

Chapter 1

Introduction

Analogies are a vital part of science. The danger arises when we become so comfortable with our analogies that we confuse the analogy with reality; when this happens we are prevented from conceiving of nature in any other way. In medieval discourse the organizing Aristotelian analogy was *nature-as-an-organism*. Natural phenomena were seen to possess a life cycle: birth, growth, old age and death. Processes in the natural world were made intelligible in this way and were understood based on the inherent desire of things to attain goals (teleology). Cats hunt mice. Why? Because a cat is a predator and it is in the nature of predators to hunt. A stone falls to the ground rather than flying up into the air. Why? Because it is in the nature of heavy objects (i.e. objects made of “earth” rather than “air”) to move down. To know the nature of a thing was to know the thing itself (Dear 2007). This analogy was replaced in seventeenth century, by people like Descartes, Galileo and Newton, with a new one: the analogy of *nature-as-a-machine*. The hand of a mechanical clock doesn’t move around the face because it is in its “nature” to do so – place the hand alone on a table and it remains stationary. Rather, the hand, when properly installed, moves because a spring turns a wheel that hits a cog. Sometime during the 17th century Nature stopped being organic and became inert matter in motion obeying physical causes. By the end of the seventeenth century the very meaning of “explanation” itself had changed from its original Aristotelian one to the new mechanistic one: to know the cause-effect sequence was to know the thing.

What is true for great scientific ideas is also true of more modest ones because these too are more easily grasped through analogy. As will be explained in Chapter 2, the analogy used by ecologists to understand community assembly has been the dominant one of nature-as-a-machine. This analogy is useful as long as the cause-effect sequences are relatively few in number and constant over time. However, when dealing large numbers of organisms in interaction with each other and with the physical environment, the cause-effect sequences are so numerous, complicated and contingent that the nature-as-a-machine analogy becomes a hindrance rather than a help. Nature is no more a machine than an organic entity. Nature is simply Nature and we must not confuse the analogy with the thing.

The analogy used in this book for the process of community assembly is *nature-as-a-biased-die* and of chance constrained by selection based on physical differences. We have no hope of accurately predicting the outcome of a single throw of a biased die but, if we know how the die is biased, we can accurately predict the behaviour of an ensemble of many throws of the die; this is explained in Chapter 4. The resulting model is one of community assembly through chance constrained by natural selection. It is, as John Herschel¹ might have called it, a model community assembly by “higgledy-piggledy”.

¹ Sir John Frederick William Herschel FRS (1792 – 1871) was one of his generation’s top English scientists. In his *Physical Geography of the Globe* (Herschel 1872) he made the condescending description of Darwin’s notion of natural selection as the “law of higgledy-piggledy” and stated: “We can no more accept the principle of arbitrary and casual variation and natural selection as a sufficient account, *per se*, of the past and present organic world, than we can receive the Laputan method of composing books (pushed *à l’outrance*) as a sufficient one of Shakespeare and the *Principia*.” The irony is that the “Laputan” method to which Herschel refers (i.e. writing books by randomly combining words) comes from Johnathan Swift’s *Gulliver’s Travels* and refers to the

In this book I compare the process of community assembly to a never-ending game of crooked dice. Imagine a large hall, thick with cigar smoke, and filled with tables of different shapes and made of different materials. There are steel tables, ones made of foam, and ones covered with Maple syrup. There are tables that are slanted and ones that vibrate or move unexpectedly. The tables are the different physical environments upon which community assembly occurs. Community assembly occurs as players constantly throw millions of many-faced dice onto these tables, placing bets, and winning or losing resources. Each face represents a different species and every time a die presents a particular face the corresponding species wins some unit of resource and so increases its abundance. Most importantly, the dice are biased by different physical properties of each face (i.e. functional traits of species) such that different species have different probabilities of winning or losing on different tables. For instance, a face might be magnetized. Whether this helps or hinders the associated species in this analogy depends on whether the table (the environment) is made of steel or wood, or if the table itself is magnetized positively or negatively, or even if other faces of other dice are magnetized as well. The dynamics of this game of crooked craps is the dynamics of community assembly. Such a view of ecology might seem bizarre right now but, I hope, it will seem less strange when you reach the end of Chapter 7.

This book is not about plant traits even though the notion of a plant trait is central to it. There are many good books dealing with physiological (Lange *et al.*

one used by the absurd and useless philosophers inhabiting the island of Laputa. This tale was a parody of the Royal Society and its "scientists" of whom Herschel was a member.

1983, Pearcy *et al.* 1991, Lambers *et al.* 1998, Larcher 2001) or functional ecology (Grime 1979, Hendry and Grime 1993, Grime 2001, Grime 2007) but this book is not one of them. The notion of vegetation structure – of which species are present at a site, at what abundances, and how these two properties of vegetation change with environments – is also central to this book but I did not write a book about vegetation structure. Rather, this book deals with how vegetation structure emerges from the stochastic interplay of individual plants and environments and how we can quantitatively predict the resulting community structure of this interplay. You will be introduced to a new mathematical model of community assembly but the empirical knowledge of functional plant ecology will dominate the mathematical formalism; this book is definitely not a treatise on mathematical modelling in ecology.

So what is this book? It is *methodological* but only if “methodology” is understood to include its conceptual and philosophical components as well as its “how-to” aspects. It is *theoretical* but only if theory is understood in its most general context of converting biological assumptions into ecological conclusions through logical inference. It is *empirical* but only in the sense that the assumptions of the logical inference are justified by repeatable and quantifiable patterns between plant traits and environmental gradients that are measured in the field. Finally it is *synthetic* since the goal is to integrate each of the approaches mentioned before.

My thoughts about the nature of this interplay first began to take form over 20 years ago during my doctoral studies but they only began to form a unified

whole during the writing of this book. To my surprise I found the conceptual link between traits, plants and communities in one of the most fundamental notions of biology: natural selection. However, to appreciate this link, and especially to cast it in a more rigorous and mathematical form, I first had to realize that I (and most other biologists) have been wearing blinders with respect to our view of natural selection. Ever since Darwin and Wallace first worked out the process of natural selection we have been trained to mentally associate “natural selection” with “evolution”. This is perfectly understandable; the primary interest of both Darwin and Wallace was, not natural selection itself, but rather its evolutionary consequences. For Darwin and Wallace natural selection was a means to a more important end.

Natural selection is a process which, when operating between genotypes belonging to the same interbreeding group, *leads* to evolution by adaptation. Natural selection is one of the causes of evolution. No modern biologist would make the mistake of claiming that natural selection *is* evolution. We all understand that there are other processes besides natural selection that contribute to evolutionary divergence. However, most biologists have made the mistake of assuming that natural selection, though not the *only* cause of evolution, is *only* a cause of evolution; that is, that natural selection has no biological consequences besides those related to evolution. Because of this it seems strange to inquire about the consequence of natural selection when it operates between genotypes of different species (reproductively isolated groups). What does natural selection look like when it is occurring between

Dandelions and Sequoia trees? Since Dandelions and Sequoia trees cannot exchange genes, is natural selection not irrelevant at the interspecific level?

It took a long time before I even asked myself this question but, once I took it seriously, it led me to the formal mathematical conclusion, given in Chapter 5, that the consequence of natural selection between genotypes of different species, far from being meaningless simply because it is irrelevant to evolution, is the key to understanding community assembly and helps to reconcile seemingly contradictory views of the nature of ecological communities that have coexisted uneasily for more than a century (Clements 1916, Braun-Blanquet 1919, Tansley 1920, Cooper 1926, Gleason 1926, Tansley 1935). The philosopher Daniel Dennett (1995) described natural selection as “Darwin’s dangerous idea” because it is a process that is not limited to evolution or even to biology. Natural selection has the ability to leak out, infect, and conceptually modify non-evolutionary domains. This is because natural selection is not really a biological phenomenon at all but is a general domain-neutral algorithm. Natural selection is a logic machine. This machine takes heritable variation of properties between replicators as input and forces these replicators, by the brute force of its axioms, to scale fitness landscapes in the “design space” of these properties. This is true whether or not² the replicators exchange heritable information about their properties. “Evolution” is what happens when this logic machine forces interbreeding genotypes possessing heritable phenotypic variation to scale fitness landscapes. “Optimisation of complex objective

² The efficiency with which natural selection can scale fitness landscapes is greater if the replicators can exchange such heritable information because this increases the variation in such properties.

functions” is what happens when this logic machine forces replicating computer algorithms³ to scale fitness landscapes (Munakata 1998). Finally, and most importantly for this book, “community assembly” is what happens when this logic machine forces reproductively isolated genotypes (i.e. individuals of different species) possessing heritable phenotypic variation to scale fitness landscapes.

In order to sketch out the development of the ideas in this book, I hope that the reader will excuse a conceit on my part so that I can describe my personal journey. I do this, not because mine is a particularly profound or even unusual one, but simply because the structure of the book reflects it and because the conclusions that I have reached are undoubtedly constrained by it.

My doctoral research began in 1983 on the topic of plant zonation along freshwater shorelines. Going from the area of a river or marsh with fluctuating water levels that is continuously inundated to the area that is never flooded, one typically sees a turnover of species (zonation) such that particular groups of species are consistently found at different levels of water depth. The motivating question for the thesis was as simple as it was naive: How do freshwater plant species assemble themselves into recognizable zones? It was a question concerning community assembly. I began, as seemed logical to me at the time, at the community level by studying the pattern of species’ boundaries as a function of water depth along the shoreline. I asked whether there really were

³ Genetic algorithms are computer programs designed to solve a complex optimization problem. They do this by starting with an initial (even random) set of algorithms and compare the output of each algorithm to the value to be optimized (thus a “fitness” function). A subset are selected to “survive” with a probability proportional to their “fitness” and then (usually) the surviving algorithm codes are “mated” resulting in new hybrid “offspring” codes, along with some random “mutations”; this generates new variation. The process is iterated many times. Even very complex optimization problems, like the famous “travelling salesman” problem for which no exact solution is known, can be quickly and efficiently solved using this process of selection.

distinct “communities” (vegetation zones) along this environmental gradient. This was, by 1983, already a classic question in plant community ecology that had occupied several generations of plant ecologists without resolution. I concluded (Shipley and Keddy 1987) that the question was unresolved because it was ill-posed with respect to its causal generating mechanisms. I took my conclusion seriously and decided that I should abandon such community-level research in favour of experimental field studies that concentrated on the mechanisms themselves. I decided to descend to the conceptual level to population biology.

I then embarked on a field experiment in which I planted individuals of three dominant species, typical of different zones along the water-depth gradient, both alone and in the presence of the natural surrounding vegetation in each zone. The idea was to determine the relative importance of competitive vs. abiotic factors affecting growth and survival and how these factors might generate vegetation zones. After two seasons of hard work I had an answer (Shipley *et al.* 1991). Unfortunately, my answer had very unsettling implications for me unless I was willing to ignore the reason for the study in the first place. The calculation was both straightforward and depressing. If it took two years for one ecologist to study competition involving three species in one site then how long would it take to understand the zonation of freshwater aquatic species in general growing in sites whose environmental conditions were almost as unique as the species themselves? If the problem is expanded to include all plant species along all environmental gradients then it was clear that I faced a problem. In fact, plant community ecology as a discipline faces the same

problem using this research strategy since there are many orders of magnitude more plant species than there are active plant ecologists. This point is the main topic of Chapter 2.

Being young and in need of an advanced degree so that I wouldn't have to do real work, I didn't give up. I decided to descend yet another conceptual level in order to tackle the question from the perspective of functional plant ecology. The logic behind this second decision was as simple as for the first one. If there are too many species interacting with too many environmental variables to hope to put together a community piece by piece, then forget about species. The appeal of plant traits is that there are many fewer important functional traits than there are plant species. Individual plants migrate or fail to disperse, they grow or die, they reproduce or fail to set seed, due to their morphological, physiological and phenological attributes in interaction with the surrounding environment (both biotic and abiotic). It is not too much of a simplification to say that traits (phenotypic properties) link genotypes to the environment. Furthermore, plant traits are generalizable to a degree that plant species are not. Knowing about the relative competitive abilities of *Acorus calamus* L. versus *Carex crinita* Lam. at a given site does not allow one to predict the competitive ability of any other two species. On the other hand, if one can relate competitive ability (or any other ecological process) between any two species to some function of their traits, then one can extend this to any pair of species possessing such traits. I therefore finished my doctoral thesis by studying the tradeoffs between a set of functional traits for a large selection of the plant species growing at my site (Shipley *et al.*

1989). One of the most important empirical results of functional plant ecology is that certain typical values of plant traits are systematically found in similar environmental contexts even though the taxonomic composition differs; this is the definition of a “plant strategy” (Grime 2007). I could see Grime’s strategies emerge from the patterns of trait correlations in my study species.

After finishing my field work, and while writing up my results, I knew that I had still not answered the motivating question of the thesis since I had no way of going from plant traits to vegetation structure except in a very vague and qualitative way. I therefore brought the following problem to Len Lefkovitch, one of my thesis advisors and a very good statistician: If we know the typical values of a set of traits that should be found in a given environment (i.e. the plant strategy) and if we know the trait values of each species that could potentially reach the site, then how can we predict which species will actually be found, and at what abundances? There is no unique mathematical solution to this problem (it is under-identified) because there are more species than there are traits, but Len pointed out that this problem was similar to those studied in statistical mechanics, and suggested that I begin reading the work of a physicist, Edwin Jaynes, who had formulated the relationship between Information Theory and statistical mechanics through the Maximum Entropy Formalism (Jaynes 1957a, b, 1971). After reading Jaynes papers I realized that this was both a potential way to solve my statistical problem and a completely different way of viewing community assembly from the dominant one based on reductionism and population dynamics. I quickly wrote up my ideas as a fourth chapter of my

thesis but was advised to remove it since there was no empirical data to test the idea and because it was incomplete as an explanation.

I followed this advice (my thesis committee was correct on both counts) and the final version of the thesis (Shipley 1987) did not contain the missing fourth chapter but I never forgot about it. Chapter 4 of this book contains the seeds of my lost thesis chapter but it has taken more than 20 years for these seeds to germinate and grow. The basic idea of applying Jayne's Maximum Entropy Formalism to community assembly continued to feel right during the intervening years even though I couldn't completely articulate why, but one important stumbling block remained. The Maximum Entropy Formalism requires one to specify "constraints" on a process of random allocation (more correctly, on our information concerning this allocation) and I could not see how to formally and mathematically derive such constraints from ecological processes, as opposed to purely mathematical ones, in any consistent and general way. I put the problem to the back of my mind for many years although I continued to think about it occasionally, usually late at night when such vague ideas seem more promising.

It was only after collaborating with Éric Garnier and Denis Vile, and after stumbling by chance onto a paper describing a new and simple way of performing the calculations of the Maximum Entropy Formalism (Della Pietra *et al.* 1997), that I allowed myself to think about the problem during the day. Éric Garnier was interested in using "community-aggregated traits" to link plant traits to ecosystem processes. A community-aggregated trait is an average trait value

in which the trait value of each species is weighted by the relative abundance of the species at the site. I realized that a community-aggregated trait was the same thing as the constraint equations of the Maximum Entropy Formalism and began to think about how population dynamics would lead to a community-aggregated trait. It was only when I removed my conceptual blindersⁱ about natural selection at the interspecific level that I began to see the link between population dynamics, natural selection, community-aggregated traits, and community assembly.

I have now come to view that the process of community assembly as a statistical mechanistic one. Chance allocations of resources at the level of microstates interact with natural selection between reproductively isolated genotypes. Natural selection constrains such random allocations due to fitness differences between genotypes, leading to repeatable community structures at the macrostate level. The analogy that emerges is of nature as an immense casino. The species play craps with loaded dice for resource payoffs. There is no guarantee of success, only a probability of success. The dice that each species uses are biased due to the unique traits that it possesses, thus weighting the probability one way or another, but whether or not the bias helps or hurts the species depends on the nature of the tables (environments) on which the dice are thrown. It is wrong to think that this view of community structure is simply an intermediate position along a conceptual axis whose limits are communities as “individualistic” assemblages and as deterministic “units”. As Sewall Wright (1967) stated: “*The Darwinian process of continuous interplay of a random and a*

selective process is not intermediate between pure chance and pure determinism, but is in its consequences qualitatively utterly different from either.”

This book is therefore an exploration of these themes. I have always avoided working on ecological theory that could not be applied to real ecological communities in the field. Similarly, I have always been dissatisfied with empirical ecology that had no generalizing theoretical structure. Therefore, I hope that the ideas that I develop in this book will be evaluated based both on their instrumental usefulness and on their explanatory unification. The greatest compliment that others can pay to these ideas will be to criticize them, expose their weaknesses, and then improve them by concentrating simultaneously on empirical application and theoretical development. Certainly, the theoretical development presented in this book is still incomplete. For instance, in this book the community-aggregated traits point downwards from communities to species but the origin of the idea of a community-aggregated trait points upwards towards ecosystems, not downwards towards species. The idea of a community-aggregated trait comes from the “biomass ratio” hypothesis put forward by Grime (1998) and proposes that “the extent to which the traits of species affect ecosystem properties is likely to be strongly related to the contributions of the species...” to community structure in terms of biomass abundance. Recent work has supported this hypothesis (Garnier *et al.* 2004, Kazakou *et al.* 2006, Vile *et al.* 2006, Garnier *et al.* 2007, Violle *et al.* 2007). Perhaps, once the ideas in this book have been properly tested and refined, they can be extended upwards to ecosystems; certainly there is still work to do in linking community assembly via

traits to ecosystem properties. I won't speculate here on the likelihood of the success of this extension can't help adding a few words of caution by quoting Banquo from Shakespeare's *Macbeth*: "*If you can look into the seeds of time, and say which grain will grow and which will not, speak then to me.*"

- Braun-Blanquet, J. (1919). Notions d'élément et de territoire phytogéographiques. *Archives des sciences physiques et naturelles* **1**:497-512.
- Clements, F. E. (1916). *Plant succession: An analysis of the development of vegetation*. Carnegie Institution of Washington, Washington.
- Cooper, W. S. (1926). The fundamentals of vegetational change. *Ecology* **7**:391-413.
- Dear, P. (2007). *The intelligibility of nature. How science makes sense of the world*. University of Chicago Press., Chicago.
- Della Pietra, S., V. Della Pietra, and J. Lafferty. (1997). Inducing features of random fields. *IEEE Transactions Pattern Analysis and Machine Intelligence* **19**:1-13.
- Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. Simon and Schuster.
- Garnier, E., J. Cortez, G. Billes, M. L. Navas, C. Roumet, M. Debussche, G. Laurent, A. Blanchard, D. Aubry, A. Bellmann, C. Neill, and J. P. Toussaint. (2004). Plant functional markers capture ecosystem properties during secondary succession. *Ecology* **85**:2630-2637.
- Garnier, E., S. Lavorel, P. Ansquer, C. H., P. Cruz, J. Dolezal, O. Eriksson, C. Fortunel, H. Freitas, C. Golodets, K. Grigulis, C. Jouany, E. Kazakou, J. Kigel, M. Kleyer, V. Lehsten, J. Leps, T. Meier, R. Pakeman, M. Papadimitriou, V. P. Papanastasis, H. Queded, F. Quétier, M. Robson, C. Roumet, G. Rusch, C. Skarpe, M. Sternberg, J.-P. Theau, A. Thébault, D. Vile, and M. P. Zarovali. (2007). Assessing the effects of land use change on plant traits, communities and ecosystem functioning in grasslands: a standardized methodology and lessons from an application to 11 European sites. *Annals of Botany* **100**:In Press.
- Gleason, H. A. (1926). The individualistic concept of the plant association. *Bulletin of the Torrey Club*. **53**:7-26.
- Grime, J. P. (1979). *Plant strategies and vegetation processes*. John Wiley and Sons, New York.
- Grime, J. P. (1998). Benefits of plant diversity to ecosystems: immediate, filter and founder effects. *Journal of Ecology* **86**:902-910.
- Grime, J. P. (2001). *Plant Strategies, Vegetation Processes, and Ecosystem Properties.*, 2nd edition. John Wiley & Sons, New York.
- Grime, J. P. (2007). *Comparative plant ecology*. Castlepoint Press, Colvend.
- Hendry, G. A. F., and J. P. Grime. (1993). *Methods in comparative plant ecology*. Chapman & Hall, London.
- Herschel, J. F. W. (1872). *Physical geography of the globe*. A. and C. Black, London.
- Jaynes, E. T. (1957a). Information theory and statistical mechanics I. *The Physical Review* **106**:620-630.

- Jaynes, E. T. (1957b). Information theory and statistical mechanics II. *The Physical Review* **108**:171-190.
- Jaynes, E. T. (1971). Violations of Boltzmann's H theorem in real gases. *Physical Reviews* **A4**:747-751.
- Kazakou, E., D. Vile, B. Shipley, C. Gallet, and E. Garnier. (2006). Co-variations in litter decomposition, leaf traits and plant growth in species from a Mediterranean old-field succession. *Functional Ecology* **20**:21-30.
- Lambers, H., F. S. Chapin, and T. L. Pons. (1998). *Plant physiological ecology*. Springer, New York.
- Lange, O. L., P. S. Nobel, C. B. Osmond, and H. Zeigler. (1983). *Physiological Plant Ecology*. Springer-Verlag, Berlin.
- Larcher, W. (2001). *Physiological plant ecology*, 4th edition. Springer, Berlin.
- Munakata, T. (1998). *Fundamentals of the new artificial intelligence: beyond traditional paradigms*. Springer, New York.
- Pearcy, R. W., J. Ehleringer, H. A. Mooney, and P. W. Rundel, editors. (1991). *Plant physiological ecology*. Chapman and Hall, London.
- Shipley, B., and P. A. Keddy. (1987). The individualistic and community-unit concepts as falsifiable hypotheses. *Vegetatio* **69**:47-55.
- Shipley, B., P. A. Keddy, and L. P. Lefkovitch. (1991). Mechanisms producing Plant zonation along a water depth gradient - a comparison with the exposure gradient. *Canadian Journal of Botany* **69**:1420-1424.
- Shipley, B., P. A. Keddy, D. R. J. Moore, and K. Lemky. (1989). Regeneration and establishment strategies of emergent macrophytes. *Journal of Ecology* **77**:1093-1110.
- Shipley, J. W. (1987). *Pattern and mechanism in the emergent macrophyte communities along the Ottawa River (Canada)*. PhD. University of Ottawa, Ottawa.
- Tansley, A. G. (1920). The classification of vegetation and the concept of development. *Journal of Ecology* **8**:118-149.
- Tansley, A. G. (1935). The use and abuse of vegetational concepts and terms. *Ecology* **16**(3):284-307.
- Van Hulst, R. (1992). From population dynamics to vegetation dynamics: modelling succession as a species replacement process. *in* Glenn-Lewin, D. C., R. K. Peet, T. T. Veblen, editors. *Plant Succession. Theory and Prediction*. Chapman and Hall., London.
- Vile, D., B. Shipley, and E. Garnier. (2006). Ecosystem productivity can be predicted from potential relative growth rate and species abundance. *Ecology Letters* **9**:1061-1067.
- Violle, C., M. L. Navas, D. Vile, E. Kazakou, C. Fortunel, I. Hummel, and E. Garnier. (2007). Let the concept of trait be functional! *Oikos* **116**:882-892.
- Wright, S. J. (1967). Comments of the preliminary working papers of Eden and Waddington. Pages 117-120 in Moorhead, P. S., and M. M. Kaplan, editors. *Mathematical challenges to the neo-Darwinian interpretation of evolution*. Wistar Institute Symposium Monograph 5. Wistar Institute Press., Philadelphia.

ⁱ An important paper in this regard was Van Hulst (Van Hulst 1992) who showed that the equations of dynamic game theory - a theory usually applied to genotypes in a single species – could be applied to model ecological succession by viewing a “species” as an asexual genotype.