history invariance rules, then why don't "cheater" species invade the communities that break these rules by escaping the life-history trade-offs? The answer is that there probably *are* cheaters from time to time that manage to break the rules. Most of the time, however, the microevolutionary change we observe in non-rule-breaking species is constrained to take place along preexisting life-history trade-offs. Thus, the mean phenotype of a species may slide up or down along the current community life-history manifold, all the while maintaining fitness invariance. The current position of the species on the life-history manifold will reflect its history of selection in past environments. For example, a tree species might gradually evolve higher growth rates as a result of a history of exposure to higher light environments, but in so doing it would be expected to lose the correlated traits for shade tolerance. Fitness invariance arising from adaptive trade-offs is not a new argument, which was perhaps earliest and best expressed by Richard Levins (1968) in his now classic book, *Evolution in Changing Environments*. Among other things, Levins argued that adaptive trade-offs are inevitable, and escaping them is very difficult, in part because organisms are limited by the amount of genetic information they can carry (preventing the evolution of superorganisms that are best in all environments), and in part because of logical

CHAPTER TEN

inconsistencies (e.g., being simultaneously large and small, many seeded and few seeded), historical-morphogenetical constraints, and ultimately, boundary conditions on the possible imposed by physico-chemical laws.

Every now and then, however, a species does manage to break partially free of the constraints of the life-history manifold currently governing its community, and this species will achieve a new level of fitness with a somewhat different set of constraints. For a time, such a species will be a superior competitor, sweeping communities free of, or at least reducing, populations of its competitors, depending on the magnitude of differences in relative fitness. Introduced exotic species that are successful invaders of biotically saturated communities are often such rule breakers in their new community settings. However, this state of affairs cannot go on for long, ecologically or evolutionarily. The existence of rule breakers establishes a new, higher fitness criterion for all other species in the community, which then come under strong selection for life-history adjustments that increase their fitness to match the rule breaker. Presumably, minor evolutionary rule breaking goes on all the time but is largely undetected. However, large breaks in the rules that are also successful are presumably increasingly infrequent the larger the break in the rules is. This is analogous to the argument behind David Raup's (1991) "kill curve" for the distribution of sizes of extinction events, but in reverse. Thus, it is only when a massive rule-breaking episode takes place in the fossil record that we identify and label it as an "adaptive radiation."

In summary, although there is no necessary reason why niche differentiation in principle should simplify community dynamics, it usually does. But why? After all, niche differentiation makes species *different*, doesn't it? I believe the answer to this question is both *yes* and *no*. The answer is *yes* in the sense that species do exhibit different life-history traits that allow, among other things, taxonomists and ecologists to distinguish one species from another. But the answer is *no* in

RECONCILING THE THEORIES

the sense that these niche differences obey life-history trade-off rules that maintain per capita fitness equivalence among the niche-differentiated species.

Niche differences follow life-history invariance rules that maintain equal relative fitnesses, but this only takes us halfway to reconciling the niche assembly with dispersal-assembly perspectives. Now we must ask: What determines the species richness and relative species abundance of the species in the community? According to the neutral theory, mostly what mechanisms remain once per capita relative fitnesses are equalized are processes of ecological drift. This implies that *life-history trade-offs and fitness invariance rules potentially decouple niche differentiation from control of the species richness and relative species abundance of communities.*

Once species' fitnesses are equalized, the degree of decoupling of niche differentiation from relative abundance will depend on the degree to which zero-sum dynamics apply. I argued in chapter 3 that the zero-sum rule must govern the dynamics of many communities of trophically similar species. I have given many examples in this book of patterns of relative species abundance that are qualitatively and quantitatively consistent with zero-sum ecological drift in a diverse array of communities. But the question remains: How generally do zero-sum dynamics apply among species on the same trophic level? The neutral theory makes the simplifying assumption that any species in a community of trophically similar species has the potential to become monodominant, occupying all available space or using all limiting resources, if no other competing species are present. In other words, no limiting resource will go unused for long if we vary the species richness and composition of a community. Thus, the zero-sum rule assumes that resources limiting the community are available to all member species of the community. This assumption will not always be true, or it may be true only on evolutionary timescales. However, once again I argue that we really have no idea because we have not

been asking the right questions. The phenomenon of "ecological release" noted on islands (e.g., Crowell 1962, 1973) may actually be far more prevalent in natural communities than we acknowledge. Indeed, it may be going on all around us all the time on a micro scale.

I believe that species' adaptability and malleability in resource use are why the assumption of zero-sum dynamics works so well in so many trophically defined communities. In many cases, any species in the community can use any of the resources that limit the total community. Resources normally used by one species are often opportunistically used by another species whenever the former species happens to be absent. Species may use some resources slightly more or less efficiently, but these differences are unlikely to seriously violate the zero-sum rule.

Opportunism probably happens far more often than current niche theory would lead us to believe. There is a pervasive tendency in ecology to treat all individuals of a given species as exhibiting the mean species phenotype, i.e., to treat species typologically as invariant entities. We routinely overlook the flexibility, opportunism, and facultative use of resources of which variable individuals in species are universally capable—not to mention potential evolutionary responses. I believe that most species will be found to be far less specialized in resource use than current theory suggests. Virtually all species have at least some capability of switching resources if their preferred resources are unexpectedly scarce. I think that the rigidity of our current niche typology has blinded us to not seeing the pervasive facultative and opportunistic use of resources by species. I therefore predict that zero-sum dynamics will apply far more generally among species on a given trophic level than one would predict based on current theory. This rigid niche typology is also what has led us to the prediction that competitive exclusion will be commonplace when in fact it is rarely if ever observed in nature.

RECONCILING THE THEORIES

Thinking typologically about species also dramatically affects our ideas about speciation and the meaning of biodiversity itself. We have seen that the unified theory generally fits patterns of biodiversity best under the assumption that species are fractal entities. According to this view, species and higher taxa are vessels of self-similar phenotypic and genotypic variation on biological scales running from the individual level all the way up to the metacommunity level of organization (chapter 8). The "point mutation" mode of speciation is more consistent with the data and with a fractal view of species, and this is the theory that generates the fundamental biodiversity number θ . The "random fission" mode of speciation arises from the classical model of allopatric speciation (Mayr 1963), but at this preliminary point in time, this mode appears to be less consistent with the data (chapter 8). If we accept a fractal view of species and of the variation that they encompass, then understanding the true meaning of biodiversity is a different and more challenging problem than we might have thought.

In order to test the unified neutral theory, therefore, we need much more research on the extent to which species in tropically well-defined communities obey the zero-sum rule, as well as research on the mechanisms of variable resource use by variable individuals within and among variable species. Such questions were once pursued vigorously more than two decades ago in studies of within- and between-phenotype niche breadths (e.g., Roughgarden 1972, Colwell and Futuyma 1971), but these are still fundamental, unresolved questions on which we need to reinvigorate research.

One of the ways that the zero-sum rule breaks down is in the case of habitat specialization. In developing the present neutral theory, I assumed that there was a completely homogeneous habitat across the metacommunity landscape, and I asked what patterns of biodiversity and biogeography would develop under ecological drift, random dispersal, and

random speciation. However, in nature real landscapes consist of environmental gradients, and very different habitats exist in a landscape mosaic. One can develop alternative theories for the maintenance of biodiversity on heterogeneous landscapes (Tilman 1988, Tilman and Pacala 1993, Ritchie and Olff 1999). Moreover, the existence of environmental gradients that affect relative fitness of species also lead to important complexities such as source and sink populations (Pulliam 1988). All of these ideas are worthy of much more theoretical development, particularly in conjunction with the theory of ecological drift, but these developments lie in the future, and are beyond the scope of the present work.

I would like to conclude this book with a discussion and response to a number of critiques that have appeared of early versions of the theory. I will respond in detail to only two of these—one empirical, and one theoretical—because my responses make some general points that are germane to the question of how to reconcile niche assembly and dispersal assembly theories.

Terborgh et al. (1996) analyzed several plots of mature tropical forest, each 1–2.5 ha, along a 40 km stretch of the Manu River in Amazonian Peru. They argued that the abundance rankings of the top twenty commonest species were too similar from plot to plot and over such a distance to be consistent with ecological drift (which they called the “nonequilibrium theory”). Pandolfi (1996) (also see Jackson et al. 1996 and Pandolfi and Minchin 1995) analyzed a continuous chronosequence of fossil reef terraces gradually uplifted along the north coast of Papua New Guinea. The record extends over a period of 95,000 years through repeated sea level and surface temperature changes, and also extends spatially for tens of kilometers along the coast. Pandolfi found relative constancy in both taxonomic composition and in species richness, punctuated by intervals of relatively rapid change, over this period of nearly 100,000 years. They attributed the long periods of relative stability to

RECONCILING THE THEORIES

the existence of limited-membership reef communities that were noninvasible (i.e., niche assembled).

The conclusions of Terborgh et al. (1996) and Pandolfi (1996) are premature, according to the unified theory, for a number of reasons. First, we know from the theory presented in previous chapters that common and widespread metacommunity species are likely to be very resistant to extinction and to persist for geologically significant periods of time. Second, in chapters 4 and 7 we showed that very modest rates of dispersal will insure that these common metacommunity species are nearly everywhere all the time. Third, increasing the rate of dispersal will increase the proportion of metacommunity diversity that is present in local communities.

Attempting to test ecological drift theory using only common metacommunity species is particularly problematic. For one thing, there is a positive correlation between the total abundance of species and their geographic ranges (Brown 1984, 1995). Limiting one's attention to the most abundant species is to bias the analysis in favor of those species that are least likely to be dispersal limited and most likely to be persistent in space and time. Therefore, one's conclusions are strongly biased toward constancy of community composition. We have seen that the time to extinction $T(N_i)$ of a very common metacommunity species is expected to be extremely long (chapters 4, 8). From θ we know that J_M is an enormous number—on the order of the inverse of the speciation rate v . But the time to extinction $T(N_i)$ for large N_i is an even bigger number, on the order of the metacommunity size times the log of metacommunity size, times the initial population size of the species (chapter 4). The immensity of this number is important because it means that common species will be very persistent members of the metacommunity and of local communities as well. This number is so large that it means that the ecological dynamics of metacommunities are temporally commensurate with

the evolutionary dynamics of speciation and extinction. Indeed, this has to be so because we have proved that a steady-state metacommunity biodiversity is established at equilibrium between speciation and extinction. Paleobiological evidence that widespread species are more resistant to extinction events has been provided by Jablonski (1995) and Jackson (1995).

These long persistence times mean that very abundant metacommunity species have plenty of time to disperse nearly everywhere. Community stability should therefore be increasingly conspicuous as one looks at aggregate species abundances at higher and higher taxonomic levels. This effect was noted by Gentry (1982) in a large number of 0.1 ha forest plots inventoried all over tropical South America. Gentry found that the familial composition of tropical forests is extremely constant, which is exactly what the unified theory would predict. This stability is shown in figure 10.2, which presents the family level composition of the BCI forest, for all of Panama, and for the world. The cumulative distributions of families are nearly identical from local to global scales. Results such as these can also be taken as support for the niche assembly perspective. If the early radiation of the angiosperms produced a rapid filling of niches at the family level, and if these niches are ubiquitous in plant communities, then this result would be expected. Thus, the problem is that both theories of community organization—dispersal assembly and niche assembly—predict the same results, and they certainly do not falsify ecological drift.

In chapter 7 I discussed incidence functions, which give the probability that a species will be present in a local community or habitat patch as a function of patch size J and the probability of immigration per birth m . Recall the relatively conservative case of a species that is only moderately common in the metacommunity, constituting just one percent of it (fig. 7.2). In spite of relative rarity in the metacommunity

RECONCILING THE THEORIES

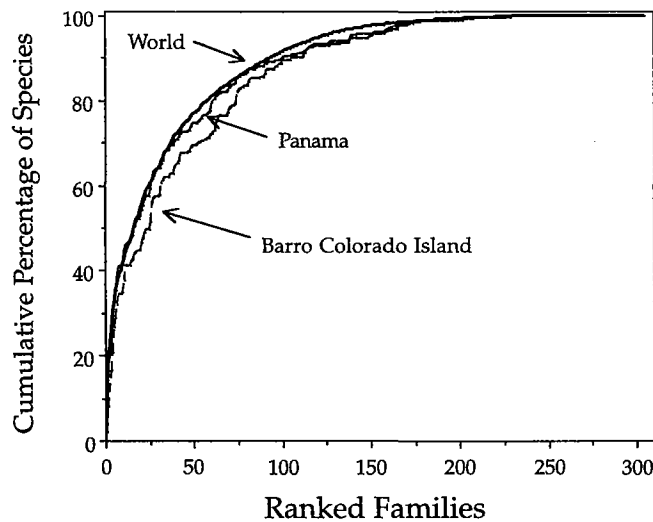


FIG. 10.2. High degree of constancy in the family-level composition of floras from local scales (BCI) to global spatial scales. Such a result is predicted by both the unified theory and by niche assembly theory, and therefore is not useful in testing them.

as a whole, this species will still be essentially always present (100%) in a local community of 10^3 individuals or more if $m \geq 0.1$, in communities of 10^4 individuals or more if $m \geq 0.01$, and in communities of 10^5 individuals or more if $m \geq 0.001$. In this context, let us reconsider the Terborgh example. Suppose that in round numbers there are 10^4 trees >10 cm in trunk diameter at breast height in a 25 ha plot of Manu forest (estimate based on BCI). Now let us conservatively imagine a corridor of forest 40 km long consisting of 1 km^2 plots along the Manu River, each containing approximately $4 \cdot 10^4$ trees. Therefore, the corridor contains about $1.6 \cdot 10^6$ trees. On BCI we estimated the immigration rate for the 0.5 km^2 plot to be close to 0.1. Suppose that for a square kilometer the value of m is half that value, or 0.05. For the moment, assume that species must migrate only within this corridor. This is a more restrictive case than reality because the forest extends in all dimensions for much greater distances. Let us now simulate the Manu metacommunity and measure the predicted presence of the twenty

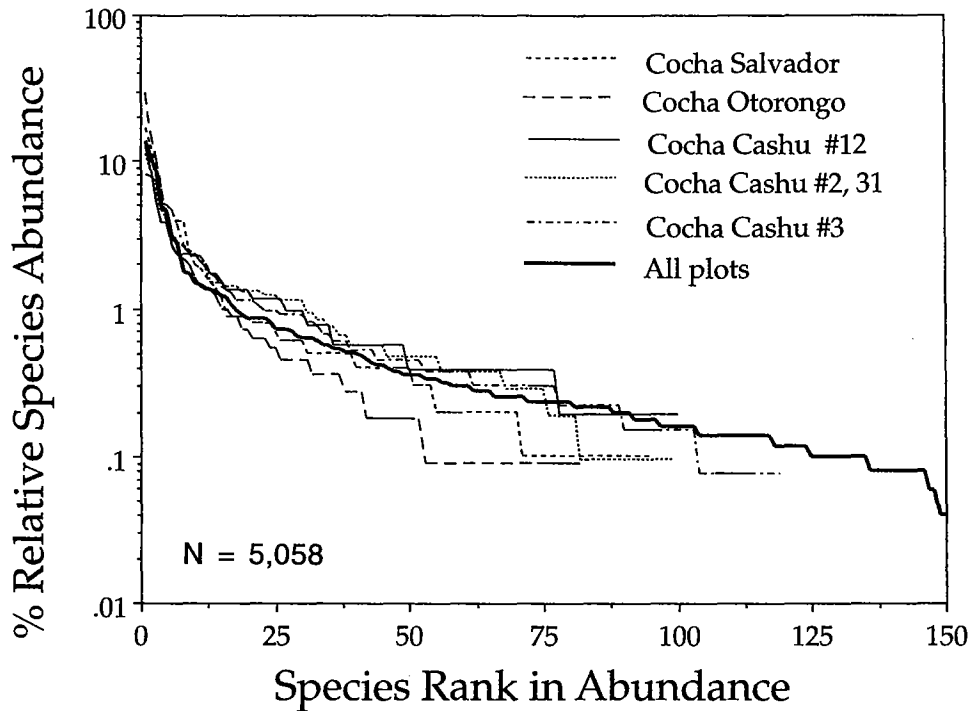


FIG. 10.3. Dominance-diversity curves for five plots of 1 to 2.5 ha in mature forest along the Manu River in Amazonian Peru. Three of the plots are nearly contiguous with one another (the plots at Cocha Cashu) and separated by about 40 km from the other two sites (Cocha Salvador and Cocha Otorongo), which were close together. The "all plots" is a composite of all five plots. These data represent all species with a total abundance of four individuals or more. Unpublished data of John Terborgh.

most abundant species in local community samples of one thousand trees.

Using unpublished raw data on Manu tree abundances provided by John Terborgh, I plotted the dominance-diversity curves for the five plots discussed in Terborgh et al. (1996) (fig. 10.3). Three of the plots were nearly contiguous at a site named Cocha Cashu and were approximately 40 km from the other two sites (Cocha Salvador and Cocha Otorongo), which were near to each other. Because of potential problems with taxonomy, I excluded from the analysis any species with fewer than four individuals in a plot, with one exception. One of the plots was less than half the size of the others (trail 12), and in this case I included all

RECONCILING THE THEORIES

species with two or more individuals (to attempt to equalize sampling intensity). In chapter 6, I demonstrated that pooling a series of small samples taken across the metacommunity should help compensate for dispersal limitation and give a better estimate of the metacommunity relative abundance distribution. Therefore, I pooled all five plots for an estimate of the species richness and relative species abundance distribution of the Manu metacommunity. Four of the five sites have steeper dominance-diversity curves than the metacommunity aggregate (fig. 10.3), as predicted by dispersal limitation in the unified theory, and of course fewer species.

The simulations show that the top sixteen species are expected to be present in all local stands essentially 100% of the time (fig. 10.4). The Manu forest is heavily dominated by two palm species, *Iriartea deltooides* and *Astrocaryum macrocalyx*, which together make up a quarter (24.5%) of all trees in the forest. The top twenty species in abundance collectively make up two-thirds (65.1%) of the forest. Also as expected from the theory, the average rank abundance of the species was the same in the local stands as in the metacommunity, although the actual rank positions in individual stands varied. The variation in rank position among local stands increased for the higher-ranked, rarer species, which is exactly what Terborgh et al. (1996) observed.

Ironically, Terborgh et al. (1996) may have been correct about the importance of niche assembly in the Manu forests for at least one reason that they did not give. They were correct in noticing that the dominants were everywhere and nearly equally dominant in rank position in all plots. But this observation alone does not contradict ecological drift, as we have just pointed out. What they did not comment on, however, was the excessive dominance itself—presumably because they had no prior statistical hypothesis of what null relative abundance distribution to expect. The unified theory detects significant levels of dominance in the seven most

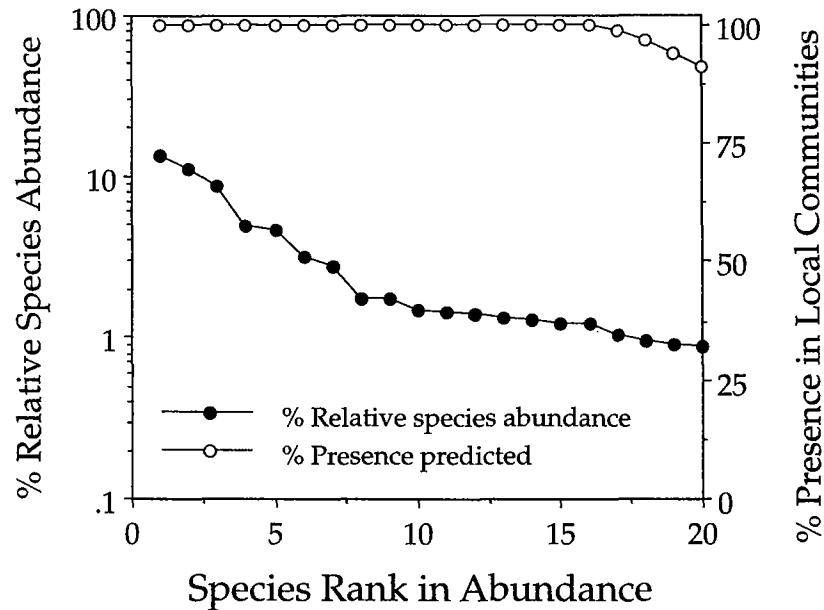


FIG. 10.4. The percentage relative species abundance of the twenty most abundant species in five mature forest plots along the Manu River in Amazonian Peru, and the percentage of plots of comparable size in which presence of the ranked species is predicted. The closed circles are the observed percentage relative species abundances. The open circles are the predicted percentages of plots that will contain the species. Unpublished data provided by John Terborgh.

abundant species over that predicted by the best-fit metacommunity relative abundance distribution. Note that there is a kink and sudden upward bend in the distribution at a rank abundance of eight (fig. 10.4). This implies that the top seven species enjoy some relative fitness advantage over the remaining species in the metacommunity. This advantage might be expressed as a reduced per capita probability of death or a higher per capita probability of birth, or in some combination of the two.

The exact biological reasons for the excessive dominance in the top seven most abundant species in the Manu forest are currently unknown, so it is not presently possible to model them. However, I can illustrate how easily in principle one can modify the stochastic equations of zero-sum ecological drift to incorporate a fitness advantage for given species. Suppose, for example, we assign each of the top

RECONCILING THE THEORIES

seven species a slight competitive advantage by increasing its per capita survival rate over other species. Under pure drift, the probability of a death in the i th species is N_i/J . Now let s_i be the survival advantage of the i th species. Then the probability that the next death in the community will occur to species i is simply $N_i(1 - s_i)/J$. Note the obvious parallels with the way in which selection is incorporated in population genetics models. One can fit the s_i by a combination of simulation and maximum likelihood procedures.

I have done this for the Manu forest example, and the results are presented in figure 10.5. Several conclusions from this case study are interesting and of general significance. The most important general result is that *incorporating modest fitness differences does not result in the rapid competitive exclusion and extinction of disadvantaged species in the metacommunity*. This result is not predicted by classical theory, which instead predicts that competitive exclusion will take place quickly, even with relatively slight competitive differences (e.g., Zhang and Lin 1997). This difference from classical theory is due to dispersal limitation, which almost indefinitely delays competitive exclusion (Tilman 1994, Hurtt and Pacala 1995). I will explain this result later. It has important general implications for community ecology, as well as for the ultimate reconciliation of dispersal assembly and niche assembly theory.

The second conclusion, which is also general, is that *relatively small fitness differences can nevertheless produce large differences in steady-state relative species abundance*. In the Manu case, for example, I calculate that the commonest species has a fitness advantage of only about 6%, but this small advantage resulted in almost a doubling of the relative abundance of the rank-1 species over what was expected under pure ecological drift (fig. 10.5). The remaining six species with elevated abundances have even smaller relative fitness advantages. In general, this and other case studies indicate that the unified theory is capable of detecting small differences in fitness.

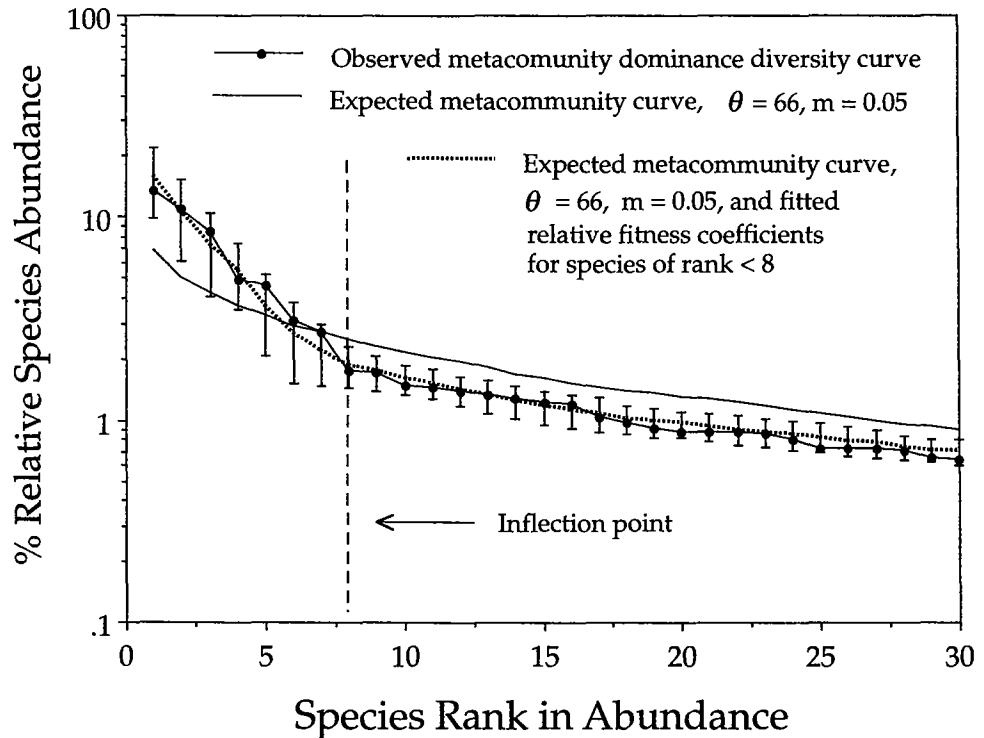


FIG. 10.5. Dominance-diversity curve for trees in old-growth forest along the Manu River in Manu National Park, Peru. The thin line without symbols is the expected metacommunity distribution under pure ecological drift. The dotted line is the expected metacommunity distribution with survival advantages for the seven most abundant species. The observed dominance diversity is the curve with filled-circle symbols. Only very slight survival advantages were required to fit the curve. The largest advantage was just 6% for the rank-1 species, which resulted in a nearly twofold increase in steady state abundance over the abundance expected under pure zero-sum ecological drift. The error bars are ± 1 standard deviation of the observed mean.

I suggest using the term *ecological dominance deviation* for cases in which dominance in ecological communities statistically exceeds that predicted under ecological drift.

The third conclusion is an empirical generalization that so far has held in most communities I have examined, and is true in the Manu case, but its acceptance awaits much further testing. I tentatively conclude that *ecological dominance deviations, if present at all, are usually detectable in only a small fraction of the species in a community*. In the Manu forest case, for example, the relative abundance distribution of the

species remaining after the most abundant seven species are taken out is statistically indistinguishable from that expected under pure community drift. It is as if the remainder of the metacommunity distribution is collectively downshifted to lower relative abundances. Note that the observed curve is almost perfectly parallel to the expected metacommunity curve under pure drift for the species of rank-8 and higher (fig. 10.4). This suggests that ecologists should explore mixed models of ecological communities in which a few species are competitive dominants that sequester most of the resources under the zero-sum game, and that the remaining species in the community are gleaners that compete for the leftovers under a more neutral zero-sum game. These ideas are reminiscent of the niche preemption theories of Motomura (1932) and Whittaker (1965). Once again, the competitive dominants cannot take over completely, according to the unified theory, because of steady-state dispersal limitation.

The final conclusion is really a corollary of the second and third conclusions, but it bears emphasizing. Even when ecological dominance deviations are detectable, the assumption of per capita equivalence among all species is rarely if ever quantitatively far from the truth. Thus, even when some species in the community have detectable competitive advantages, nevertheless these advantages are often quite small and therefore often do not seriously compromise the assumptions of neutrality or the zero-sum game.

Before leaving the subject of detecting ecological dominance deviations, I offer one theoretical example. Figure 10.6 shows the effect of giving the rank-1 species a 25% and a 50% survival advantage, respectively, over all other species in the community. I examined a case of a local community of size $J = 1600$, which receives immigrants with probability $m = 0.1$ per birth from a metacommunity with a fundamental biodiversity number $\theta = 25$. I assume that the dominant species is locally dominant only, so that




FIG. 10.6. The effect of ecological dominance on the theoretical distribution of relative species abundance in a local community of size $J = 1600$ under moderate dispersal limitation ($m = 0.1$). Ecological dominance in this case is due to a survival advantage in the rank-1 species of either 25% or 50%. Note that the curves for the higher-rank species remain parallel to the expected curve under pure drift. This is expected if all remaining species engage in a neutral zero-sum game.

its relative abundance in the metacommunity is as expected under pure drift. With no advantage, the expected local abundance of the rank-1 species with moderate dispersal limitation ($m = 0.1$) is a little over 10%. This percentage relative species abundance rises to 28% for a survival advantage of 25%, and to 71% for an advantage of 50% (fig. 10.6).

There are other kinds of significant ecological dominance deviations that might also be observed, not just in the most abundant species, but also in species of middle abundance rank or rare species. For example, in the previous chapter we illustrated two cases of such deviations, one for the bird community in 10 ha of tropical forest in Manu National Park (Terborgh et al. 1990; fig. 9.13), and the other for a rain-forest-canopy beetle community (Stork 1997; fig. 9.17). In the case of the Manu bird community, there

was a hump in the dominance-diversity curve at intermediate abundances, as well as excess dominance in the most abundant species. In the case of the canopy beetle community, there were too many species represented by singletons. In each case, we were able to detect significant departures from the neutral expectation.

I turn now to a theoretical paper that is critical of the theory of ecological drift. I will argue that understanding and answering the authors' objections goes a long way toward solving the larger problem of reconciling the dispersal assembly and niche assembly perspectives on the organization of ecological communities.

Zhang and Lin (1997) argued that the long coexistence times that are expected for identical species under zero-sum ecological drift (called "community drift" in Hubbell and Foster 1986b) is a fragile result. They introduced small differences in fecundity into the absorbing case of the ecological drift model, and found that relatively small differences in fecundity led to a dramatic shortening of the expected mean lifespans of species. They studied a two-species community in which the species were identical in all respects except for per capita fecundity. Let w be the per capita fecundity factor of the focal species relative to the nonfocal species. Thus, if $w = 1$, then the per capita fecundities of the two species are identical.

Now consider the modified transition probabilities for the focal species that take into account the differential per capita fecundity; note that the equations reduce to the absorbing case discussed in chapter 4 if $w = 1$.

$$\Pr\{N_i - 1|N_i\} = \left(\frac{N_i}{J}\right) \left(\frac{J - N_i}{(N_i - 1)w + J - 1}\right)$$

$$\Pr\{N_i|N_i\} = 1 - \Pr\{N_i - 1|N_i\} - \Pr\{N_i + 1|N_i\}$$

$$\Pr\{N_i + 1|N_i\} = \left(\frac{J - N_i}{J}\right) \left(\frac{N_i w}{J - 1 + N_i(w - 1)}\right).$$

To simplify notation, let $p_n = \Pr\{N_i + 1|N_i\}$ and $q_n = \Pr\{N_i - 1|N_i\}$. The mean time to extinction or complete dominance for the focal species, starting at abundance N_i , can be found analytically (Zhang and Lin 1997):

$$T(N_i) = \frac{\sum_{k=0}^{J-1} R_k}{\sum_{k=0}^{J-1} e_k} \sum_{k=0}^{N_i-1} e_k - \sum_{k=0}^{N_i-1} R_k \quad \text{for } 0 < N_i < J,$$

where

$$e_k = \frac{q_1 q_2 \cdots q_k}{p_1 p_2 \cdots p_k}, e_0 = 1,$$

$$R_k = \frac{1}{p_k} \left[1 + \frac{q_k}{p_{k-1}} + \cdots + \frac{q_k \cdots q_2}{p_{k-1} \cdots p_1} \right], \text{ and } R_0 = 0.$$

Zhang and Lin (1997) then studied how the time to extinction or complete dominance in a two-species community was affected by the differential fecundity factor w . There were rapid decreases in the time to “fixation” as w increased to about 1.2. However, for larger values of w , the rate of decline in species lifetimes was much slower (fig. 10.7). Zhang and Lin focused their attention on the fast decline in the time to fixation for small differences in relative fecundity, but they failed to note that the species lifetimes remained very large. More important, there are still huge increases in the time to fixation as community size J increases. They only computed lifetimes in small communities below 4096 individuals. As community size grows as large as a naturally occurring metacommunity, average species lifespans will become very long, even under the condition of unequal fecundities. This is one of the main conclusions of chapters 5 and 8, which show that the times to extinction will be so long that a diversity equilibrium will be established with the speciation rate.

Perhaps the most substantive criticism of Zhang and Lin’s (1997) result is that their model fails to take dispersal and dispersal limitation into account. In chapter 4 we also considered a model of local community dynamics in the absence of dispersal limitation. In that model, as in Zhang and Lin’s,

RECONCILING THE THEORIES

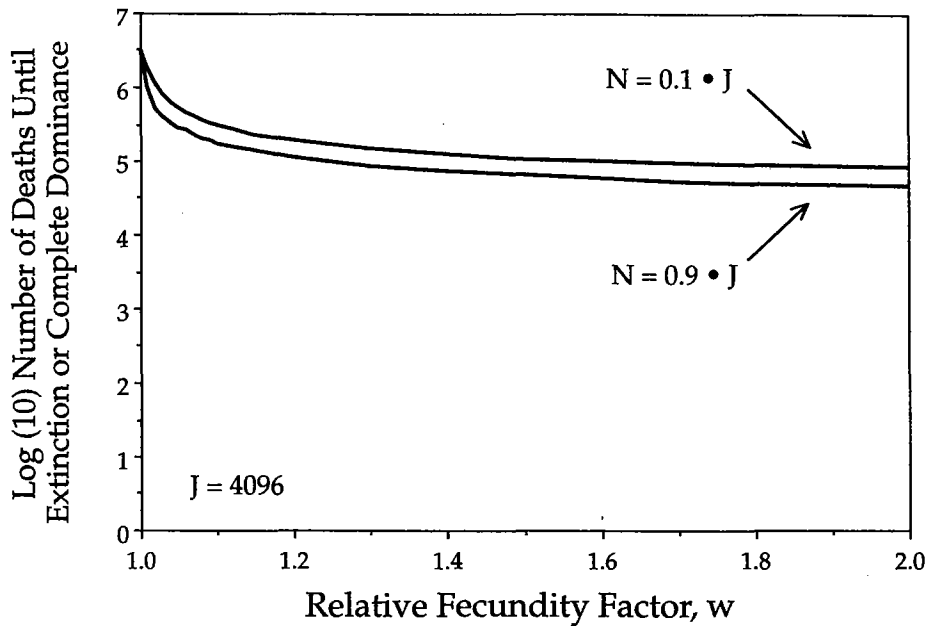


FIG. 10.7. Effect of competitive asymmetry (unequal fecundity) on the expected time to extinction or monodominance in the model community of Zhang and Lin (1997). Note the rapid decrease in mean lifespan with increasing difference in relative fecundity, especially for $w < 1.2$. Note also, however, that the rate of decline in lifespan slows down considerably for higher w . Two cases are drawn, one for a species having an initial abundance of 10% of the community, and one for a species having an initial abundance of 90% of the community. Redrawn from Zhang and Lin (1997).

each individual was equally likely to give birth to the replacement for any dying individual, regardless of whether the death occurred nearby in the same local community or far away in some other community in the metacommunity. However, chapters 6 and 7 of this book have largely been devoted to exploring the consequences of dispersal limitation under ecological drift for equilibrium species-area relationships and metacommunity biodiversity. According to the unified theory, these large-scale spatial patterns of biodiversity arise through an interaction of slow speciation rates, long extinction times of abundant and widespread metacommunity species, and steady-state limitation in the rate of dispersal of species over the metacommunity landscape.

The significance of dispersal limitation to the coexistence of species in ecological communities has been studied explicitly in two important theoretical papers by Tilman (1994) and Hurtt and Pacala (1995). Tilman (1994) showed that if there was a strict, ordered trade-off between competitive ability and dispersal ability, such that the best competitor was the worst disperser and vice versa, then in principle there is no upper limit to the number of coexisting species. Hurtt and Pacala (1996) went further and demonstrated that dispersal limitation all by itself can overwhelm the effects of even strong competitive differences, especially in species-rich communities. They analyzed a model of a community of niche-differentiated species partitioning an environmental axis. Each species was the absolutely best competitor in some region of the axis. Without dispersal limitation, each species as expected won those sites in which it was the best competitor. With dispersal limitation, however, many sites were won by "forfeit" by species that were not the absolutely best competitor for the site.

Hurtt and Pacala (1995) also demonstrated that a runaway process can occur between species richness and community-wide dispersal limitation. As communities become richer in species, the average population density of individual species falls, increasing the mean steady-state level of dispersal limitation in the community. This increases the number of forfeited sites and the number of species that win by default. The result is that the exclusion of inferior competitors becomes nearly indefinitely delayed. This runaway process is more likely to start if competitive differences among species are not large, but once it has begun, the effects of dispersal limitation can overwhelm even large competitive differences. The limits to community species richness and relative species abundance are then set, not by local niche differentiation, but by regional biogeographic processes—speciation, extinction, migration, and drift on the metacommunity landscape.

RECONCILING THE THEORIES

The true significance of these two papers is that they provide yet another mechanism that largely decouples niche differentiation from the control of species richness and relative species abundance in ecological communities. Among other things, these papers imply that a complete continuum is possible in the quantitative degree to which niche assembly versus dispersal assembly controls local community structure and relative species abundance.

In this book I hope that I have achieved my three original goals, which were, first, to explore how far we can get with a formal neutral theory in community ecology; second, to show that both ecological drift and biotic interactions are important to the assembly and dynamics of ecological communities; and, finally, to reduce the pervasive resistance of ecologists to neutral theory by demonstrating its considerable power and the richness of its testable predictions. Much theoretical and empirical work remains to be done, however, to achieve a truly synthetic theory of biodiversity and biogeography, and I look forward with considerable excitement to continued progress on both fronts.

At the end of their monograph on the theory of island biogeography, MacArthur and Wilson (1967) offered a brief prospect for the future of the science of biogeography. "Biogeography appears to us," they wrote, "to have developed to the extent that it can be reformulated in terms of the principles of population ecology and genetics To achieve this restatement . . . requires both theory and experiments that must be in large part novel. Simultaneously it demands a cultivation of population and community ecology in a way that contains much more evolutionary interpretation than has been traditional."

It is my hope that the unified neutral theory of biodiversity and biogeography that has been presented here—a theory that incorporates speciation and relative species abundance into the original theory of island biogeography—represents

CHAPTER TEN

a significant step, close to 40 years later, toward achieving MacArthur and Wilson's ambitious goal.

SUMMARY

1. The niche differentiation of virtually all species in natural communities obeys a small set of scaling laws and fitness invariance rules expressed as a series of shared life-history trade-offs.
2. These life-history trade-offs equalize the per capita relative fitness of species in the community, which sets the stage for ecological drift.
3. The equalized per capita relative fitnesses of species potentially decouple niche differentiation from control of species membership and relative abundance in the community.
4. The degree of decoupling and the importance of ecological drift then depend on the degree to which zero-sum dynamics apply to the species in the community.
5. Critics who conclude that natural communities are too stable for the theory, or conversely, that theoretical communities are too unstable under ecological drift, underestimate the importance of dispersal and dispersal limitation in structuring communities.
6. The present unified neutral theory of biodiversity and biogeography is only a beginning. There are many promising new theoretical and empirical paths to explore for achieving the ultimate reconciliation and unification of niche-assembly and dispersal-assembly perspectives in community ecology.