


Lectures for the ICTP Spring College in the Physics of Complex Systems

Introduction to Microbial Ecology



By Amir Erez,
Racah Institute of Physics,
The Hebrew University of Jerusalem, Israel.

March 13, 2025

Contents

1	Consumer-Resource Models: The Chemostat	3
1.1	Background	3
1.2	The Chemostat	3
1.3	Competition Between Two Species	8
1.4	Competitive Exclusion and the Paradox of the Plankton	10
1.5	Metabolic Tradeoffs and Diversity	12
1.6	Homework questions	12
2	May's Approach to Stability of Networks	13
2.1	Background and Main Questions	13
2.2	Mathematical Approach	14
2.3	Random Matrix Theory Background	14
2.4	Implications and Follow-Up Research	19
2.5	Conclusions	20
2.6	Homework questions	20
3	Generalized Lotka-Volterra Models	21
3.1	Introduction and Motivation	21
3.2	The Classical Lotka-Volterra Model	21
3.3	Incorporating Logistic Growth	24
3.4	Generalized Lotka–Volterra (gLV) Models	26
3.5	Emergence of Structured Dynamics from Random Interactions	30
4	Some Measures of Biodiversity and Their Relationship to Entropy	32
4.1	Motivation	32
4.2	Basic Entropy Measures and Effective Number of Species	33
4.3	Rényi Entropies and Hill Numbers	34
4.4	Rao's diversity index	36
4.5	Incorporating Species Similarity via a Similarity Matrix	37
4.6	Average Nucleotide Identity (ANI) and its limitations	42
4.7	UniFrac: Phylogenetic Measures of Community Similarity	42
4.8	Summary	43

Chapter 1

Consumer-Resource Models: The Chemostat

1.1 Background

- When attempting to model microbial dynamics, one popular approach is to model directly the microbial species abundance and the nutrient levels in the system.
- Essentially, we have to types of ‘players’ in the game: consumers (microbes, can also be producers), and resources (consumed for growth, can be toxic to certain species, etc.).
- This framework was formalized by MacArthur (1970), and together with generalized Lotka-Volterra models, is one of the main contemporary modeling approaches.
- Consumer-resource models extend well beyond microbial ecology and can be used to model ecosystems, economic systems, and more.
- It is not a coincidence that ecology and economics sounds similar. They both stem from the Greek word “oikos”, which means household. Many concepts from one field can apply to the other.
- For example, we recently considered the human concept of ‘debt’ and how it might apply to microbes!

1.2 The Chemostat

- The chemostat is a controlled bioreactor that maintains microbial populations in a steady state by continuously supplying nutrients while removing culture fluid at an equal rate (Figure 1.1).
- Originally introduced by Novick and Szilard (1950), the chemostat is a fundamental model for understanding microbial growth dynamics under nutrient-limited conditions.
- Its applications range from microbial ecology to industrial fermentation.

System Description

Consider a well-stirred chamber of constant volume V . The key features of the chemostat are as follows:

- **Nutrient Influx:** Nutrient solution enters the chamber at a rate Q (volume per unit time) with nutrient concentration c_{in} .
- **Dilution/Outflow:** To maintain constant volume, the medium is removed at the same rate Q . The outflow has nutrient concentration $c(t)$ and bacterial density $\rho(t)$.
- **Growth and Consumption:** Bacteria grow at a **per-capita growth rate**, $k_g(c)$,

$$k_g(c) = k_{g,\text{max}} \frac{c}{K + c},$$

where $k_{g,\text{max}}$ is the maximum growth rate and K is the half-saturation constant, which we discuss later in the chapter.

- **Nutrient Consumption and Yield:** As bacteria incorporate the nutrient into biomass, a yield factor ν is introduced, representing the number of nutrient molecules required per bacterium formed.

Model Dynamics

The dynamics of the chemostat involve a balance between bacterial growth and dilution.

Bacterial Population Dynamics

In a short time interval dt :

- **Outflow:** A volume $Q dt$ is removed, carrying away $\rho(t) Q dt$ bacteria.
- **Growth:** The bacterial population increases by $\rho(t) V k_g(c) dt$.

Thus, the net change in bacterial density is determined by the competition between these two processes.




Figure 1.1: Schematic of a chemostat. (Credit: Wikipedia).

Nutrient Dynamics

During the same interval dt :

- **Inflow:** The incoming nutrient adds $c_{\text{in}} Q dt$ amount of nutrient.
- **Outflow:** Nutrient is removed at the current concentration $c(t)$, yielding a loss of $c(t) Q dt$.
- **Consumption:** Bacteria consume nutrient at a rate proportional to their growth, dictated by the yield ν .

The conservation of mass for the nutrient leads to a differential equation coupling $c(t)$ and $\rho(t)$.

Dimensional Form

$$\begin{aligned}\frac{d(cV)}{dt} &= Q(c_{\text{in}} - c) - \nu\rho V k_g(c), \\ \frac{d(\rho V)}{dt} &= \rho V k_g(c) - Q\rho,\end{aligned}$$

with

$$k_g(c) = k_{g,\text{max}} \frac{c}{K + c}.$$

Dimensionless Formulation

To simplify the analysis, we introduce dimensionless variables and parameters:

- **Time Scale:** Define the inflow time (The time it takes to fill the empty tank or empty the full tank) as

$$T = \frac{V}{Q}.$$

The dimensionless time is then given by

$$\bar{t} = \frac{t}{T}.$$

- **Concentration and Density:** Define the dimensionless nutrient concentration as

$$\tilde{c} = \frac{c}{K}.$$

Similarly, rescale the bacterial density as

$$\tilde{\rho} = \frac{\rho\nu}{K}.$$

- **Growth Parameter:** A dimensionless maximal per-capita growth (per unit time filling the tank)

$$k_{g,\max} T = \gamma$$

With these definitions, the six original parameters $\{k_{g,\max}, K, Q, V, c_{\text{in}}, \nu\}$ enter the system only through the two dimensionless combinations: \tilde{c}_{in} (the rescaled nutrient input) and γ . This reduction simplifies the mathematical analysis. The equations become:

$$\begin{aligned}\frac{d\tilde{c}}{d\tilde{t}} &= \tilde{c}_{\text{in}} - \tilde{c} - \tilde{\rho} \gamma \frac{\tilde{c}}{1 + \tilde{c}}, \\ \frac{d\tilde{\rho}}{d\tilde{t}} &= \tilde{\rho} \left[\gamma \frac{\tilde{c}}{1 + \tilde{c}} - 1 \right].\end{aligned}$$

Steady-State Conditions

The system is expected to approach a steady state where bacterial growth exactly compensates for dilution.

We first calculate the nullclines: by setting one or the other time derivatives to zero.

1. $d\tilde{c}/d\tilde{t} = 0 \longrightarrow \tilde{\rho} = (\tilde{c}_{\text{in}} - \tilde{c})(1 + \tilde{c})/(\gamma\tilde{c})$
2. $d\tilde{\rho}/d\tilde{t} = 0 \longrightarrow \tilde{c} = 1/(\gamma - 1)$ or $\tilde{\rho} = 0$

The intersection of the nullclines gives the fixed points of the system. The phase field and steady state are shown in Figure 1.2.

1. Steady population size: $\tilde{c}_1 = \frac{1}{\gamma-1}$ and $\tilde{\rho}_1 = \tilde{c}_{\text{in}} - \frac{1}{\gamma-1}$.

Notice that the per-capita growth rate $\gamma \frac{\tilde{c}}{1+\tilde{c}} = 1$, or, $k_{g,\max} \frac{c}{c+K} = 1/T = Q/V$. Therefore, at this fixed point, the per-capita growth rate equals the dilution rate of the chemostat. This makes sense because **the bacteria need to grow at the same rate they are diluted to maintain steady state**.

2. $\tilde{\rho}_2 = 0$ and $\tilde{c} = \tilde{c}_{\text{in}}$.

This means no bacteria are in the tank depleting the nutrient, and therefore the same nutrient density provided is the one inside the tank.

Key observations include:

- **Stability Mechanism:** If the population temporarily falls below the steady state, the per capita nutrient availability increases, boosting growth and restoring the density. Conversely, an excess population leads to nutrient depletion and slower growth.
- **Washout Scenario:** If the bacterial growth rate is insufficient relative to the dilution rate, the bacteria will be washed out, and the population will collapse to zero.




Figure 1.2: Phase diagram for $\gamma = 3, c_{in} = 1$ (left) and $\gamma = 0.9, c_{in} = 1$ (right).

- **Fixed Points:** The existence of fixed points (steady states) is determined by solving the coupled equations for $\rho(t)$ and $c(t)$. Only nonnegative fixed points are physical. One must have

$$\gamma > \frac{1}{\bar{c}_{in}} + 1,$$

where γ is the dimensionless growth parameter. Adding back the dimensions,

$$k_{g,\max}T > \frac{K}{c_{in}} + 1 \longrightarrow k_{g,\max} > \frac{1}{T} + \frac{K}{c_{in}T}$$

At infinitely high input concentration, $c_{in} \rightarrow \infty$, we have $k_{g,\max} > \frac{1}{T}$. Therefore, at high input nutrient concentration, washout implies that the maximal per-capita growth rate is smaller than the dilution rate of the chemostat. The bacteria can't keep up with the dilution!

Discussion of Liebig's Law

- Liebig's law of the minimum: the growth of an organism is controlled not by the total amount of resources available, but by the scarcest (limiting) resource.
- This notion of a limiting resource is critical in understanding the competitive interactions among species.
- In the chemostat, if a single nutrient limits growth, then the microbial growth rate depends solely on the concentration of that nutrient. (Originally: nitrogen. Often: carbon. In tropical rainforest soils: phosphorous).
- The Monod form,

$$k_g(c) = k_{g,\max} \frac{c}{K + c},$$

provides a mathematical representation consistent with Liebig's law when a single nutrient is limiting. The form and its limiting behavior are shown in Figure 1.3.




Figure 1.3: Comparison of linear, Monod, and Hill $h = 2$ nutrient consumption curves. Equivalent to Holling’s type I, II and III functional response, respectively. Only the Monod curve is linear at $c \ll K$ while saturating at $c \gg K$. (From Lopez, Hein, Erez 2024)

- At low concentrations ($c \ll K$), the growth rate is approximately proportional to c , indicating that even small decreases in c can significantly limit growth.
- At high concentrations ($c \gg K$), the growth rate saturates at $k_{g,\max}$, demonstrating that further increases in c do not enhance growth.

Additional Context and Discussion

While the chemostat model provides essential insights into consumer–resource dynamics, it rests on some simplifying assumptions:

- The model assumes a well-mixed environment, thereby ignoring spatial heterogeneity.
- It considers a single limiting nutrient, whereas natural systems often have multiple interacting resources.
- Parameters are assumed constant, though in real systems rates may vary with time or environmental conditions. One classical example is “serial-dilution” where nutrients are added at the beginning of a “batch” and everything is diluted at the end of a “batch” (Figure. 1.4). The steady state of such a system is the steady state of the discrete map. The dynamics within the batch are non-trivial and lead to interesting effects that do not exist in the chemostat where $d/dt = 0$ at steady state.

1.3 Competition Between Two Species

In many experimental and natural scenarios, two microbial species may compete for the same limiting nutrient in a chemostat. We now extend the one–species model to include two competing species, denoted as Species A and Species B. Species A is characterized by a higher maximal growth rate but a higher half–saturation constant, while Species B exhibits a lower maximal growth rate but a lower half–saturation constant.




Figure 1.4: Schematic of a serial dilution system. At the beginning of each batch, a microbes are introduced to the system together with a nutrient bolus. After growth, a fraction of the culture is taken to inoculate the next batch. Steady state is when the species abundance distributions at the beginning and end of a batch are identical. From Erez, Lopez, Weiner, Meir, Wingreen (2020) under the Creative Commons license.

Two–Species Model Equations

Dimensional Form

We assume that the biomass grows according to Monod kinetics, while the nutrient is consumed in proportion to biomass production. The dimensional system is then given by

$$\begin{aligned}\frac{dc}{dt} &= (c_{\text{in}} - c)/T - \nu_A k_g^A(c) \rho_A - \nu_B k_g^B(c) \rho_B, \\ \frac{d\rho_A}{dt} &= \rho_A \left[k_g^A(c) - 1/T \right], \\ \frac{d\rho_B}{dt} &= \rho_B \left[k_g^B(c) - 1/T \right],\end{aligned}$$

with $T = V/Q$. The Monod functions for each species are defined as

$$k_g^A(c) = k_{g,\text{max}}^A \frac{c}{K_A + c}, \quad k_g^B(c) = k_{g,\text{max}}^B \frac{c}{K_B + c}.$$

Discussion of Competitive Outcomes

In the two–species model, both species compete for the same limiting nutrient.

- Initially, when nutrient c is abundant, Species A, with its higher maximal growth rate, may grow faster.
- As biomass increases the nutrient is depleted, driving c to lower levels. At low c , the species with the lower half–saturation constant (Species B) maintains a higher effective growth rate because it can utilize nutrient more efficiently.
- Consequently, at steady state, the ambient nutrient level may drop below K_A , curtailing the growth of Species A and allowing Species B (the K–strategist) to dominate.

- This outcome is consistent with Tilman's R^* theory (1982), which predicts that the species capable of reducing the resource to the lowest concentration at steady state will outcompete the other (Figure 1.5).
- If $K_A = K_B$ then the species with the higher $k_{g,\max}$ will take over.
- We've seen that the break-even point for a species is $k_g(c^*) = k_{g,\max} \frac{c^*}{c^* + K} = 1/T \rightarrow c^* = \frac{K}{k_{g,\max} T - 1}$. A nutrient density $c < c^*$ would lead that species to wash out. Therefore, in two-species competition, the equilibrium point is $c_A^* = c_B^*$, or $\frac{K_A}{k_{g,\max}^A T - 1} = \frac{K_B}{k_{g,\max}^B T - 1}$. The species with the lowest c^* can persist while the other species, unable to maintain a positive net growth rate at the lower nutrient level, is outcompeted.

1.4 Competitive Exclusion and the Paradox of the Plankton

- The competitive exclusion principle asserts that two species competing for the exact same limiting resource cannot coexist in a stable environment. (Georgy Gause, 1932).
- In our chemostat models, under steady-state conditions, the species with the lower c^* , typically the species with the lower effective half-saturation constant will dominate.
- If there are several nutrient types, there could be as many coexisting species as nutrient types.
- This was proved mathematically (Armstrong and McGehee 1980), establishing that any coexistence state is unstable.
- However, observations in natural systems, such as plankton communities in aquatic environments, appear to contradict this principle. Despite competing for a limited number of nutrients, a remarkably high diversity of plankton species is commonly observed.

This discrepancy is known as the **paradox of the plankton** (Hutchinson 1961).

Several mechanisms have been proposed to resolve this paradox:

- **Temporal Variability:** Environmental fluctuations can prevent the system from reaching a steady state, allowing species with different R^* values to coexist over time.
- **Spatial Heterogeneity:** Variations in resource availability across space can create microhabitats where different species have competitive advantages.
- **Predation and Other Interactions:** Top-down control by predators or inter-species interactions may further promote coexistence.
- **Niche Differentiation via Metabolic Tradeoffs:** As discussed below, differences in metabolic strategies resulting from enzyme allocation tradeoffs can facilitate coexistence by allowing species to collectively shape the nutrient environment.




Figure 1.5: R-k competition showing how the biomass grows until the nutrient level falls well below the R (high- K) species, which then goes extinct.

1.5 Metabolic Tradeoffs and Diversity

- Recent work by Posfai, Taillefumier, and Wingreen (2017) showed that metabolic tradeoffs can promote diversity in ecosystems.
- Species are characterized by fixed metabolic strategies, which specify how they allocate a limited enzyme budget across various nutrient uptake pathways.
- All species are assumed to have the same enzyme budget: whether specialists (consume just one type of resource) or generalists (consume several substitutable resources).
- This “fixed enzyme” condition is the loophole which allows to violate Armstrong and McGehee’s competitive exclusion theorem. It is roughly correct because bacteria operate near their biophysical limits—an excess of one enzyme would come at the expense of another.
- This tradeoff allows the species to collectively shape the nutrient environment, driving nutrient concentrations to values that are equally favorable to all species.
- As a result, an unlimited number of species can coexist even when the number of nutrients is limited, thereby bypassing the traditional limit imposed by the competitive exclusion principle.

1.6 Homework questions

1. Show that in the 1-species chemostat at steady state, if the flow rate is halved ($Q \rightarrow Q/2$) then the per-capita growth rate is also halved.
2. What happens to the per-capita growth rate if $K \rightarrow 0$?
3. In a very large tank, with very slow flow, what do we expect $c(t)/c_{\text{in}}$ at steady state? Find a leading order approximation of $\tilde{\rho}/\tilde{c}_{\text{in}}$ in this limit. What does it mean regarding the conversion of input nutrient to biomass? What does it mean regarding the outflow of nutrient in the system?

Chapter 2

May's Approach to Stability of Networks

In his seminal work, Robert May revolutionized our understanding of the stability of large, complex systems by applying random matrix theory to ecological networks. May's approach provided a framework to assess whether a community of many interacting species is likely to be stable when the interactions among species are assembled at random or emerge without systematic structure.

2.1 Background and Main Questions

Prior to May's work, classical ecological theory focused on low-dimensional, structured models.

- It was generally thought that since natural ecosystems are complex and stable, then stability arises from complexity.
- Indeed, when an ecosystem loses species, it often destabilizes, consistent with that prior intuition.
- Gardner and Ashby (1970) showed that even simple systems could lose stability with increasing complexity.
- May (1972) extended this idea by posing the key question: *What is the probability that a large, randomly assembled community is stable?*
- As the number of species n increases, and as the overall connectance C (the fraction of nonzero interactions in the community matrix) and the typical interaction strength (with standard deviation σ) vary, how does the likelihood of achieving local stability change?

2.2 Mathematical Approach

May considered a general dynamical system described by

$$\frac{d\mathbf{x}}{dt} = \mathbf{F}(\mathbf{x}),$$

with a fixed point \mathbf{x}^* satisfying $\mathbf{F}(\mathbf{x}^*) = 0$. Linearizing around this fixed point yields

$$\frac{d\delta\mathbf{x}}{dt} = \mathbf{J} \delta\mathbf{x},$$

where the Jacobian (or interaction) matrix is

$$\mathbf{J} = \left. \frac{\partial \mathbf{F}}{\partial \mathbf{x}} \right|_{\mathbf{x}^*}.$$

In ecological contexts, the off-diagonal elements of \mathbf{J} represent interspecific interactions while the diagonal elements typically reflect self-regulation (often taken as -1). This negative value signifies that each species has a built-in density-dependent control mechanism that curbs its own growth.

May assumed that the off-diagonal elements are independent random variables with zero mean and variance σ^2 , and that only a fraction C of them are nonzero (i.e. the connectance). This allowed the use of results from random matrix theory.

2.3 Random Matrix Theory Background

Two classical results in random matrix theory underpin May's analysis:

Wigner's Semicircle Law

- Wigner's semicircle law describes the eigenvalue distribution for large symmetric random matrices with independent entries with standard deviation σ (subject to symmetry).
- According to this law, the density of eigenvalues is given by a semicircular distribution: $\rho(\lambda) = \frac{1}{2\pi\sigma^2} \sqrt{4\sigma^2 - \lambda^2}$. The eigenvalues range in $\lambda^2 \leq 4\sigma^2$. This is shown in Figure 2.1.
- Wigner developed the semicircle theorem as part of his effort to understand the statistical properties of energy levels in complex atomic nuclei. In the 1950s, the intricate nature of nuclear interactions made it impractical to compute the exact eigenvalues of the nuclear Hamiltonian.




Figure 2.1: Histogram of eigenvalues computed from a normalized 1000×1000 random symmetric matrix $M = A/\sqrt{N}$, where A is drawn from the Gaussian Orthogonal Ensemble (GOE). Off-diagonal elements of A were independently sampled from $\mathcal{N}(0, \sigma = 1)$ and the diagonal elements from $\mathcal{N}(0, \sigma)$. The red curve represents the theoretical semicircular density $\rho(\lambda) = \frac{1}{2\pi\sigma^2} \sqrt{4\sigma^2 - \lambda^2}$, $\lambda \in [-2\sigma, 2\sigma]$, demonstrating the convergence of the eigenvalue distribution of M to Wigner's semicircle law.

Ginibre and Girko's Circular Law

- **Extension to Non-Hermitian Matrices:** For random matrices with Gaussian distribution of entries (the Ginibre ensembles), the circular law was established by Jean Ginibre (1965). **The matrix need not be symmetric, as the interaction matrix generally isn't.**
- Vyacheslav Girko (1984) further generalized. The theorem applies to matrices with i.i.d. entries whose eigenvalues can be complex.
- **Limiting Eigenvalue Distribution:** Shows that the eigenvalues of the normalized matrix $\frac{1}{\sqrt{n}}X$ converge to a uniform distribution over a circle of radius σ .
- **Universality:** The limiting behavior is independent of the specific distribution of the entries, given they satisfy mild moment conditions.
- **Significance:** Provides a foundational framework in random matrix theory for understanding the spectral properties of large, non-self-adjoint matrices.
- In May's context, aside from the diagonal (self-regulation) elements, the off-diagonal entries of the Jacobian are modeled as random with variance σ^2 and connectance C . Consequently, the eigenvalues (ignoring the diagonal shift) lie within a circle of radius $\sigma\sqrt{nC}$ in the complex plane. Let's prove this.

The radius $\sigma\sqrt{nC}$ in Ginibre and Girko's circular disk

A single row of the interaction matrix is denoted by x_1, x_2, \dots, x_n . In the sparse model, each x_i is independently nonzero with probability C and then is drawn from a Gaussian distribution of zero mean and variance σ^2 . Therefore,

$$x_i = \eta_i y_i$$

with η_i a Bernoulli random variable that is 1 with probability C and 0 with probability $1 - C$. The component $y_i \in \text{Gaussian}[\mu = 0, \sigma]$. The variance of x_i is given by

$$\text{Var}[x_i] = E[x_i^2] = E[\eta_i^2]E[y_i^2] = C\sigma^2$$

since $\eta_i^2 = \eta_i \rightarrow E[\eta_i] = C$ and $E[y_i^2] = \sigma^2$. Summing the entries in a row, $S = \sum_{i=1}^n x_i$, the variance of the sum (using the fact that the x_i are independent) is

$$\text{Var}[S] = \sum_{i=1}^n \text{Var}[x_i] = nC\sigma^2$$

Thus, we derived May's extension of the Ginibre disk, $\text{stdev}[S] = \sigma\sqrt{nC}$.




Figure 2.2: Left: Eigenvalue distribution of a normalized 2000×2000 real matrix X from normal distribution, where each entry is an i.i.d. Gaussian (with variance 1) and X is scaled by $1/\sqrt{N}$. The eigenvalues adhere to Girko's circular law, uniformly filling the unit disk (red circle). Right: The same simulation with the diagonal entries replaced by -1 , which shifts the eigenvalue distribution. The red circle in this panel, centered at $(-1, 0)$, indicates the theoretical support for the shifted eigenvalues. Both simulations utilize identical off-diagonal parameters and matrix size ($N = 2000$). Dots that appear outside Girko's circle are due to finite-size effects.

Linear algebra reminder: the eigenvalues of $M - I$

Let M be an $n \times n$ matrix and suppose that $v \neq 0$ is an eigenvector of M with eigenvalue λ , i.e.,

$$Mv = \lambda v.$$

Then,

$$(M - I)v = Mv - Iv = \lambda v - v = (\lambda - 1)v.$$

Thus, v is also an eigenvector of $M - I$ with eigenvalue $\lambda - 1$.

Eigenvalue Spectrum and Stability Criterion

When self-regulation is incorporated by setting the diagonal elements to -1 , the eigenvalue disk shifts leftward, centering at $(-1, 0)$. In this case, the eigenvalues are approximately contained within

$$|\lambda + 1| \leq \sigma\sqrt{nC}.$$

Local stability of the fixed point requires that all eigenvalues have negative real parts. Because a real positive eigenvalue will cause the difference from steady state, δx to grow exponentially. Stability, i.e., the decay of δx to zero, is ensured if the rightmost point of the disk, located at $-1 + \sigma\sqrt{nC}$, remains in the left half-plane:

$$-1 + \sigma\sqrt{nC} < 0 \quad \implies \quad \sigma\sqrt{nC} < 1.$$

Thus, May's stability criterion is

$$\sigma\sqrt{nC} < 1.$$

This threshold implies that as the complexity of the system increases (through n or C) or as the interactions become stronger (σ increases), the probability of maintaining stability rapidly decreases. This is shown in Figure 2.3.

Importantly: May's result does not assume symmetry, and therefore uses Girko's unit disk and not Wigner's ± 2 support. In general, one may apply the *elliptic law* stating that for the covariance of entries a_{ij} and a_{ji} of a random matrix, $E[a_{ij}a_{ji}] = \frac{\rho\sigma^2}{n}$, we have the support of the eigenvalue spectrum,

$$\frac{(\operatorname{Re} \lambda)^2}{\sigma^2(1 + \rho)^2} + \frac{(\operatorname{Im} \lambda)^2}{\sigma^2(1 - \rho)^2} \leq 1$$

For a symmetric matrix (with only real eigenvalues), $\rho = 1$ and we have $\lambda^2 \leq 4\sigma^2$ which is Wigner's law. For an uncorrelated matrix $\rho = 0$ we have $(\operatorname{Re} \lambda)^2 + (\operatorname{Im} \lambda)^2 \leq \sigma^2$ which is the Ginibre Girko law.

May's result is sharp in the limit $n \rightarrow \infty$, the apparent smoothness in the transition in Figure 2.3 arises from the finite- n transition.




Figure 2.3: Probability of instability as a function of the number of species n for community matrices of the form $M = -I + A$ where A has off-diagonal entries drawn independently from $\mathcal{N}(0, \sigma^2)$ with probability C (and zeros otherwise), and $\sigma = 0.2$. For each n (ranging from 4 to 120 in steps of 4) and for connectance values $C = 1.0$ and $C = 0.5$, 500 trials were performed to estimate the fraction of matrices with at least one eigenvalue having a positive real part. The vertical dashed lines indicate the May stability threshold, given by $n = \frac{1}{\sigma^2 C}$, which corresponds to the condition $\sigma\sqrt{nC} = 1$. The results demonstrate that the probability of instability increases with n , and that the critical value of n shifts to lower values as the connectance C increases.

Block Decomposition and Stability Enhancement:

If the large interaction matrix is decomposed into m independent blocks—each corresponding to a module or compartment of size n/m —then the stability condition for each block becomes

$$\sigma \sqrt{\frac{n}{m}} C < 1.$$

Since $\sqrt{n/m} < \sqrt{n}$ for $m > 1$, each block has a smaller effective spectral radius. This modular structure confines perturbations to individual blocks, reducing their potential to propagate across the entire system. Thus, breaking a large matrix into independent blocks enhances overall stability. **Caution:** by breaking M to blocks we're replacing previous nonzero interactions with zero, changing effectively C .

2.4 Implications and Follow-Up Research

May's work provided several key insights:

- **Paradigm Shift:** Prior to May, the common belief was that more complex ecosystems—those with a higher number of species and interactions—were inherently more stable. This was inspired by observing natural ecosystems and their apparent stability. May's analysis, however, demonstrated that increased complexity can lead to instability if the interaction strengths and connectance exceed a certain threshold.
- **Complexity-Stability Tradeoff:** There is a fundamental tradeoff between complexity (measured by n and C) and stability. High complexity can lead to instability unless interactions are sufficiently weak.
- **Relation to Natural Ecosystems:** in nature, though ecosystems are stable, it is often the case that either a species interacts strongly with few other species, or interacts weakly with many others. This is consistent with May's analysis.
- **May's result in the context of human systems:** Consider how social networks destabilize political systems that previously appeared robust. Consider how COVID19 spread so quickly and wrecked such havoc within weeks. Consider how global supply chains mean that a catastrophe in one country can cause shortages in the other side of the world. These are all examples of how the increased connectance and interaction strength in human relations creates instabilities.
- **Random vs. Structured Networks:** Although randomly connected communities tend to be unstable, empirical networks often display structured interactions (e.g., modular or hierarchical organization) that enhance stability.
- **Broad Impact:** May's approach has influenced various fields, including systems biology, neuroscience, control theory, and network science, by providing a quantitative framework to study the interplay between complexity and stability.

- Importantly, in many complex systems, the nonlinear dynamics may support the existence of **multiple equilibria**. Thus, even if the steady state chosen for analysis is unstable, **there may exist a nearby alternative steady state that is stable**. May's criterion, while providing critical insights into the conditions under which a given equilibrium loses stability, **does not rule out the possibility that the overall system may settle into a different, stable configuration**.

Subsequent research has built on May's framework by exploring how non-random (structured) interactions, nonlinearity, and external environmental drivers can buffer against instability, allowing complex ecosystems to persist despite high connectivity and strong interactions.

2.5 Conclusions

May's pioneering work demonstrated that high complexity in a system, quantified by the number of species, the connectance of interactions, and the strength of these interactions, can lead to instability when interactions are random. His stability criterion, $\sigma\sqrt{nC} < 1$, provided a clear, quantitative measure of the conditions under which a large complex system can be expected to be stable. This paradigm has profoundly influenced our understanding of ecological stability and continues to inspire research in diverse fields dealing with complex networks.

2.6 Homework questions

1. What happens if you take an $n \times n$ random matrix with $C = 1/2$ chosen randomly, and convert it to a block matrix with two blocks, $n/2 \times n/2$ that do not communicate with each other but each has internal connectance $C = 1$?
2. What is the meaning of the imaginary part of the eigenvalues in Girko's disk? Why don't we consider the imaginary part for the stability calculation?
3. In May's framework, self-regulation is modeled by setting each diagonal entry to -1. Suppose instead that the diagonal entries are $-d$ with $d > 0$. Derive how the stability condition is modified.
4. May's analysis assumes that the off-diagonal entries a_{ij} are independent. In many ecological networks, however, the reciprocal interactions a_{ij} and a_{ji} may be correlated.
 - Discuss qualitatively how a positive correlation ($\rho > 0$) versus a negative correlation ($\rho < 0$) between a_{ij} and a_{ji} might alter the eigenvalue spectrum.
 - Explain what implications these correlations might have on the stability condition derived by May.

Chapter 3

Generalized Lotka-Volterra Models

When we considered consumer-resource models, the species only interacted with the nutrients, and therefore indirectly with each other through the nutrient abundance. One can go the opposite direction, and consider only direct species-species interactions without modeling the nutrient dynamics at all.

3.1 Introduction and Motivation

- Ecological networks in nature are not random.
- Empirical food webs exhibit structured interactions: (i) clear trophic hierarchies; (ii) specialized predator–prey relationships. These shape the species-species interactions.
- In contrast to the random interaction matrices, real ecosystems exhibit nonrandom, often modular or compartmentalized structures.
- These structured interactions play a critical role in determining community dynamics and stability.

3.2 The Classical Lotka-Volterra Model

The classical Lotka–Volterra (LV) model describes the dynamics of two interacting species: a prey species $x(t)$ (e.g., rabbits) and a predator species $y(t)$ (e.g., lynx). The model is given by

$$\begin{aligned}\frac{dx}{dt} &= \alpha x - \beta xy, \\ \frac{dy}{dt} &= \delta xy - \gamma y,\end{aligned}$$

where

- α is the intrinsic growth rate of the prey,
- β is the predation rate coefficient,

- γ is the mortality rate of the predator,
- δ is the efficiency of converting consumed prey into predator offspring.

Fixed point analysis

A *fixed point* (or equilibrium) of the system is a pair (x^*, y^*) for which the time derivatives vanish; that is, $\frac{dx}{dt} = 0$ and $\frac{dy}{dt} = 0$.

For the prey, we set

$$\frac{dx}{dt} = \alpha x - \beta xy = 0.$$

Thus, either

$$x = 0 \quad (\text{trivial solution})$$

or

$$\alpha - \beta y = 0 \quad \implies \quad y = \frac{\alpha}{\beta}.$$

For the predator, we set

$$\frac{dy}{dt} = \delta xy - \gamma y = 0.$$

Thus, either

$$y = 0 \quad (\text{trivial solution})$$

or

$$\delta x - \gamma = 0 \quad \implies \quad x = \frac{\gamma}{\delta}.$$

To obtain the nontrivial fixed point where both x and y are nonzero, we discard the trivial solutions $x = 0$ and $y = 0$. Therefore, we have

$$x^* = \frac{\gamma}{\delta} \quad \text{and} \quad y^* = \frac{\alpha}{\beta}.$$

Fixed Points and Linear Stability

To analyze its stability, we compute the Jacobian matrix J evaluated at (x^*, y^*) :

$$J = \begin{pmatrix} \frac{\partial}{\partial x}(\alpha x - \beta xy) & \frac{\partial}{\partial y}(\alpha x - \beta xy) \\ \frac{\partial}{\partial x}(\delta xy - \gamma y) & \frac{\partial}{\partial y}(\delta xy - \gamma y) \end{pmatrix}.$$

Calculating the partial derivatives,

$$\begin{aligned} \frac{\partial}{\partial x}(\alpha x - \beta xy) &= \alpha - \beta y, & \frac{\partial}{\partial y}(\alpha x - \beta xy) &= -\beta x, \\ \frac{\partial}{\partial x}(\delta xy - \gamma y) &= \delta y, & \frac{\partial}{\partial y}(\delta xy - \gamma y) &= \delta x - \gamma. \end{aligned}$$

At the fixed point,

$$\begin{aligned}\frac{\partial}{\partial x}(\alpha x - \beta xy) \Big|_{(x^*, y^*)} &= \alpha - \beta \left(\frac{\alpha}{\beta} \right) = 0, \\ \frac{\partial}{\partial y}(\alpha x - \beta xy) \Big|_{(x^*, y^*)} &= -\beta \left(\frac{\gamma}{\delta} \right) = -\frac{\beta\gamma}{\delta}, \\ \frac{\partial}{\partial x}(\delta xy - \gamma y) \Big|_{(x^*, y^*)} &= \delta \left(\frac{\alpha}{\beta} \right) = \frac{\alpha\delta}{\beta}, \\ \frac{\partial}{\partial y}(\delta xy - \gamma y) \Big|_{(x^*, y^*)} &= \delta \left(\frac{\gamma}{\delta} \right) - \gamma = 0.\end{aligned}$$

Thus, the Jacobian becomes

$$J = \begin{pmatrix} 0 & -\frac{\beta\gamma}{\delta} \\ \frac{\alpha\delta}{\beta} & 0 \end{pmatrix}.$$

The eigenvalues λ satisfy

$$\lambda^2 + \alpha\gamma = 0 \quad \implies \quad \lambda = \pm i\sqrt{\alpha\gamma}.$$

Linearization about this fixed point yields a Jacobian matrix with purely imaginary eigenvalues, indicating that the system exhibits neutrally stable closed orbits (oscillatory dynamics) around the fixed point. **Essentially, any initial condition would result in its own closed orbit** (Figure 3.1). The closed orbits are generally not perfect circles, and this asymmetry reflects the time lag between the peaks of the two populations.

Derivation of the Constant of Motion

Starting from the classical Lotka–Volterra equations

$$\begin{aligned}\frac{dx}{dt} &= \alpha x - \beta xy, \\ \frac{dy}{dt} &= \delta xy - \gamma y,\end{aligned}$$

we eliminate the time variable by dividing the second equation by the first:

$$\frac{dy}{dx} = \frac{\frac{dy}{dt}}{\frac{dx}{dt}} = \frac{\delta xy - \gamma y}{\alpha x - \beta xy} = \frac{y(\delta x - \gamma)}{x(\alpha - \beta y)}.$$

Rearrange this equation to separate variables:

$$\frac{\alpha - \beta y}{y} dy = \frac{\delta x - \gamma}{x} dx.$$

We now integrate both sides:

$$\int \left(\frac{\alpha}{y} - \beta \right) dy = \int \left(\delta - \frac{\gamma}{x} \right) dx.$$

Performing the integrations, we obtain

$$\alpha \ln y - \beta y + C_1 = \delta x - \gamma \ln x + C_2,$$

where C_1 and C_2 are constants of integration. Combining these into a single constant C yields

$$\delta x - \gamma \ln x + \beta y - \alpha \ln y = C.$$

Thus, the function

$$H(x, y) = \delta x - \gamma \ln x + \beta y - \alpha \ln y$$

remains constant along trajectories of the system, i.e., $\frac{dH}{dt} = 0$. This conservation law confines the dynamics to the level curves $H(x, y) = C$, which are closed orbits in the phase plane.

Rabbit-Lynx Population Oscillations

The classical Lotka–Volterra model is historically linked to the observed oscillatory dynamics of the Canadian lynx and snowshoe hare populations. In this context:

- When predators (lynx) are scarce, the prey (hares) grow nearly exponentially.
- As the hare population increases, it supports a growing lynx population due to higher food availability.
- The increasing lynx numbers then suppress the hare population through predation.
- With fewer hares, the lynx population eventually declines, allowing the hare population to recover, and the cycle repeats.

The constant of motion $H(x, y)$ defines closed orbits in the phase space, mathematically describing the periodic oscillations observed in the field. While the classical model predicts neutrally stable cycles (with the amplitude of oscillations depending on initial conditions), in reality additional factors such as environmental variability, density dependence, and seasonal forcing tend to modify these dynamics, leading to more realistic damped or sustained oscillatory behavior.

In summary, while the classical LV model captures the essential predator–prey interactions leading to oscillations, its extensions via logistic growth and the generalized LV framework allow for a richer description of ecological communities, incorporating structured interactions and multiple species. This provides a powerful basis for understanding the complex dynamics observed in natural ecosystems.

3.3 Incorporating Logistic Growth

To account for resource limitations, logistic growth is incorporated into the prey dynamics:

$$\frac{dx}{dt} = r x \left(1 - \frac{x}{K}\right) - \beta xy,$$

where




Figure 3.1: Phase plane of the Lotka–Volterra system with parameters $\alpha = 1.0$, $\beta = 0.1$, $\gamma = 1.5$, and $\delta = 0.075$. The arrow directions indicate the instantaneous flow of the system, while the colors represent the speed (magnitude) of the vector field at each point. The nontrivial nullclines: the horizontal (red dashed) line $y = \alpha/\beta$ and the vertical (blue dashed) line $x = \gamma/\delta$. These lines intersect at the nontrivial fixed point. Trajectories from several initial conditions (ICs: [10, 5], [30, 40], [40, 10], and [20, 20]) integrated over the time interval [0, 50]. Each trajectory remains confined to a closed orbit determined by the constant of motion $H(x, y) = \delta x - \gamma \ln x + \beta y - \alpha \ln y$, which is preserved along solutions. Different initial conditions generally correspond to different constant values, and hence, to distinct closed orbits in the phase plane.

- r is the intrinsic growth rate of the prey, and
- K is the carrying capacity of the environment.

In this formulation, when $y = 0$, the prey population follows logistic growth and asymptotically approaches the carrying capacity K .

The predator dynamics remain unchanged:

$$\frac{dy}{dt} = \delta xy - \gamma y.$$

The nontrivial attractor (interior fixed point) is

$$x^* = \frac{\gamma}{\delta}, \quad y^* = \frac{r}{\beta} \left(1 - \frac{\gamma}{\delta K}\right),$$

provided that $K > \frac{\gamma}{\delta}$ (so that $y^* > 0$).

Impact on Long-Term Dynamics

- In the classical LV model, trajectories are neutrally stable closed orbits; the specific orbit followed is determined by the initial conditions.
- With the introduction of logistic growth, the system is pulled toward a unique attractor.
- For limit-cycles, one can use the Rosenzweig-MacArthur predator-prey model (1963)

$$\begin{aligned} \frac{dx}{dt} &= r x \left(1 - \frac{x}{K}\right) - \frac{a x y}{1 + a h x}, \\ \frac{dy}{dt} &= b \frac{a x y}{1 + a h x} - d y, \end{aligned}$$

where:

- $x(t)$ is the prey density,
- $y(t)$ is the predator density,
- r is the intrinsic growth rate of the prey,
- K is the carrying capacity,
- a is the predation rate coefficient,
- h is the handling time (the average time a predator spends capturing, subduing, consuming, and digesting an individual prey item),
- b is the conversion efficiency, and
- d is the predator mortality rate.

3.4 Generalized Lotka–Volterra (gLV) Models

To model communities with many interacting species, the LV framework is extended to the generalized Lotka–Volterra model. For n interacting species, the gLV equations take the form

$$\frac{dN_i}{dt} = N_i \left(r_i - \frac{N_i}{K_i} - \sum_{j \neq i} a_{ij} N_j \right), \quad i = 1, \dots, n,$$

where:

Phase Plane of the Lotka–Volterra System with Logistic Growth Nullclines, Vector Field, and Attractor




Figure 3.2: Phase plane of the logistic Lotka–Volterra system with prey dynamics given by $dx/dt = r x(1 - x/K) - \beta xy$ and predator dynamics $dy/dt = \delta xy - \gamma y$. The prey nullcline, $y = r/\beta(1 - x/K)$, and the predator nullcline, $x = \gamma/\delta$, are shown as red and blue dashed lines, respectively. The nontrivial fixed point (attractor) is located at $x^* = \gamma/\delta$ and $y^* = (r/\beta)(1 - \gamma/(\delta K))$, and trajectories from several initial conditions converge to this attractor. The vector field arrows, colored by their speed (magnitude), illustrate the local direction and rate of change in the phase plane.

- $N_i(t)$ is the population density of species i ,
- K_i is the carrying capacity of species i in the absence of interactions,
- r_i is the intrinsic growth rate of species i , and
- a_{ij} quantifies the effect of species j on species i (with $a_{ij} > 0$ for facilitation/mutualism and $a_{ij} < 0$ for competition or predation).

The interaction matrix $\mathbf{A} = (a_{ij})$ encapsulates the structure of the ecological network. By appropriately choosing \mathbf{A} , the carrying capacity K_i and the growth rates r_i , the gLV model




Figure 3.3: Time series of species abundances in a generalized Lotka–Volterra model with 10 species. The intrinsic growth rates r_i are drawn uniformly from $[0.5, 1.5]$ and the interaction matrix A is generated from a normal distribution with self-regulation (diagonal elements fixed to -1). The dynamics are governed by $dN_i/dt = N_i \left(r_i + \sum_{j=1}^{10} a_{ij} N_j \right)$. This simulation shows that, over time, the system converges to a steady state (or attractor) regardless of initial conditions. The final species abundances provide insight into the structure of the ecological interactions, which can be analyzed to infer whether interactions are predominantly competitive, mutualistic, or otherwise.

can capture a wide range of dynamics observed in natural communities, from (i) stable coexistence (Figure 3.3) to extinction of some species and oscillatory dynamics (Figure 3.4) and chaotic behavior.

There is a small industry of physicists applying advanced techniques from the statistical physics of disordered systems to calculate stability conditions if the A matrix is random, averaging over quenched disorder.

Inference of Interaction Parameters

- **Challenges:** Inferring the parameters r_i and a_{ij} from empirical data is challenging due to the high dimensionality of ecological communities, the presence of measurement noise, and environmental variability. Time-series data or controlled perturbation experiments are often required to obtain reliable estimates.
- **Ecological Insights:** Once the interaction matrix is inferred, its entries can be ana-




Figure 3.4: Two more examples of gLV dynamics, where one species can go extinct and where 3 competitors reach a stable limit cycle.

lyzed to determine the nature of species interactions, for example, whether most interactions are mutualistic (indicated by $a_{ij} > 0$) or competitive (indicated by $a_{ij} < 0$).

Types of Interspecific Interactions and the Signs of α_{ij} and α_{ji}

In the generalized Lotka–Volterra framework, consider two interacting species, i and j , with interaction coefficients α_{ij} and α_{ji} . The signs of these coefficients characterize the nature of the interaction between the species (Figure 3.5):

- **Mutualism:** Both species benefit from the interaction,

$$\alpha_{ij} > 0 \quad \text{and} \quad \alpha_{ji} > 0.$$

- **Competition:** Both species are harmed because they compete for the same limited resources,

$$\alpha_{ij} < 0 \quad \text{and} \quad \alpha_{ji} < 0.$$

- **Predation/Parasitism:** One species benefits at the expense of the other. If species i preys on species j , then

$$\alpha_{ij} > 0 \quad (\text{benefit to the predator } i) \quad \text{and} \quad \alpha_{ji} < 0 \quad (\text{harm to the prey } j).$$

- **Commensalism:** One species benefits while the other is unaffected. This interaction is typically modeled by having one nonzero coefficient and one zero coefficient,

$$\alpha_{ij} > 0 \quad \text{and} \quad \alpha_{ji} = 0.$$

- **Amensalism:** One species is harmed while the other is unaffected,

$$\alpha_{ij} < 0 \quad \text{and} \quad \alpha_{ji} = 0, \quad \text{or vice versa.}$$

These categorizations, while idealized, provide a useful framework for interpreting empirical estimates of interaction matrices.




Figure 3.5: The six possible types of symbiotic relationship, from mutual benefit to mutual harm. Source: wikipedia.

3.5 Emergence of Structured Dynamics from Random Interactions

- Recent studies have shown that even if an ecosystem starts with random interactions, the long-term dynamics may converge to a state that is effectively described by a structured generalized Lotka–Volterra (gLV) system.
- In this view, the initial random interactions—when subjected to ecological and evolutionary dynamics—can self-organize into an effective network of interactions that is nonrandom in structure.
- For example, work by Guy Bunin and collaborators has demonstrated that under certain conditions, large ecosystems with randomly assigned interaction coefficients tend to reach a steady state where the effective interactions become sparse and structured.
- In such scenarios, the complex dynamics of a high-dimensional, random system collapse onto a lower-dimensional manifold where species interactions are effectively described by a structured gLV model. This emergent structure often reflects underlying ecological constraints, such as resource partitioning and the need for mutualistic or predator–prey relationships that promote coexistence and stability.
- Going back to our discussing of Robert May’s stability criterion, because social networks happened very fast, human interaction didn’t have time to evolve and self-organize to stable structures...

- The analysis of gLV systems typically requires numerical simulations or approximations (e.g., mean-field or replica methods) to explore the full nonlinear dynamics. Although linearization around an equilibrium can recover May's type of condition, the full model can exhibit behaviors (e.g., limit cycles, multiple attractors) that the linearized approach does not capture.
- While elegant, May's result is inherently local; it tells us about the immediate response to small perturbations but does not predict the long-term dynamical behavior of the ecosystem.

Homework questions

1. Jacobian for the Logistic LV Model: Derive the Jacobian matrix, evaluate it at the fixed point, compute the eigenvalues at the fixed point and determine the local stability of the system.
2. The functional response in the Rosenzweig–MacArthur model is given by $\frac{ax}{1+ahx}$. Rewrite it in Monod form and identify the maximum predation rate and half-saturation constant. Comment on how this relates to the predation rate coefficient a and the handling time h .
3. Consider a system of n species described by the generalized Lotka–Volterra equations:

$$\frac{dN_i}{dt} = N_i \left(r_i - \frac{N_i}{K_i} - \sum_{j \neq i} a_{ij} N_j \right), \quad i = 1, \dots, n.$$

- (a) Derive the general form of the Jacobian matrix J evaluated at a fixed point $\mathbf{N}^* = (N_1^*, N_2^*, \dots, N_n^*)$.
- (b) Express the diagonal and off-diagonal elements of J in terms of N_i^* , K_i , and a_{ij} . Show that the Jacobian matrix J at the fixed point \mathbf{N}^* is given by

$$J_{ij} = \begin{cases} -\frac{N_i^*}{K_i}, & \text{if } i = j, \\ -a_{ij} N_i^*, & \text{if } i \neq j. \end{cases}$$

Chapter 4

Some Measures of Biodiversity and Their Relationship to Entropy

So far we have studied models of ecosystems, which may or may not support multi-species coexistence. But how is the biodiversity quantified? This is a big question and deserves a big answer. We will see there is a strong relationship between information, entropy and diversity (Jost, 2006).

Astonishingly, even though everyone heard about chemostats, May's stability criterion, and gLV equations, this is not the case for modern diversity metrics. However, I believe they are extremely important.

4.1 Motivation

- Biodiversity is much more than a simple count of species.
- Modern measures account not only for the number of species (species richness) but also for the evenness of their abundances and the similarity (or overlap) between them.
- Information theory provides a natural language for this task, with entropy emerging as a central concept.
- In this chapter, we review several measures of biodiversity, explain their relationship with entropy, and discuss how additional information—such as phylogenetic (evolutionary history) similarity and genetic overlap—can be incorporated into these measures.

4.2 Basic Entropy Measures and Effective Number of Species

Consider a community with N species where the relative abundance (or probability) of species i is denoted by p_i (with $\sum_{i=1}^N p_i = 1$). The *Shannon entropy* is defined as

$$H = - \sum_{i=1}^N p_i \ln p_i.$$

The exponential of Shannon entropy,

$$D_1 = \exp(H),$$

is often interpreted as the *effective number of species* or the Hill number of order 1. This measure adjusts species richness by taking into account the evenness of the distribution; for a perfectly even community with N species, $D_1 = N$.

Examples.

- Say we have two species of equal abundance, $p_1 = p_2 = 0.5$. The Shannon entropy is then $H = -\frac{1}{2} \ln \frac{1}{2} - -\frac{1}{2} \ln \frac{1}{2} = -\ln \frac{1}{2} = \ln 2$ and the effective number of species is $D_1 = e^{\ln 2} = 2$.
- If, however, one species became extinct, so $p_1 = 1, p_2 = 0$, then $H = 0$ and $D_1 = 1$. So we naturally extend from 1 to 2 effective species in this example, depending on their relative abundance.
- For 3 equally abundant species, $p_1 = p_2 = p_3 = 1/3$, we get $H = \ln 3$ and $D_1 = 3$.

How is this entropy related to information? Consider how surprised you are when you randomly sample an organism from the ecosystem.

- In information theory, the surprise (or information content) associated with observing species i is quantified as $-\ln p_i$. The Shannon entropy $H = -\sum_{i=1}^N p_i \ln p_i$ is then the expected surprise (i.e., the average information content) when sampling a random individual.
- A higher entropy indicates that the outcome is more uncertain—each observation carries more information because the species abundances are more even.
- A lower entropy implies that a few species dominate the community, making the outcome more predictable and the sampling experiment less informative.
- Maximal entropy is when all species are equally abundant: then the surprise is maximized when sampling a random individual. Therefore, the maximum entropy is simply $\ln N$ with N the number of species in the system. (Stat mech: the microcanonical ensemble).

4.3 Rényi Entropies and Hill Numbers

More generally, a one-parameter family of diversity measures is obtained from the Rényi entropies. For any order $q \geq 0$ (with $q \neq 1$), the Rényi entropy is defined as

$$H_q = \frac{1}{1-q} \ln \left(\sum_{i=1}^N p_i^q \right),$$

with the corresponding Hill number given by

$$D_q = \exp(H_q) = \left(\sum_{i=1}^N p_i^q \right)^{\frac{1}{1-q}}.$$

In the limit as $q \rightarrow 1$, D_q converges to D_1 ; in the limit as $q \rightarrow 0$, $D_0 = N$ recovers species richness. These different values of q emphasize different aspects of diversity: low q values give more weight to rare species, while high values emphasize common species.

Proof that $\lim_{q \rightarrow 1} H_q = H$

We have

$$H_q = \frac{1}{1-q} \ln \left(\sum_{i=1}^N p_i^q \right).$$

As $q \rightarrow 1$, both the numerator and denominator approach 0. Applying L'Hôpital's rule,

$$\frac{d}{dq} \ln \sum_i p_i^q = \frac{1}{\sum_i p_i^q} \frac{d}{dq} \sum_i e^{q \ln p_i} = \frac{\sum_i p_i^q \ln p_i}{\sum_i p_i^q} \Bigg|_{q=1} = \sum_i p_i \ln p_i,$$

since $\sum_i p_i = 1$. Thus,

$$\lim_{q \rightarrow 1} H_q = - \sum_{i=1}^N p_i \ln p_i = H.$$

Hence, the Hill number of order 1 is $D_1 = \exp(H)$.

Proof that $\lim_{q \rightarrow 0} D_q = N$

When $q = 0$, note that for any $p_i > 0$, $p_i^0 = 1$. Thus,

$$\sum_{i=1}^N p_i^q \xrightarrow{q \rightarrow 0} \sum_{i=1}^N 1 = N.$$

It follows that

$$\lim_{q \rightarrow 0} H_q = \lim_{q \rightarrow 0} \frac{1}{1-q} \ln \left(\sum_{i=1}^N p_i^q \right) = \ln N,$$

and therefore,

$$D_0 = \exp(\ln N) = N.$$

The Case $q = 2$ and its relation to the Gini Index

For $q = 2$, the Hill number is defined as

$$D_2 = \left(\sum_{i=1}^N p_i^2 \right)^{-1}.$$

- This quantity is the inverse of the Simpson index, $\lambda = \sum_{i=1}^N p_i^2$, which represents the probability that two individuals randomly drawn from the community belong to the same species.
- The complement, $1 - \lambda$, is known as the Gini–Simpson index and reflects the probability that two individuals belong to different species.
- When one species dominates the community—say, it comprises 99% of the total abundance—then λ becomes large, and consequently $1 - \lambda$ is small. In this context, a small Gini–Simpson index indicates low diversity because most individuals belong to the same species.
- Thus, for $q = 2$, the effective number of species, $D_2 = \lambda^{-1}$, emphasizes the role of dominant species. A low D_2 (or, equivalently, a high λ and low $1 - \lambda$) indicates that a single or a few species account for most of the community.

$q \rightarrow \infty$: ‘Min-entropy’ only considering the dominant species

The limit as $q \rightarrow \infty$ places all the emphasis on the most abundant species.

Derivation: Let

$$p_{\max} = \max\{p_1, p_2, \dots, p_N\}.$$

When q is very large, the term p_i^q for any $p_i < p_{\max}$ becomes negligibly small compared to p_{\max}^q . Therefore, we have

$$\sum_{i=1}^N p_i^q \approx p_{\max}^q.$$

Substituting this approximation into the expression for D_q ,

$$D_q \approx (p_{\max}^q)^{\frac{1}{1-q}} = p_{\max}^{\frac{q}{1-q}}.$$

Observe that as $q \rightarrow \infty$,

$$\frac{q}{1-q} \rightarrow -1,$$

so that

$$D_\infty = \lim_{q \rightarrow \infty} D_q = p_{\max}^{-1} = \frac{1}{p_{\max}}.$$

Interpretation:




Figure 4.1: Renyi entropy of a random variable with two possible outcomes against p_1 , where $P = (p_1, 1 - p_1)$. Shown are H_0 , H_1 , H_2 and H_∞ , with the unit on the vertical axis being the shannon. Source: wikipedia.

- The limit $D_\infty = 1/p_{\max}$ represents the *effective number of species* when only the most abundant species is considered.
- In other words, this measure is entirely determined by the dominant species, ignoring all others. It is sometimes referred to as the *Berger-Parker index*.
- Notice that for the case of N equally abundant species, we still retrieve $D_\infty = N$ as in the Shannon case. Indeed, this is true for all q values.
- A low value of D_∞ (i.e., a high p_{\max}) indicates that one species dominates the community, whereas a higher D_∞ suggests a more even distribution among the most abundant species.
- The min-entropy, in information theory, is the smallest of the Renyi entropies, and the most conservative way of measuring the unpredictability of a set of outcomes.
- The various Renyi entropies are all equal for a uniform distribution, but measure the unpredictability of a nonuniform distribution in different ways (Figure 4.1).
- $q = 0$ counts all species regardless of their abundance. $q > 1$ emphasizes the more abundant species over the less abundant ones. $q = 1$ is a special point where each species is measured according to its 'surprise' or information content derived from observing it. But q does not have to be integer. Sometimes we want to emphasize contributions from rare species, but still account for the abundance. Then we can use, say, $q = 1/2$.

4.4 Rao's diversity index

Not all species are similar or different to each other the same way.

- Rao’s quadratic diversity index (1982) is a measure of biodiversity that extends traditional diversity indices by incorporating information about the pairwise differences between species.
- Rather than considering only the relative abundances, it accounts for how distinct species are from each other, whether in terms of functional traits, phylogenetic relationships, or other measures of dissimilarity.
- For p_i as the relative abundance of species i , and d_{ij} as a dissimilarity (or distance) measure between species i and j , (and $d_{ii} = 0$), Rao’s quadratic Q is defined as

$$Q = \sum_i \sum_j p_i p_j d_{ij} = \sum_{i \neq j} p_i p_j d_{ij}$$

Since $\sum_i p_i = 1$ and therefore $(\sum_i p_i)^2 = 1 = \sum_i p_i \sum_j p_j = \sum_i p_i^2 + \sum_{i \neq j} p_i p_j$ we have

$$\sum_{i \neq j} p_i p_j = 1 - \sum_i p_i^2.$$

Therefore, for the naïve case where every species is equidistant from every other species, $d_{ij} = 1$, we have

$$Q|_{d_{ij}=1} = \sum_{i \neq j} p_i p_j = 1 - \sum_i p_i^2.$$

The expression on the right is the Gini-Simpson index. **Therefore, the Rao index generalizes the Gini-Simpson index to consider species-species similarity.**

Can we define a similarity-sensitive measure on the other Hill numbers?

4.5 Incorporating Species Similarity via a Similarity Matrix

- Traditional diversity measures typically assume that species are completely distinct. However, many species overlap in their genetic or functional characteristics.
- To account for this, we introduce a similarity matrix, following Leinster and Cobbold (2012) and Leinster (“Entropy and diversity the axiomatic approach”, 2020).
- These diversity measures will be almost completely neutral as to what ‘similarity’ means or how it is quantified.
- The similarity Z_{ij} between two species is measured on a scale of 0 to 1, with 0 representing complete dissimilarity and 1 representing identical species.
- The similarity-sensitive diversity measure D_q^Z generalizes the Hill numbers by weighting species abundances according to their pairwise similarities. It is defined as

$$D_q^Z = \left(\sum_{i=1}^N p_i \left(\sum_{j=1}^N Z_{ij} p_j \right)^{q-1} \right)^{\frac{1}{1-q}},$$




Figure 4.2: The major ranks: domain, kingdom, phylum, class, order, family, genus, and species, applied to the red fox, *Vulpes vulpes*. From: wikipedia.

where p_i is the relative abundance of species i (with $\sum_{i=1}^N p_i = 1$).

- The similarity Z_{ij} between two species can be interpreted as percentage genetic similarity.
- Functional similarity can also be quantified. For instance, suppose that we have a list of k functional traits satisfied by some species but not others. We can then define the similarity between two species as j/k , where j is the number of traits possessed by either both species or neither.
- Similarity can also be measured phylogenetically, that is, in terms of an evolutionary tree. For instance, Z_{ij} can be defined as the proportion of evolutionary time before the two species diverged, relative to some fixed start time.
- In the absence of better data, we can measure similarity crudely using taxonomy. For instance, we could define the similarity such that,

$$Z_{ij} = \begin{cases} 1, & \text{if the species are the same,} \\ 0.8, & \text{if the species are different but of the same genus,} \\ 0.5, & \text{if the species are of different genera but of the same family,} \\ 0, & \text{otherwise.} \end{cases}$$

For a picture of taxonomic ranks, see Fig. ref:ranks.

- More crudely, one could simply define

$$Z_{ij} = \begin{cases} 1, & \text{if the species are the same,} \\ 0, & \text{otherwise.} \end{cases}$$

Although this appears preposterous, in fact, the diversity measured defined using the Renyi entropies assume exactly that!

- **Reduction to Hill Numbers:** If species are assumed to be completely distinct (i.e., $Z_{ij} = \delta_{ij}$, where δ_{ij} is the Kronecker delta), then

$$\sum_{j=1}^N Z_{ij} p_j = p_i,$$

and the similarity-sensitive diversity measure reduces to

$$D_q^Z|_{Z_{ij}=\delta_{ij}} = \left(\sum_{i=1}^N p_i p_i^{q-1} \right)^{\frac{1}{1-q}} = \left(\sum_{i=1}^N p_i^q \right)^{\frac{1}{1-q}},$$

which is the usual Hill number of order q . One can take the logarithm of this expression to get a similarity-adjusted Renyi entropy.

- Given any metric of distance between species, d_{ij} , we can define $Z_{ij} = e^{-d_{ij}}$.

This adjusted formulation provides a more nuanced view of biodiversity, as it reflects both the species' abundances and the degree of overlap (or similarity) between species.

Application of the Leinster Cobbold Z matrix

If the diagonal entries of the Z matrix a 1 as in the examples above ($d_{ii} = 0$), then,

$$\tilde{p}_i = \sum_j Z_{ij} p_j \geq p_i$$

- This inequality states that a species appears more ordinary when the similarities between species are recognized than when they are ignored.
- Any species that is highly abundant is also highly ordinary: large p_i implies large \tilde{p}_i .
- But even if species i is rare, its ordinariness \tilde{p}_i will be high if there is some common species very similar to it.
- The ordinariness of species i will even be high if it is similar to several species that are each individually rare, but whose total abundance is large. This makes sense: the more thorny bushes a region contains, the more ordinary any thorny bush will seem, even if its particular species is rare.

We can now consider some of the Hill numbers adjusted for similarity.

- The case of $Z_{ij} = 1$ retrieves the naïve Hill formulations.
- For a general Z_{ij} matrix, the diversity of order 0 becomes,

$$D_0^Z = \sum_{i, p_i > 0} \frac{p_i}{\tilde{p}_i},$$

with $\tilde{p}_i \geq p_i$. The quantity p_i/\tilde{p}_i is large when the i 'th species is unusual.

- The min-entropy becomes,

$$D_{\infty}^Z = \frac{1}{\max_i \tilde{p}_i}$$

The order ∞ can be interpreted as before, except that is low not only if there is a single highly abundant species, but also if there is some highly abundant cluster of species.

- **Relation to the Rao index.** For any Z similarity matrix, we can define a dissimilarity matrix $d_{ij} = 1 - Z_{ij}$. It is easy to show that in this case, the the $q = 2$ case converges to the Rao index $\sum_{ij} p_i d_{ij} p_j$.
- **Phylogenetic diversity.** Our discussion would be incomplete without acknowledging the contributions of Chao, Chiu, and Jost (2010,2014,2015) to the field. They suggested a similar method of weighing phylogenetic distances in the Renyi entropies. The field is still very much alive and surely many more people do and will contribute to our understanding of how to quantify biodiversity.

How does it matter in practice?

Consider two examples shown in Figure 4.3.

- The first row consider the diversity of butterflies (of a particular subfamily) at the canopy and understory in the Ecuadorian rainforest.
- Note the need to consider a spectrum of q values to obtain a picture of the diversity.
- Using the naïve similarity matrix, the diversity profiles cross at a high q value. the diversity profile of the canopy lies above that of the understory until about $q = 5$, from which point they are almost identical. So for any sensitivity value, the canopy is more diverse than, or as diverse as, the understory.
- Using a similarity matrix that give $Z_{ij} = 0.5$ for different species from the same genus, the diversity profiles tell a different story: it is the understory that is more diverse. It is easy to see why. Most of the population in the canopy is from one genus, whereas the understory population is spread more evenly between genera.
- All diversity values drop when similarity is taken into account.
- The second row compares the gut microbial communities of lean and overweight humans, using the DNA sequences to estimate distances. Naïvely, the profiles cross at $q \approx 1$, suggesting that the gut microbiota of the lean child has a greater variety but is less evenly distributed than that of the overweight mother.
- However, using genetic similarity, the diversity of the lean child is higher at all values of q . Note the change in values for the effective species count.




Figure 4.3: Illustration of different similarity matrices: diversity profiles of six butterfly species using (a) the naive similarity matrix, and (b) a taxonomic similarity matrix. The bottom two panels show diversity profiles of the gut microbiomes in a lean child (TS1) and an overweight mother (TS3), using (c) naive and (d) genetic similarity matrices. Note the different scales. From: Leinster and Cobbold (2012).

4.6 Average Nucleotide Identity (ANI) and its limitations

Average Nucleotide Identity (ANI) is a widely used metric for assessing the genetic similarity between microbial genomes, by averaging the identities of aligned regions, thereby providing an effective means to assess species relatedness. While ANI is useful for delineating species boundaries, it has limitations:

- Hard to compare very different organisms: no easy way to align.
- ANI measures the average identity at the nucleotide level and may not capture functional or ecological differences between organisms.
- The thresholds used to define species (often around 95–96% ANI) are somewhat arbitrary and do not always correspond to ecological or phenotypic differences.
- ANI does not incorporate the evolutionary history or phylogenetic relationships among organisms, which are often critical for understanding community structure.

4.7 UniFrac: Phylogenetic Measures of Community Similarity

UniFrac is a phylogenetic metric used to compare microbial communities by incorporating evolutionary relationships among species rather than relying solely on species presence and abundance. In constructing a phylogenetic tree, branch lengths represent the evolutionary distance or divergence between taxa—longer branches indicate greater differences in genetic or evolutionary history.

Calculation of UniFrac

Given a phylogenetic tree that encompasses the taxa present in two communities, UniFrac (Lozupone and Knight, 2005) is computed as follows:

- **Unweighted UniFrac:** This metric is based solely on the presence or absence of taxa. For each branch in the tree, one checks if its descendant taxa appear exclusively in one community or are shared between both. Unweighted UniFrac is then defined as

$$\text{UniFrac}_{\text{unweighted}} = \frac{\text{total branch length unique to one community}}{\text{total branch length of the tree}}.$$

In other words, it is the fraction of the total evolutionary history (as measured by branch lengths) that is unique to one community.

- **Weighted UniFrac:** This version extends the idea by incorporating the relative abundances of the taxa. Here, each branch's length is weighted by the difference in the relative abundances of its descendant taxa between the two communities. This allows the metric to account quantitatively for how much each branch contributes to the overall dissimilarity, combining phylogenetic and abundance information.

Example Calculation of Unweighted UniFrac

Consider a simplified example with a phylogenetic tree that has three branches:

1. **Branch A:** Length = 0.3, leading to taxa found only in Community 1.
2. **Branch B:** Length = 0.5, leading to taxa found in both communities.
3. **Branch C:** Length = 0.2, leading to taxa found only in Community 2.

The total branch length of the tree is $0.3 + 0.5 + 0.2 = 1.0$. The branches unique to one community are Branch A and Branch C, which sum to $0.3 + 0.2 = 0.5$. Thus, the unweighted UniFrac is computed as:

$$\text{UniFrac}_{\text{unweighted}} = \frac{0.5}{1.0} = 0.5.$$

This result indicates that 50% of the total evolutionary history represented in the tree is unique to one of the communities, reflecting a moderate level of phylogenetic dissimilarity.

UniFrac has become a standard tool in microbial ecology for assessing community differences, especially when integrated with multivariate analyses such as principal coordinate analysis (PCoA) and clustering methods.

4.8 Summary

In this part, we have shown how biodiversity measures such as the Hill numbers D_q relate to entropy. In particular, we proved that as $q \rightarrow 1$ the Hill number converges to the exponential of Shannon entropy, while as $q \rightarrow 0$ it converges to species richness. We have also discussed the importance of incorporating species similarity through a similarity matrix Z , leading to the similarity-sensitive measure D_q^Z . Additionally, we provided an overview of the UniFrac metric, which leverages phylogenetic information to compare microbial communities, and we discussed the limitations of average nucleotide identity (ANI) in capturing the full complexity of microbial diversity. Together, these tools offer a robust and nuanced framework for quantifying biodiversity.

Homework questions

1. What are the limitations of using species richness ($q = 0$) as the sole measure of biodiversity?
2. How do Hill numbers provide a more nuanced understanding of community diversity?
3. How does incorporating a similarity matrix Z alter diversity estimates in communities with high genetic overlap?
4. Derive the limit $\lim_{q \rightarrow 1} H_q = H$ using L'Hôpital's rule.