

Assessing the sensitivity of microbial communities to external perturbations

A Thesis

submitted to

Indian Institute of Science Education and Research, Pune
in partial fulfillment of the academic requirements for the
BS-MS Dual Degree Programme

by

Krishna Samkaran Girish



Indian Institute of Science Education and Research Pune
Dr. Homi Bhabha Road,
Pashan, Pune 411008, INDIA.

May, 2024

Supervisor: Serguei Saavedra, Massachusetts Institute of Technology

Co-supervisor: Amit Apte, IISER Pune

© Krishna Samkaran Girish 2024

Certificate

This is to certify that this dissertation entitled ‘Assessing the sensitivity of microbial communities to external perturbations’ submitted towards the partial fulfillment of the BS-MS degree at the Indian Institute of Science Education and Research, Pune represents original research carried out by Krishna Samkaran Girish at Massachusetts Institute of Technology and Indian Institute of Science Education and Research, Pune, under the supervision of Professor Serguei Saavedra from Massachusetts Institute of Technology and Professor Amit Apte from Indian Institute of Science Education and Research, Pune, during the academic year May 2023 to March 2024.



Serguei Saavedra
Associate Professor
Department of Civil and
Environmental Engineering
Massachusetts Institute of
Technology



Amit Apte
Professor
Department of Data Science
Indian Institute of Science
Education and Research, Pune

Declaration

I, Krishna Samkaran Girish, hereby declare that the matter embodied in the report titled ‘Assessing the sensitivity of microbial communities to external perturbations’ is the results of the investigations carried out by me at the Department of Civil and Environmental Engineering, Massachusetts Institute of Technology and the Departments of Biology and Data Science, Indian Institute of Science Education and Research, Pune, from the period 01-07-2023 to 01-04-2024 under the supervision of Professor Serguei Saavedra and Professor Amit Apte and the same has not been submitted elsewhere for any other degree. Wherever others contribute, every effort is made to indicate this clearly, with due reference to the literature and acknowledgement of collaborative research and discussions.



KRISHNA SAMKARAN GIRISH
ROLL NUMBER 20191004
BS-MS
IISER PUNE
DATE: 26/03/2024

പരാനപേക്ഷം പ്രാണിക്കമരാൻ പശുതില്ലൊരിടത്തും
പരൻ പുമാനും പ്രകൃതിസഹായൻ പ്രപഞ്ചഘടനത്തിൽ
“Paraanapeksham praanikyamaraan pazhuthilloridathum
Paranpumaanum prakritisahaayam prapanchaghathanatthil.”

“No living being can exist without the help of others.
Even the supreme deity needed the help of mother nature in the creation of the universe.”
-Ulloor S. Parameswara Iyer, Premasangeetham

This thesis is dedicated to all the other living beings who help others exist, in any capacity.

Acknowledgements

First and foremost, none of this would have been possible without the support of my advisors, Prof. Serguei Saavedra and Prof. Amit Apte. Their support, guidance, advice and suggestions have been invaluable throughout the entire duration of the project. This has extended well beyond project-related mentorship (usually in the form of a barrage of ‘quick questions’ from me at the entrance of Serguei’s office) to more larger-scale facets, like how to be a good scientist, how to think about these kinds of complex problems without getting too lost in specific details, and how to be an effective, kind, and understanding mentor. I am lucky to have found such role models on my scientific journey. I also would like to extend my gratitude to Chengyi Long and Jie Deng, who have been a critical part of defining my experience at MIT. This has ranged from being a sounding board for several of my ridiculous ideas and practical suggestions on interpreting my results, to of course, finding solutions to the endless dilemma of where to go for lunch if not Flour or Stata. Outside the little room of 1-240 where so much happens, I am also grateful to the Zombie Squad (Charlie Neufeldt and Kenneth Gee) for all those problem-solving sessions that felt more like a friendly chat, the Magic Mates (Josh/Zach/Jordan/Michael) for all those weekends spent losing our minds over a sequence of cards, and MIT Quizbowl for wonderful weekly sessions, full of terrible jokes, cool diversions, and so many incredibly dubious Mystery Packets. Cambridge was a wonderful time because of all of you, thank you.

Before this, my scientific trajectory has been moulded and shaped by innumerable people to whom I cannot express my gratitude enough. Prof. Umesh Srinivasan has been a cornerstone of my development as a scientist, supporting me and constantly helping me achieve my ambitions throughout the pandemic and beyond. Prof. Deepak Barua has also been an incredibly important mentor to me, both in terms of the ways of thinking about ecology I picked up in his two courses, and in terms of seeing beyond all the complex math and statistics to derive core ecological insights. Prof. Leelavati Narlikar’s patient advice to a bumbling undergrad learning how to write good code was also instrumental. Beyond faculty mentorship, my personality, my aspirations and my growth as a person have all been driven subtly but surely by my wonderful set of friends. There are perhaps far too many to name within this small section, across batches, colleges and countries, but if we have spent meaningful time together, I want you to know it has enriched my life immensely. My time at IISER would not have been the same without all of you, from all the chats over the tables at

Shivsagar (but not V-Square, yuck), random 2 AM walks, and all the ridiculous business we pulled off in Hostel-1 and the seminar rooms of the Main Building.

Finally, all of this started one fateful morning in October 2006 when an exasperated father looking for ways to entertain his hyperactive 5-year old decides to take him birdwatching, and the child getting absolutely obsessed. Since then, it's been a journey of learning more and more about nature and the world and how to explain it. Fortunately, I have had support every stage of the way from my family. I can imagine it's not easy to raise a kid whose dream vacation destination is a remote sector of the Himalayas and/or the Andes, and decides that he's going to eschew all conventional norms and go to college to study ecology, a famously nebulous career path. Despite this, I never faced any resistance, and instead have always turned to my parents for practical, meaningful advice whenever situations have turned rough, whether academic or not. My gratitude is unbounded.

Contributions

Contributor Name	Contributor Role
Krishna S. Girish, Serguei Saavedra	Conceptualization
Krishna S. Girish, Serguei Saavedra	Methodology
Krishna S. Girish	Software
Krishna S. Girish	Validation
Krishna S. Girish	Formal analysis
Krishna S. Girish	Investigation
Serguei Saavedra	Resources
Krishna S. Girish	Data Curation
Krishna S. Girish	Writing - Original Draft
Krishna S. Girish, Serguei Saavedra	Writing - Review and Editing
Krishna S. Girish	Visualization
Serguei Saavedra, Amit Apte	Supervision
Krishna S. Girish, Serguei Saavedra	Project Administration

This contributor syntax is based on the Journal of Cell Science CRediT Taxonomy.¹

¹<https://journals.biologists.com/jcs/pages/author-contributions>

Abstract

The composition of microbial communities, in terms of relative abundances of constituent taxa, is constantly subjected to external perturbations, largely from changes in environmental parameters affecting dynamical trajectories. Understanding the sensitivity of community composition to such perturbations is crucial, given their role in maintenance of ecosystem function at scales from human health to global biogeochemical cycles. In this thesis, I unify statistical techniques from empirical dynamical modeling for nonlinear dynamical reconstruction with causal inference to evaluate how abundance and interactions affect the sensitivity of microbial communities to environmental parameter perturbations. Using simulated communities alongside longitudinal abundance data of taxa in three human microbiota, I develop a novel approach based on sampling small communities to detect and track the effect of external perturbations. Then, using causal inference to identify the effects of abundance and interactions on sensitivity, I identify 4 sub-community profiles based on causal effect signs. Simulated communities with different interaction strengths display a continuous transition in frequencies of different profiles, with real communities displaying frequencies similar to simulated weak interaction regimes. Detecting perturbations within each profile, I find that rarer profiles in competitive communities have consistently more sensitive responses, while in weak interaction regimes, this identity is variable. Real microbiota appear to show inconsistent sensitivity patterns as well. Despite the individualized nature of microbial compositions and the variety of possible responses, our study shows the possibility to synthesize the nonlinear nature of microbiota dynamics. This approach can be relevant for personalized approaches in monitoring and designing optimal therapeutic treatments.

Contents

Front Matter	i
Abstract	xiv
Contents	xviii
1 Introduction	1
1.1 The organization of this thesis	1
1.2 Microbial systems and microbiota	3
1.3 Understanding ecological sensitivity	7
1.3.1 Structural Stability	10
1.3.2 Non-equilibrium Notions of Stability	12
1.3.3 Other Stability Metrics	15
2 A primer on empirical dynamic modeling and causal inference	17
2.1 Problems with parametric models	17

2.2	Introduction to Empirical Dynamic Modeling	19
2.2.1	Takens' Embedding Theorem	19
2.2.2	Simplex Projection	22
2.2.3	The S-Map Algorithm	24
2.2.4	Structural Sensitivity and the S-map	28
2.3	A brief foray into nonparametric causal inference	32
2.3.1	Probabilistic causal inference	34
3	Methods: Understanding Microbial Communities	39
3.1	Generating Sample Communities	40
3.2	Developing the Similarity Statistic	47
3.3	Causal Inference on Sub-Sampled Communities	51
3.4	Parameters for the Pipeline	53
3.5	Sourcing Microbial Data	56
3.6	Metadata Processing and Causal Discovery	60
4	Results and findings	63
4.1	Synthetic Data: Generating Time Series and Jacobians	63
4.2	Synthetic Data: Similarity Statistic	64
4.3	Synthetic Data: Causal Profiling	70
4.3.1	Sensitivity of Causal Profiles	74
4.4	In Vivo Microbiota Data	77
4.4.1	Family Composition and Structure	77
4.4.2	Similarity Statistic on Real Data	79
4.4.3	Causal Discovery and Profile Assignment	81

4.4.4	Similarity Statistic Across Profiles	84
4.5	Other Datasets	88
4.5.1	Moving Pictures of the Human Microbiota	88
4.5.2	Antibiotic Recovery	89
5	Discussion, summary and outlook	93
5.1	Synthesizing results, briefly	94
5.2	Implications	96
5.2.1	Ecological Implications	96
5.2.2	Implications for Clinical Interventions and Microbial Therapeutics . .	100
5.3	Limitations	102
5.4	Future Directions	105
	Appendices	111
A	Robustness Checks	111
A.1	Simplex Projection to fill in missing values	112
A.2	Other Taxonomic Levels	115
A.3	Differing Rare Species Cutoffs	117
A.4	Larger Cluster Sizes	118
B	Additional parameter value plots	123
B.1	Families included in each microbiota	123
B.2	Abundance and Mean Interaction Strength Variation	125
B.3	Simulated Time-series Plots: Other Parameter Values	127

B.4 Distribution of Natural Causal Effects	128
--	-----

References	131
-------------------	------------

Chapter 1

Introduction

We think too big, we think ourselves as one whole thing, and we claim that this collection is a name and is a being; but deep inside, where every cell divides, it sets upon the rule that self-interest is divine.

Danny Schmidt, *This Too Shall Pass*

1.1 The organization of this thesis

Much like the intricate network of a microbiota, with each species entwined in noisy clockwork with the others, this thesis, too, consists of multiple interacting sections. In my opinion, the optimal way to understand the contents of this thesis is by reading through all sections, the same way that a complete picture of an ecological community can be observed by sampling all components, both biotic and abiotic.

However, if you are interested in a particular subset of sections or suite of methods, this section briefly highlights the contents and aims of each of the subsequent sections. Fortunately, the ideas introduced in one segment are not exclusive of the others or the overarching aim of

this thesis- drawing an ecological synergy again, as I highlight in Chapter 2 about Empirical Dynamical Modeling, we can reconstruct quite a great deal of information about the system even if we sample from just a few key components. I highlight the contents of the subsequent chapters below-

- In the remainder of this chapter, the introduction, I highlight the microbiota as a complex, highly diverse, dynamic (and endlessly fascinating) ecological system, and the overarching aim to understand its sensitivity to external environmental perturbations. Subsequent to this, I change tracks to briefly review ecological notions of stability, how they have been quantified in the past, and illustrate how the notion of structural stability as a potent metric to assess environmental perturbation sensitivity in these types of communities. While far from a complete synthesis of dimensions of stability in ecological communities, this section highlights potential pitfalls with many usual metrics that necessitate novel approaches.
- Chapter 2 focuses on the modeling frameworks utilized to analyze these microbiota, specifically empirical dynamic modeling and causal inference. This section is intended to provide a primer on these techniques, building them up from first principles up towards the specific applications used in this analysis. For readers familiar with these methods already, or interested in major biological insights without going deep into the mathematical framework behind them, I also provide abridged summaries of the details of the specific tools we borrow from these formalisms. This section is intended as a brief review of these methods, to provide context and justify why these are optimal approaches in our system, and reviewing the state of the art and shortcomings inherent in these methodologies.
- Chapter 3 takes these tools and develops them further towards novel techniques and insights to understand community sensitivity in microbiota specifically, describing the various facets and considerations that go into this. This is similar to a Methods section in spirit, and the first chapter where I bring in my new, original work.
- In Chapter 4, I highlight the results of these techniques I have developed, building up and synthesizing the outcomes from the application of different approaches and methods towards two unifying and general results- one developing a system-agnostic metric of detecting structural sensitivity in complex ecosystems, and one about differential

frequencies and structural sensitivity of certain sub-communities of microbiota based on interaction strengths.

- Chapter 5, similar to a discussion section, contextualizes these results in a broader sense, drawing both from its applicability to understanding these results in the setting of general ecosystem stabilization, future ecological sampling study design, as well as towards clinical intervention modeling. Critically, I also highlight shortcomings in this approach, and what could be done in the future to scale this up to tangible effects on ecosystem restoration and gastrointestinal health maintenance.
- Finally, the appendices act as a sort of supplementary information to all of the above. This includes replications of these results with modified assumptions to check for robustness, and further plots (and their interpretations) to gain a more complete picture of the system.

Each of these sections are largely self-contained, all while orbiting around the main thematic question of what drives the susceptibility of microbiota to changes in the environment. I hope this thesis allows you to appreciate the importance of these systems, the rich dynamics occurring in them, and the potency and significance of the statistical tools used to analyze them.

1.2 Microbial systems and microbiota

Microbial communities in nature have a profound impact on sustaining ecosystem function across various biological scales. These functions encompass the maintenance of host health, wherein trillions of commensal microbes residing in the organs of larger animals and plants fulfill vital metabolic roles (Ursell et al., 2012; Hou et al., 2022). Additionally, these microbial communities contribute to soil fertility enrichment and plant growth (Panke-Buisse et al., 2015; Wagg et al., 2019) and are integral to global biogeochemical cycles through their involvement in fundamental chemical transformations (Falkowski et al., 2008). Despite their microscopic scale, these microbial communities have been estimated to constitute 15% of global biomass, comprising an estimated one trillion different species. (Locey and Lennon, 2016; Bar-On et al., 2018).

These ubiquitous communities are typically highly complex, diverse, and dynamic. In particular, the human alimentary canal contains highly diverse microbial communities (termed as

microbiota) in the oral cavity and all along the gut lumen, which provide a range of metabolic and immune functions to the host (Marchesi et al., 2016). These microbial communities have been implicated in several organismal functions through the synthesis of metabolites such as short-chain fatty acids, which may then mediate metabolic pathways that affect inflammation, body weight, and neural function, among other things (Tan et al., 2014). The gut microbiota has been shown to contain about 300-500 species (Guarner and Malagelada, 2003). The oral microbiota tends to have higher diversity than the gut (over 700 species (Wade, 2013)), and provides immune support by forming protective biofilms and preventing colonization by pathogens. In addition, each of these bacterial species may contain a multitude of strains, further increasing the extent of genetic and phenotypic diversity, as well as possible functional outcomes in the community. Therefore, these microbiota constitute a complex ecological system, with many functional roles and interactions.

Recent development of sampling techniques based on high-throughput sequencing of 16S rRNA (Eckburg et al., 2005; Janda and Abbott, 2007; Caporaso et al., 2010; Navas-Molina et al., 2013) has allowed us to sample human microbial ecosystems and its associated changes over time and space at an unprecedented scale. This also permits design of controlled trials to assess how specific interventions affect the organization and abundances of different taxa in these systems. Such sampling experiments have revealed significant variation in the species present and their relative abundances observed when comparing the microbiota across different individuals- this variation can be seen at a smaller scale across individuals within a population, as well as larger-scale differences when comparing between populations (Manor et al., 2020). Sampling gut microbial compositions across different regions of the world have shown fundamental differences in organization of microbiota across different regions of the world in terms of dominant genera of bacteria (Suzuki and Worobey, 2014; He et al., 2018)- based on this, researchers proposed the idea of enterotypes (MetaHIT Consortium (additional members) et al., 2011), a classification of individuals on the basis of the dominant bacteria present in their gut microbiota. This variation has been linked with differences in environmental conditions and long-term local lifestyle and diet choices (Wu et al., 2011).

In addition, the structure of these microbiota can change drastically within a single individual over time. These changes are driven by continuously changing external environments, as well as endogenous ecological dynamics in the microbial community. Environmental perturbations that restructure microbiota may emerge from a wide range of exogenously-

mediated factors, such as pathogenic infection (Pérez-Cobas et al., 2014), antibiotic usage (Dethlefsen et al., 2008; Chng et al., 2020), and changes in food intake (David et al., 2014b). In accordance with this idea of continuous perturbation, sampled time-series of the states of microbiota have been observed to display high levels of nonlinearity and stochasticity, with several instances of transitions between alternative stable states and many random fluctuations (David et al., 2014a; Faust et al., 2015; Jones et al., 2022; Fujita et al., 2023). This complexity and the high-dimensional nature of these systems makes temporal dynamics and the effect of intervention hard to predict.

At a community level, this wide spectrum of continuous perturbations to the environment could affect the structure of microbiota through addition/removal of species, directly affecting species' relative abundances, changing physico-chemical parameters of the environment affecting growth rates and interactions, and modulating the sensitivity of the system to future stochastic disturbances that may facilitate or induce critical transitions to alternative stable states (Sommer et al., 2017; Amor et al., 2020; Nguyen et al., 2021a; Fujita et al., 2023). Changes in physiological parameters in particular could have significant and rapid effects on microbiota on account of their large spatial extent compared to the size of bacterial communities leading to pervasive effects across populations in discrete patches. These physiological perturbations are of particular concern when they operate at a large scale and have persistent effects on the structure of the microbial community. Such extensive reorganizing events of the ecological structure of the microbiota, termed as dysbiosis, have been linked with the onset of several illnesses (Carding et al., 2015; Bäumlér and Sperandio, 2016), including intestinal disorders, diabetes, neurological conditions, and onset of cancer (Hartstra et al., 2015; Gilbert et al., 2016; S. Zou et al., 2018; Ke et al., 2023). These dysbiotic events are of particular concern when the microbiota settles in this disturbed state, making it extremely hard to recover back to a healthy state (Van De Guchte et al., 2020; Fujita et al., 2023) and causing persistent disruption of ecological function.

Given the crucial functions provided by microbiota in maintaining homeostatic equilibrium in the human body, understanding the determinants of the sensitivity of its composition (in terms of the relative abundances of constituent species) to exogenous perturbations changing the ecological dynamics is of paramount importance. Clinical studies have shown a great deal of heterogeneity in the sensitivity of response of different subjects to similar perturbations, both in terms of immediate effect on the gut microbial structure as well as the

period of persistence of community changes (and often, physiological symptoms) after the departure of the perturbation, indicating the extent of dysbiosis. Inter-individual microbiota variation is known to be an important driver of the sensitivity of response (Rashidi et al., 2021); however, studies have also isolated other important factors determining susceptibility. These include the instantaneous state of the microbiota prior to perturbation (Raymond et al., 2016), diet (Brown et al., 2012; David et al., 2014b), priority effects (Sprockett et al., 2018), and geographical location (De Filippo et al., 2010).

The return to a healthy baseline state after perturbations arising from antibiotic treatment has been found by large-scale longitudinal screening to be associated with the re-establishment of certain specific bacterial taxa (Chng et al., 2020). Other studies have also identified potential candidate taxa that could contribute to long-term stability of the gut microbiota (eg. O’Callaghan and Van Sinderen, 2016; Dempsey and Corr, 2022). While these give us a deep mechanistic understanding of components of the microbiota, scaling up these studies towards whole-microbiota stabilization carries two major limitations. One is that robustly identifying keystone taxa across multiple microbiota that always promote or deplete stability is extremely challenging considering the high standing diversity, including subspecific strains (Wang et al., 2023; Weiss et al., 2023; Wolff et al., 2023). Secondly, the extent and rate of introduction of these taxa in the microbiota following the perturbation is contingent on several complex, interrelated and nonlinearly varying factors. Therefore, to gain a generalizable understanding of microbiota stability that accounts for the large variability over subjects and over time, and the potentially infinite kinds of perturbations that could affect the system, it is perhaps more feasible to seek system-level properties that promote certain outcomes instead of isolating specific beneficial taxa.

Therefore, accurately reconstructing and assessing the relative roles of drivers of community-level sensitivity could be fruitful in facilitating the development of optimal clinical interventions to alleviate a wide range of gastrointestinal perturbations. This is the overarching question we will be confronting- **what factors determine the sensitivity of microbiota to environmental perturbations?** Once we gain this theoretically-motivated understanding, the results could manifest into real-world functional outcomes as targeted probiotic or antibiotic therapies to catalyze recovery towards a baseline, or steps to improve resilience of the microbiota ecosystem to future perturbations. To do this, we first need to develop a means of tracking the stability of the community over time, to understand the effect of disturbance,

and what determines the extent of its effect.

1.3 Understanding ecological sensitivity

In this section, I briefly review notions of ecosystem stability and sensitivity and how to compute them. In order to understand variation in stability of the gut microbiota as a function of environmental changes, it is critical to select an appropriate metric of stability that is able to capture the complex nature of responses in the system. Broadly, ecological stability is concerned with the capability of an ecosystem to remain in its current state after the appearance of a disturbance. However, this broad definition can be interpreted in several nuanced and distinct ways depending on the notion of “capability to remain in a state” and “disturbance” that are used. The purpose of this section is twofold- given the large number of definitions and notions of stability in the ecosystem literature (Grimm and Wissel, 1997), it clarifies the interpretations of terminologies I will be using in the remainder of this document. Second, I trace the historical development of different concepts of stability, and justify why structural stability could potentially be an optimal way to describe the sensitivity of ecosystem dynamics to environmental perturbations, especially in complex nonlinear systems such as the microbiota. Finally, I also critically highlight shortcomings of the typical notions of structural stability preventing direct application of many traditional approaches to our system, as a means of motivating the developments that follow in the subsequent chapters.

Before proceeding to highlight conceptualizations of stability and select what metrics could be most illustrative, I first define sensitivity. For the purposes of this thesis, I define sensitivity of an ecological system as a measure of the *instability* of the system- therefore, a system that is more stable (in whatever sense that stability is defined) will also be less sensitive, and vice versa- one way to conceptualize this could be the inverse of stability. Therefore, outside this section, when referring to sensitivity, I mean the extent of stability of the system. The reason for this is that the metric I use (expanded in the next section) is inversely related to stability, and both intuitively and mathematically, it becomes easier to think about the effect of perturbations to the microbiota as being mediated by the sensitivity of the system- a more sensitive system is more affected by a perturbation. I now proceed to trace the development of notions of stability in ecological dynamics.

The notion of ecological stability, or the capacity for an ecosystem to remain in a particular observed state in the face of disturbances, has been one that has been of interest since

the development of ecology as a formal science, beginning with the ideas of Gleason pertaining to the role of stochastic environmental effects affecting trajectories of ecological succession and recovery, making changes in communities difficult to predict (Gleason, 1926). Around this period, mathematical models of ecosystem dynamics also came to the forefront with the work of Lotka and Volterra (Lotka, 1925; Volterra, 1926), who developed coupled differential equation based models for assessing predator-prey dynamics. These basic two-species predator-prey models predicted oscillatory behavior of the species involved about a central equilibrium point.

The Lotka-Volterra model was quickly generalized into community structures beyond predator-prey systems in the form of the framework termed as the generalized Lotka-Volterra (GLV) equations. The GLV equation models large interacting communities of species, by fixing intrinsic growth rates of each constituent species and their interaction strength with other species. Then, this allows one to track abundances of species over time as their growth rates are modulated by the presence of other species. The standard formulation of the GLV equation takes the following form-

$$\frac{d\vec{x}}{dt} = D(\vec{x}(t)) \times (r + A\vec{x}(t))$$

Here, $\vec{x}(t)$ denotes a column vector containing the densities of each of the species in the community at time t , r represents the intrinsic growth rates of each species in the community (that is, growth rates when grown in isolation), and A is the interaction matrix with dimensions equal to the diversity of the community, with element A_{ij} indicating the effect of species i on the growth rate of species j .

Depending on parameter choices, these generalized Lotka-Volterra models were able to demonstrate a wide range of possible system dynamics, ranging from convergence to a single fixed point of either competitive exclusion or coexistence (Levin, 1970), limit cycles (Smale, 1976), and chaotic behavior (Vano et al., 2006). Along with these developments, it was observed that simple models of population dynamics, even at the single-species level without the complexity of the GLV framework (such as the logistic map or the Ricker model (Ricker, 1954)) could display extremely complex dynamics by the appearance of bifurcations and limit cycles as internal parameter values were changed (May and Oster, 1976; May, 1976). Given these rich dynamics, a major line of inquiry that emerged alongside these theoretical studies was whether such trajectories were actually observed in real ecosystems. As more and more long-term ecological time-series were generated across a range of ecosystems and taxa,

these questions became of great interest. Experimental work (beginning with Gause, 1934 working on competitive exclusion), as well as field studies (Hassell et al., 1976; W.R. Thomas et al., 1980) showed a diverse pool of qualitatively different trajectories, including chaotic behavior (reviewed in Turchin and A.D. Taylor, 1992; Hastings et al., 1993; Rogers et al., 2022).

With these major developments in understanding the complex dynamics of ecosystems both real and simulated, parallelly, understanding the stability of these observed dynamics became a major area of focus. Broadly, the idea of assessing ecosystem stability is quantifying the capability of an ecosystem to return to its original state after a disturbance- i.e. the tendency of the ecosystem state to recover back to the original state after departure from it. These questions of stability became pertinent in light of both explaining the current observed states of ecosystems, as well as exploring the fate of these systems in the future under conditions of anthropogenically-induced environmental and ecological changes. Ecological communities in nature are constantly subjected to perturbations arising from both exogenous environmental changes as well as endogenous fluctuations from demographic noise (Bjørnstad and Grenfell, 2001)- this is true for communities in the microbiota as well (Nguyen et al., 2021a). These disturbances could have various effects on community structure, either in terms of modifying the population densities of constituent species and changing the ecosystem state (for example, depletion of abundance of a single species), or by perturbing the values of governing parameters (for example, a change in temperature affecting the growth rates) and thereby changing the shape of the vector field along which the population dynamics moves.

Today, stability is one of the most contentious topics in ecology, with many conflicting theories, ideas and definitions proposed over many studies. Some reviews have attempted to enumerate, group and unify these notions along common axes (Grimm and Wissel, 1997; Domínguez-García et al., 2019; Van Meerbeek et al., 2021). These reviews show many metrics that have been used to quantify stability in both experimental and theoretical work are based on responses of ecosystems to specific types of perturbations, or in terms of distance to a threshold of some critical transition. These metrics encompass a wide range of possible responses of the system to a diverse pool of changes, with many metrics developed specifically for a particular ecosystem or perturbation type.

One of the first major analyses of ecosystem stability using a theoretically motivated framework was done by Robert May (May, 1972). This was done by considering randomly generated

ecosystems (in terms of randomly sampled interactions), and given an equilibrium state of the constituent species, seeing whether a perturbation to species abundances from this equilibrium state will decay back to the original equilibrium using the local Jacobian of the equilibrium. Notably, May’s study found that as diversity increases, the likelihood of having a stable community tends to decline. However, empirical data from field and experimental work has shown that real ecosystems, with appreciable amounts of diversity, are indeed stable (Tilman, 1996; Ives and Carpenter, 2007; Pennekamp et al., 2018). Considering the counter-intuitiveness of this result in light of the high amount of diversity observed in nature, this sparked a large body of research in its wake on how large complex systems could be stable, referred to as the diversity-stability debate (McCann, 2000). This notion of stability is based on the capacity of the ecological community to return to a given attractor after a perturbation to the abundances of the constituent species (states)- in this sense, a stable system will return to the initial state after being perturbed, while an unstable system will diverge away from the state towards another one. This can be assessed by linearizing the system and finding its Jacobian at a given fixed point, and then finding its eigenvalues. If the system’s eigenvalues are all negative, then any perturbations in any direction will decay back down to the state at a rate specified by the eigenvalues. Thus, this indicates the stability of an ecosystem state to perturbations to the state variables (typically abundances)- we term this as **dynamical stability**.

One intuitive way of thinking about stability in a very general sense is to imagine the distribution of values in the state space taken by the system over time. This approach is particularly useful as it is extensible to cases where the dynamics are not single fixed points or limit cycles that take on a small set of values but rather chaotic complex attractors (of course, for a single fixed point, the distribution collapses to a single point of probability 1, with 0 everywhere else on the domain). Given this idea, we may think of a dynamically stable system as one whose distribution is unchanged by perturbations to the state variables (ignoring transient dynamics of settling at the fixed point). A dynamically unstable system which diverges from the fixed point will then display drastically different distributions.

1.3.1 Structural Stability

Another pertinent notion of stability is the system’s response to changes in the topology of the vector field that determines its trajectory through state space, a notion termed as **structural stability**. This is distinct from dynamical stability as structural stability measures the propensity of the system’s behavior to remain unchanged upon (smooth) external

perturbations acting on the *vector field*, as opposed to on the state variables. A structural perturbation changes the configuration of the vector field of the system; therefore, if the system is structurally stable, a small (smooth) structural perturbation to one of the parameters would result in a correspondingly very slight change in the topology of the vector field, and thus in the positions of attractors in the state space, leading to small deviations in the state accordingly. However, in a system with low structural stability, a small parameter change alters the vector space produced drastically, and thereby modifies the dynamics to be drastically different. In an ecological sense, a structural perturbation typically constitutes fluctuations in the parameters of the model (in a GLV framework, these constitute the growth rates and the interaction strengths, which in turn may be determined by abiotic factors). Given that environmental parameters that dictate ecological dynamics are constantly in flux due to a host of factors, structural stability also takes precedence as a critical notion of ecosystem dynamics that needs to be understood.

Once again, invoking the previously described idea of the distribution of states is very useful to gain an intuitive understanding here. A structurally stable system will have some amount of change in the distribution of states as the attractor gets deformed, but this will be a minor amount of change. If the system is structurally unstable, then a perturbation to the parameters will reorganize the distribution by an appreciable amount, possibly both in terms of position and shape.

The mathematical formalism for a dynamical system to be structurally stable was developed initially in the form of criteria for the qualitative behavior of a dynamical system to remain unchanged by smooth structural perturbations (Andronov and Pontryagin, 1937; Peixoto, 1962). Other significant mathematical theorems such as the Hartman-Grobman Theorem (Grobman, 1959; Hartman, 1960) developed conditions for when an equilibrium state is structurally stable to linearization, that is, its qualitative behavior remains unchanged. The development of bifurcation and catastrophe theory, where a model parameter is tuned up and down so as to observe changes in the number, type and stability of equilibria (Zeeman, 1977; Thom, 1994; Arnold, 1986), led to further advancements in our understanding of structural stability, including in biological systems (Lewontin, 1969; Recknagel, 1985; Perfect et al., 1990; Solé and Valls, 1992). These ideas have shown that the Lotka-Volterra predator-prey model is structurally unstable; for example, if even infinitesimally small second-order terms of self-limitation are added, the cyclical dynamics observed breaks, and the system

spirals down to a single stable point. While in general, structural stability is harder to assess than dynamical, with most previous work concerned with qualitatively ascribing stability/instability as a binary property of an attractor, there are many metrics that have been recently developed to assess it (such as the stability radius defined by Hinrichsen and Pritchard, 1986), especially in ecological contexts as the size of the feasibility domain or the local vector field divergence (Rohr et al., 2014; Saavedra et al., 2017; Cenci and Saavedra, 2018a; Song and Saavedra, 2018; Cenci and Saavedra, 2019). Recent work has also shown that the two notions of structural and dynamical stability are intimately interlinked in hyperbolic systems (Arnoldi and Haegeman, 2016; Ushio et al., 2018; Medeiros et al., 2021).

Going briefly back to our system of interest, the microbiota, recall that our aim is to understand the factors determining the sensitivity of microbial dynamics to *environmental* perturbations specifically, since these are representative of large lifestyle-related changes (as opposed to smaller-scale regular dynamical perturbations occurring as changes in abundances of specific taxa either due to demographic noise or via consumption of particular foods). Therefore, developing a methodology of assessing structural stability of microbiota states would be a natural and highly tractable direction to proceed in. Differing structural stability of microbiota over time and across systems could mediate differing sensitivities to environmental perturbations.

1.3.2 Non-equilibrium Notions of Stability

However, many pre-existing notions of structural stability cannot be directly applied to many real ecosystems, including that of the microbiota. This is on account of the fact that these metrics evaluate the stability of a given equilibrium attractor in the state space to perturbations. Given the continuously varying range of governing environmental parameters, many natural systems are never truly at equilibrium (Bjørnstad and Grenfell, 2001; Benincà et al., 2008; Sugihara et al., 2012), including the gut microbiota (Caporaso et al., 2011; Lozupone et al., 2012; Nguyen et al., 2021a), and given continuous topological deformations to the state space, identifying a single equilibrium state around which these perturbations are occurring is extremely challenging. Therefore, in order to understand the effect of external perturbations on the gut microbiota, we need to develop a notion of local structural stability of a state that does not explicitly require the assumption of equilibrium. Another further complication is the highly nonlinear and state-dependent nature of microbial dynamics- though linearization is often effective as an approximation, it does not work outside hyperbolic

attractors, and the state dependence of parameters means that parametrization from real data can often be prohibitively difficult.

Non-equilibrium notions of stability have also been developing over time in parallel with the other notions grounded in dynamical systems theory highlighted above. One commonly used approach that is agnostic to equilibria, used especially in the analysis of metacommunity models, is that of persistence (MacArthur, 1955; Levins, 1969). At the level of an ecological community, persistence may be thought of as the probability that a community can sustain positive abundances for all its constituent species over time. These early notions of persistence were further complemented by empirical work on energetic fluxes through ecosystems (eg. Odum, 1953). However, developing quantitative measures of persistence in complex ecosystems has proven difficult, especially given continuous structural and dynamical fluctuations. Though persistence has been linked with structural perturbations using techniques of structural feasibility to develop metrics of community stability to environmental change (eg. Saavedra et al., 2020), these are difficult to apply to diverse ecosystems, especially those spatially extended ecosystems structured to have a high immigration rate like the gut microbiota, wherein taxa that have undergone local extirpations can often be reintroduced very rapidly through consumption of food (Dey and Joshi, 2013). Further, persistence is majorly concerned with the feasible coexistence of this species pool- however, dysbiosis has been linked to changes in abundance of common gut flora, without necessarily implying local extinction (Winter and Bäuml, 2023). Therefore, measures of persistence may not necessarily give us the requisite insights into the sensitivity of microbiota.

Another frequently used non-equilibrium measure is the idea of ecological resilience, pioneered by the work of Holling (C S Holling, 1973). This refers to the system’s ability to withstand to remain in the domain of attraction of a particular attractor state (therefore quantifying in some sense the size of the attractor’s basin). It is worth highlighting that in ecological studies, resilience has been used to refer to both Holling’s definition as well as the rate at which a system variable returns to an initial state after a perturbation, called “engineering resilience” (Pimm, 1984). Since engineering resilience is similar to the notions of dynamical stability, henceforth, any references to “resilience” will be exclusively to “ecological resilience” (*sensu* Holling). Some studies have also termed ecological resilience (*sensu* Holling) as resistance. Resilience is a non-equilibrium conceptualization of stability, based on the ability of the system state to *withstand* structural perturbations, encoded as the size of

the domain of the attractor (as opposed to the previous notions of stability as the rate of the system returning to an equilibrium state). This nonequilibrium framework permits assessment of systems with multiple equilibria and various different types of attractors, by seeking conditions for when the system moves away from the observed basin.

Historically, resilience was typically qualitatively analyzed by building stability landscapes to describe dynamical behavior- however, recent work has made these ideas more quantitative, and has adapted to data from real ecosystems with a great degree of effectiveness in the form of attempts to develop early-warning signals (EWSs) that can anticipate when the system is in a state of lower resilience. Such systems will be more sensitive to shifting to an alternative stable state upon appearance of a sufficiently strong perturbation to environmental conditions (Scheffer et al., 2001; Scheffer et al., 2009). Given an ecological time-series of state variables in a continuously fluctuating environment, tracking early warning signals such as autocorrelation (Dakos et al., 2012), variance (Carpenter and Brock, 2006), skewness (Guttal and Jayaprakash, 2008), and kurtosis (Biggs et al., 2009) allows one to predict declines in system resilience before a critical shift to alternative stable states occur. These EWSs have been shown to predict tipping points in ecosystems with varying levels of success (O’Brien et al., 2023a), with different metrics having varying degrees of effectiveness in different systems. While computing such signals would certainly be one means of tracking the effects of environmental perturbations in community structure (as these track the system state observed as the structural parameters of external conditions vary), they are not always effective for multivariate environmental conditions or at the community level (Dai et al., 2015; Weinans et al., 2021; Baruah et al., 2022; O’Brien et al., 2023a). Despite these potential limitations, early warning signals could potentially be one way of tracking the effects of structural perturbations in the system.

However, it is worth noting that while EWSs are capable of detecting critical transitions in state space caused by structural perturbations, not every structural perturbation necessarily leads to a critical transition. Depending on the nature of the system’s attractors, a structural perturbation may lead to major structural reorganization of the state space via deformation of the attractor without necessarily leading to a transition of the state between basins of attraction. Therefore, in some sense, structural stability is a more generic metric of stability to parameter perturbation than resilience, since the system may not actually contain multiple equilibria. To reiterate, resilience tells us how much perturbation a particular ecological state

can tolerate before movement to the domain of attraction of another state, while structural stability informs us as to the extent of change of the topology of the field upon a small structural perturbation. A perturbation that leads to the state moving into the basin of another attractor will have very low resilience as well as structural stability; however, a system can potentially have very low structural stability but high resilience. As an example, consider an ecosystem such as microbes growing in extreme and highly stable environments such as deep under the ocean, with dynamics converging to a single stable fixed point of coexistence. While even minor changes in the environmental conditions could have adverse effects on population densities (indicating low structural stability), it is unlikely that the equilibrium state of the ecosystem as being populated with bacteria will change (indicating high resilience- the basin may shift and move, but the system remains in the basin). Therefore, especially since we have no a priori information whether perturbations in the gut are the result of deformations in the topology of a complex attractor manifold or critical transitions between alternative attractors, if possible, it would be more useful to track structural stability over resilience if measures of local structural stability that do not necessarily presuppose equilibrium can be used. Indeed, many recently developed tools for early warning signals have been utilizing metrics of structural stability as a signal alongside other classical metrics (O’Brien et al., 2023b; Grziwotz et al., 2023; Huang et al., 2024).

1.3.3 Other Stability Metrics

Therefore, to summarize, among the myriad metrics of stability, quantifying structural stability, or the extent of change in state space upon perturbations in the topology of the underlying vector field (manifesting from changes in the conditions of the external environment), is a promising direction forward. This could, however, prove to be challenging, on account of the high degree of nonlinearity and state-dependence seen in microbial systems, with an appropriately parameterized model formulation likely very complex and high-dimensional. Another significant challenge is to overcome the non-equilibrium and highly noisy nature of the dynamics. I describe methods to overcome these potential pitfalls, and develop a technique of assessing perturbations that is agnostic to model type and equilibrium state in the following section.

Regardless of the specific metric of stability used, there are many outstanding questions about the system dynamics that remain to be resolved. One particularly pertinent question that has attracted quite a lot of debate, especially in microbial systems, is the effect of community-level

biotic parameters such as abundance and interaction strength on the stability of microbial systems, given that abiotic environmental factors are constantly in flux and could potentially be the dominant force in structuring these ecosystems' responses. Assessing these ecological factors as well as the role of potential exogenous variables such as food intake will allow us to gain a deeper insight into the determinants of sensitivity to external perturbations. To answer these, we couple methods taken from the emerging field of empirical dynamic modelling along with tools from causal inference to robustly identify perturbations and highlight sources of variability in the responses of microbial communities.

Chapter 2

A primer on empirical dynamic modeling and causal inference

If you master mathematics sufficiently, you will find that you can make extraordinary predictions.

Noctae, Owlboy

2.1 Problems with parametric models

To understand and quantify the stability of microbial communities to environmental perturbations, it is important to reconstruct the ecological parameters of the system. With this, one may analyze how changes in the state and parameters of the system are associated with perturbations. Typically, this can be done by assuming the dynamics of the system follows a specified parametric model that approximates mechanistic processes in the system. Then, given observed data, identify the parameter values that best describe the observed data. Given the random noise present in natural systems, it is difficult to exactly determine the true values- however, several techniques can be used that provide the best fit of parameters that minimize error in reconstruction and prediction. This approach has been used with several microbial community studies (Descheemaeker and De Buyl, 2020; Joseph et al., 2020;

Gibson et al., 2021; Guk et al., 2022), with many using the generalized Lotka-Volterra (GLV) framework; some also have used the consumer-resource model (Ho et al., 2022). These simple models, grounded in ecological principles of interaction and growth, have been somewhat effective in describing broad patterns of species organizations and abundances in the system. However, given the complex and interrelated set of variables determining specific temporal trajectories, the high diversity, and the range of highly nonlinear responses across a range of scales, especially in the face of environmental perturbations, constructing low-dimensional parametric models of dynamics for such complex systems can often generate large errors (Wood and M.B. Thomas, 1999). One way to remedy this is to create extremely high-dimensional models with many covariates and interactions, which may be able to capture the full dynamics in a better way.

However, parameterizing such high-dimensional models accurately requires tremendous amounts of data, which may not typically be available for real ecosystems, especially if these parameters (such as growth rate or interaction strengths) constantly fluctuate over time, affecting dynamics in a nonlinear manner (Nguyen et al., 2021a). These driving parameters may also be interrelated in nonlinear fashions, further intensifying the complexity of the dynamics and making Lotka-Volterra pairwise approaches often ineffective (Momeni et al., 2017; Dedrick et al., 2023). The non-equilibrium nature of these systems also makes other approaches based on finding equilibria and modeling dynamics as perturbations away from that difficult, making efforts such as McBurney et al. (2019) to define a healthy microbiota state infeasible, especially if environmental perturbations make the position and nature of these equilibria constantly in flux. Even more complications emerge when considering which underlying model structure to use- in the absence of any other information, situations where the true functional form of the underlying dynamics is known are extremely rare, and selecting a model structure or complexity in the face of many mechanistically justifiable forms is difficult (Perretti et al., 2013).

Another major problem present in microbial systems in particular (depending on the model functional form selected) is the important effects caused by the presence of higher-order interactions (Ansari et al., 2019; Ludington, 2022; Morin et al., 2022), which adds another large set of parameters to be identified to accurately model and predict the dynamics- for a community of S species, there are S^2 possible pairwise interactions and S^3 third-order interactions. Naturally, for a high value of S , this creates a dizzying amount of parameters

to identify, which correspondingly requires often unrealistically long time-series. It is also worth noting that especially when using parametric equilibrium models on short time-series of nonlinear systems, the influence of drivers on the system can potentially be misidentified due to transient correlations (or lack thereof) that emerge when nonlinear behavior is seen, an effect termed as state dependency (Sugihara et al., 2012). For example, for the Lorenz system of ordinary differential equations (Lorenz, 1963), a 3-dimensional coupled system of variables X , Y and Z with a chaotic attractor (sometimes called the butterfly attractor on account of its shape), variables X and Z are positively correlated when the system state is in the “right lobe”, but this correlation inverts direction in the “left lobe” of the attractor. Given this, linear correlative approaches often produce incorrect or inconsistent results, especially if the system sampling is incomplete.

2.2 Introduction to Empirical Dynamic Modeling

In light of all of these challenges, using parametric models to analyze microbial community dynamics presents a major problem. One approach that has emerged as a means of combating these shortcomings is the idea of empirical dynamic modeling (EDM). This is a suite of “equation-free” non-parametric approaches to reconstruct system dynamics, that makes minimal structural assumptions about the system. Instead, it uses only the time series of observed ecosystem states to infer and forecast the dynamics of the system (C.-W. Chang et al., 2017). The term equation-free (or model-free) comes about due to the fact that there is no underlying parametric model structure assumed for the system- instead, the system is assumed to have some unknown underlying model guiding the dynamics, which manifests as empirical observations of the time series. This lack of structural constraints allows the bypassing of several of the defects with parametric models highlighted above. In this subsection, I extensively trace out the intuition and model framework used in EDM, building it up as a primer towards its utility in this study. Readers familiar with EDM or uninterested in the model framework can skip directly to section (2.2.4) for the link to structural stability and connection to the main question of factors affecting microbial community sensitivity.

2.2.1 Takens’ Embedding Theorem

The development of EDM as a technique that could be used to model high-dimensional nonlinear natural systems began with the development of the idea of state-space reconstruc-

tion. This stems from the observation that the attractor manifold of the system, the shape in state space traced out by the time-series of all state variables, was diffeomorphic (i.e. topologically equivalent under smooth transformations) to a “delay manifold” (or “shadow attractor”) constructed using time-lags of other measured variables as proxy state variables, a result known as Takens’ Theorem (Takens, 1981). While the proof of the theorem requires quite a bit of machinery from topology, I will justify the intuition and its relevance to modeling ecological dynamics here, and highlight the original paper by Takens as well as Deyle and Sugihara (2011) for proofs of a generalized set of theorems for state-space reconstruction beginning with Takens’ Theorem.

Intuitively, when looking at a (deterministic) n -dimensional dynamical system where state variables are coupled (say, abundances of interacting species), one may think about plotting the state variables as axes on a coordinate system- thus, the time-evolution of the system may be described by plotting vectors of the values of each of the state variables at different times in this n -dimensional state space. This traces out the surface of the system’s attractor, which can be described as a manifold (indicating that the system is locally Euclidean). It is also worth noting that projections of this manifold onto one of the coordinate axes will generate the corresponding time-series for that state variable (Packard et al., 1980). Now, because the state variables are coupled, the values taken by one particular state variable at a given timestep depend on the values of other state variables at prior timesteps. Therefore, the values taken by a single state variable over time are also, in some sense, encoded in the values of the other state variables. This observation forms the basis of Takens’ theorem. Since the variables are coupled, this allows the information on the time-evolution of one state variable to be represented as time-lags of other state variables. Takens’ Theorem shows that while the manifolds (shadow attractors) that may be reconstructed by using time-lags of state variables as coordinates in lieu of other state variables are not exactly the same geometrically, they can be smoothly (and invertibly) transformed by some coordinate change into each other (topologically equivalent). This means that local trajectories on the original manifold map bijectively onto local trajectories in the shadow attractor: the only difference is that they have been transformed by a smooth function.

As an example, consider an ecosystem where 3 bacterial species (say A , B , and C) are interacting with each other in a fixed, isolated environment, and we have measurements of only one bacterial species’ abundance over time (say, A). In this case, by lagging the

time series of this one recorded bacterial species by a fixed sampling interval τ and using E of its lagged values as coordinates $(A_t, A_{(t-1)\tau} \dots A_{(E-1)\tau})$, we may generate a new state space, and trace out a new manifold. By selecting appropriate sampling intervals (in most cases, anything positive suffices) and an appropriate number of time-lagged coordinates E (defining the dimension of the new state space), Takens' Theorem asserts that certain essential mathematical properties of the ecosystem dynamics can be preserved in this new manifold, accounting for the missing state variables. This includes the Lyapunov exponent, indicating that the qualitative dynamical behavior of the system remains preserved. The number of coordinates chosen, E , defines the "embedding dimension" of the shadow manifold. The optimal embedding dimension for the lagged state space may not necessarily be equal to the original number of state variables or the true dimension of the system- in fact, as I will highlight further, there are ways to select optimal embedding dimensions (Sugihara and May, 1990). While the original proof of Takens' Theorem was only shown for the case of time-invariant smooth manifolds (those that are locally similar to a vector space), further work extended these ideas to be true for the more general setting of forced fractal and stochastic attractors as well (Stark, 1999; Stark et al., 2003). Further generalizations have established the machinery for Takens' Theorem to be applicable when considering multiple datasets and continuous values (Sauer et al., 1991; Deyle and Sugihara, 2011; Gutman, 2015), making it viable to generate heterogeneous embeddings comprising multiple state variables with different lags.

This is an exceptionally powerful result in terms of understanding the dynamics of high-dimensional complex systems (even if the systems are dissipative), especially for those where it is challenging to specify the underlying dynamical function of the system. In real systems, it is often impossible to know the real dimensionality or the dynamical function of the system, and all of its driving variables (for example, the gut microbiota could be driven by complex interrelated variables pertaining to food intake and physiological conditions in addition to endogenous ecological dynamics). However, by Takens' Theorem, even from partial observations of at least a single state variable, it is possible to embed lagged coordinates of the state and reconstruct a topologically equivalent attractor describing the system trajectory. This is especially relevant in large ecological systems, where not only is the dimensionality high and very hard to measure.

Crucially, the construction of this attractor manifold allows the assessment of the local

trajectory at each time point of the system, allowing nonlinear nonequilibrium inference and forecasting. The first nonlinear forecasting methodology developed for empirically reconstructed state spaces is termed as “simplex projection” (Sugihara and May, 1990), and exploits the idea that since most manifolds of ecological dynamics are smooth (or approximable to smooth manifolds with a noise term, as per Munch et al., 2020), trajectories that begin from nearby points on the attractor remain close together for a short period of time. Therefore, simplex projection allows prediction of future values from the behavior of similar values in the past. While I do not explicitly use simplex projection in the analyses of microbial community dynamics to follow, building intuition for simplex projection allows a more nuanced understanding of further results that use the S-map algorithm, a closely related and more general method.

2.2.2 Simplex Projection

Simplex projection is a forecasting algorithm based on nearest-neighbours that entails tracking the forward evolution of geometrically similar points in an embedded manifold (paraphrased from Hsieh et al. (2005)). That is, past events that are similar to the current state in the topology of the reconstructed manifold are used to forecast future states. Suppose we have a point $\vec{X}(t)$, and aim to predict the state after some time T , that is, $\vec{X}(t + T)$. Then, we select the collection of n historical states nearest to $\vec{X}(t)$: let these be defined as $\{\vec{X}(t_1), \vec{X}(t_2), \vec{X}(t_3) \dots \vec{X}(t_n)\}$. Then, we may compute the states of this collection of points T steps into the future as $\{\vec{X}(t_1 + T), \vec{X}(t_2 + T), \vec{X}(t_3 + T) \dots \vec{X}(t_n + T)\}$. Given these future states of nearby points, we may now predict the future state of our focal point. The simplest way to do this is to simply average over the future states- thus, $\vec{X}(t + T) = \sum_{i=1}^n \vec{X}(t_i + T)$. Another alternative and more effective technique is to take the weighted average of each of these predictions based on the Euclidean distance of each of the neighbouring states to the focal point, so that more similar states affect the prediction more. This distance is rescaled into a distance kernel: one such popular kernel is exponential. If we have n nearest neighbours (call them $\vec{x}_1, \vec{x}_2 \dots \vec{x}_n$), we can compute their distances ($d_1, d_2 \dots d_n$) to the focal point, and then add a weight w_i of the i^{th} point as $e^{-d_i / \langle d \rangle}$, where $\langle d \rangle$ denotes the average distance from the focal point to the neighbourhood points. Therefore, $\vec{X}(t + T) = \sum_{i=1}^n w_i \vec{X}(t_i + T)$. We may think of this prediction as a locally weighted linear map for that focal point on the attractor. Also, in principle, by repeating this over a wide range of initial states, we get multiple estimates of the mapping from states $\vec{X}(t)$ to $\vec{X}(t + T)$ for different values of t . Interpolating among these, we may get an empirically derived discrete-time model for the system.

One important facet to look into is the number of nearest neighbours to select. Typically, for a manifold embedded in dimension E , $E + 1$ nearest neighbours are chosen (as this always defines a simplex polytope in E dimensions). Along these lines, another critical requirement is the appropriate choice of embedding dimension- though the method is not particularly sensitive to choosing a varying number of nearest neighbours, the distance between points is impacted by the choice of embedding dimension. Additionally, the choice of time-lag τ could also impact the reconstruction (though for discrete ecological time-series, picking a time lag of 1 usually suffices, especially when sampling intervals tend to be coarse). The idea here is that suitable τ and E values successfully “embed” the attractor, indicating that its trajectories in the reconstructed state space do not cross. This crossing can happen if the parameters are too small, and trajectories that are actually far away in higher-dimensional space are compressed closer together in low-dimensional space, producing false neighbours (Kennel et al., 1992; Abarbanel and Kennel, 1993). The optimal value of E can be selected by performing simplex projection over a range of values of E on a measured time series, then finding the optimal embedding dimension by comparing predictive skill (as a correlation ρ or mean absolute error) on the data. Briefly, this can be done as follows- the time-series can be split into a training and test set either by some fixed ratio or by leave-one-out cross-validation. Then using the training data, reconstruct the attractor, then make out-of-sample predictions on the unseen test data.

Simplex projection has been used as an extremely effective tool to forecast ecosystem states. For example, one notable use has been to predict fish populations for ecosystem and fisheries management applications (Ye et al., 2015). Another important application of simplex projection is to distinguish chaotic time-series from stochastically-driven time series (Sugihara and May, 1990). If a deterministic process produces chaotic behaviour, the deterministic nature of the dynamics produces patterns of response that simplex projection can successfully pick up on. This is not the case for systems driven solely by noise. The diagnostic test for chaotic behaviour is that as one increases the duration in the future at which to make predictions, the prediction skill should decline. This is not the case for stochastic systems, where the random nature of the dynamics dictates that prediction accuracy is independent of prediction interval. This ability to distinguish nonlinearity from linear stochastic systems is useful in terms of modeling nonlinear behaviours, as well as the aim of building a mechanistic model, which cannot be done if the system is linear and stochastic.

2.2.3 The S-Map Algorithm

One of the shortcomings of the simplex projection method is that the number of data points that it utilizes is limited, potentially constraining the predictive ability of the system. Moreover, if the noise in the system is autocorrelated, then the method of characterizing nonlinear and stochastic systems fails. One simple way to extend simplex projection is to remove the restriction of using only a small pool of nearest neighbours, and instead extend it to include all points on the attractor, all weighted by their distance to the focal point. This is the fundamental idea of the S-map algorithm (Sugihara, 1994): the S-map is a technique for sequentially global locally weighted linear regression. This can be thought of as stepping up the locally constant “zeroth-order” simplex projection algorithm into a first-order regression. In effect, the S-map computes this weighted linear map at each timestep on the attractor by weighting every other point on the attractor by its distance from the state at the given timestep. The S-map has one critical additional component absent in simplex projection: this is the addition of another parameter, θ , in the kernel along with the distance. The exponential distance kernel for the S-map takes the form $e^{-\theta d_i / \langle d \rangle}$. The exponential kernel is typically the most appropriate, considering information on the state of the system does tend to decay exponentially as one moves away from it in the case of Lyapunov instability (Cenci et al., 2019). However, depending on the nature of the nonlinearity of the system, other kernel types such as the Epanechnikov or tricubic kernel could also be usable (Hastie et al., 2001).

The θ parameter of the S-map is a metric for the nonlinearity of the system, encoding the extent of state-dependency of the system. If θ were set to 0, all points in the attractor are treated equally independent of distance- this can be shown to be mathematically equivalent to a linear autoregressive model. However, if $\theta > 0$, then the points have different contributions to the model modulated by their distance and rescaled by θ . As θ increases, only the nearest points have appreciable contributions to the forecast- thus, θ may be thought of as a “locality” parameter of sorts. This feature allows a more robust testing of nonlinearity versus stochastic behavior than simplex projection- by testing prediction skill for different values of θ , the optimal value can be ascertained. If this optimal value is 0, then the system is not distinguishable from a linear stochastic model (Deyle et al., 2013). This feature has been used to validate the belief that ecological dynamics in nature are intrinsically nonlinear (Hsieh et al., 2005; Glaser et al., 2014; Dakos et al., 2017; Clark and Luis, 2019).

The S-map constitutes an incredibly powerful framework for prediction, producing more accurate results than simplex projection in many cases. In fact, even when the correct parametric model is known for an ecological system, given the difficulty in fitting parameter values, the S-map algorithm can outperform the parametric approach (Perretti et al., 2013). Along with this, it also carries a significant and potent implication for reconstructing system state in addition to its utility as a prediction tool. Intuitively, the S-map algorithm fits a hyperplane to the system state at each point on the attractor manifold, connecting the focal point to its predicted future value. Formally, we may think of the S-map algorithm as solving the following minimization problem on each point on the attractor:

$$\hat{c} = \operatorname{argmin}_{c \in \mathbb{R}^d} \frac{1}{n} (Y - Xc)^T K_\theta(X, X^*) (Y - Xc)$$

Here, c denotes a row of coefficients of hyperplane slopes (which may be thought of as a Jacobian, as we will see), X is the data matrix and Y is the variable to be predicted. The $K_\theta(X, X^*)$ term is the distance kernel function ascribing weights, accounting for both the distance of points from each other as well as the “locality parameter” of the S-map, θ . It is important to recall that the distance here refers to positions in the state space and not proximity in time. This can be explicitly solved as an ordinary least-squares problem as $\hat{c} = (X^T K X)^{-1} X^T K Y$, where K is the weight matrix. Another way to think about the S-map is as a problem of singular value decomposition (SVD), with the coefficients C being the SVD solution to the linear model $B = A \cdot C$. Here, B is the vector of weighted future values, with $B_k = w_k x_i(t_k + 1)$, and A being the weighted current values as an $(n \times E)$ matrix: $A_{kj} = w_k X_j(t_k)$. While this is a simple framework based on regression, the sequentially localized fitting makes it highly tractable for these complex nonlinear systems. Other more complex techniques based on different model fitting techniques have been developing recently as well. Among these include generalized additive models (Benincà et al., 2015) and Bayesian frameworks such as Gaussian Process Regression (Munch et al., 2017). Additionally, older techniques used for nonlinear functional interpolation have been adapted to be more broadly applicable to the EDM framework, such as neural network based methods (Nychka et al., 1992) and polynomial response-surface models (Turchin and A.D. Taylor, 1992). The choice of model fitting technique depends on several data-determined factors such as the length, noise, and attractor shape. However, in ecological systems, the S-map has taken precedence as the primary forecasting framework on account of its minimal number of free parameters (only E and θ for a fixed sampling interval) and therefore higher transparency

and interpretability. Additionally, one feature that has been recently added to the S-map framework is regularization (Cenci et al., 2019) to increase the robustness of the S-map regression to process noise (Kugiumtzis et al., 1998)- the regularization scheme selected depends on the modeling objective, as well as the kernel used. With regularization, one may try a grid of kernel and regularization functions and identify which model structure is the most parsimonious while retaining a high level of predictive ability. The methodology of the S-map is summarized in Figure 2.1.




Figure 2.1: The S-map algorithm. Given time-series of state variables (here, abundances), we embed them as an attractor in state space. Then, by performing a locally weighted sequential regression, we may find the local Jacobian at each time point as regression coefficients.

From an ecological point of view, the S-map also has another major advantage in terms of system state reconstruction. Specifically, for an accurately fit S-map, the hyperplane coefficients denote the slope of one state variable with respect to another. This constitutes a Jacobian matrix- in this setting of discrete-time manifold reconstruction, for state variables $\{x_1, x_2 \dots x_n\}$, the Jacobian is a square $(n \times n)$ matrix with elements defined as follows:

$$J_{ij} = \frac{\partial x_i(t+1)}{\partial x_j(t)}$$

In practice, since the regression is done for each timepoint, the S-map generates a different Jacobian matrix for every point in the time series. Looking at the structure of a single element of a particular Jacobian matrix, it specifies the rate of change in the future state of the i^{th} state variable given the current value of the j^{th} state variable. Therefore, in an intuitive sense, this encodes the (varying) effect of the state variable j on the time-evolution of the state variable i over time. If we envision state variables as species abundances, this can

therefore denote a metric of interaction strength between the two species, since it indicates how much one species' population density is affected by changes in the other's density (Deyle et al., 2016). The Jacobian matrix at a system's equilibrium state has been used as an interaction matrix for many ecological studies as a metric of direct effect, beginning with Robert May's work on random ecosystem models (May, 1972), and continued into larger structured food webs at equilibrium (Bender et al., 1984; Yodzis, 1988; Schoener, 1993): the S-map thus provides a convenient way to model the time-varying interaction matrix between the constituent species along the attractor without necessitating the establishment of a single equilibrium state.

Before proceeding, it is important to note several distinctions from typical Lotka-Volterra community matrices that arise from using the Jacobian matrix in this setup-

- The Jacobian measures the direct effect of one species on the total *population* of another species. This is in contrast to the Lotka-Volterra type of interaction matrix, which denotes the average direct effect that a single individual of one species has on a single *individual* of another species (Laska and Wootton, 1998). Thus, the Jacobian looks at a population-level effect, versus the per capita effect of the Lotka-Volterra terms.
- The Jacobian matrix's interaction strength terms have been shown to be aggregate effect values of one species on another summing over all pathways, including higher-order effects; that is, interactions where two or more species affect another together (Song and Saavedra, 2021). However, the typical generalized Lotka-Volterra community matrix includes only pairwise interactions (though there have been extensions to higher-order settings, eg. Gibbs et al. (2022), these have separate pairwise and higher-order terms).

Another important distinction between classical equilibrium interaction models and the S-map Jacobian reconstruction in particular is that the values estimated in equilibrium models are constant over time, as the system is assumed to be fluctuating about an equilibrium state. However, the stepwise nature of the Jacobian allows computing this as a sequentially moving value over time- thus, interactions between species are not assumed to be static and can move over time, as is seen in natural systems (Deyle et al., 2016; Ushio et al., 2018; O.R. Liu and Gaines, 2022). Therefore, the Jacobian allows us to track species interactions as they change over time, providing a more complete picture of non-equilibrium nonlinear ecosystem dynamics.

2.2.4 Structural Sensitivity and the S-map

The reconstruction of the Jacobian from the S-map also provides us with an approach for assessing the system’s structural stability quantitatively. Recall from section 1.3.1 that we define structural stability as the deviation in the system’s state upon a perturbation to the topology of the underlying vector field. One illustrative conceptualization of this idea comes from thinking about the divergence of flows on the vector field (Cenci and Saavedra, 2019). At points where the local divergence of the vector field (that is, the volume density of the outward flux from the point) is low, a structural perturbation to the topology of the vector field will not have a major effect on the flow, as the extent of the flow “emanating” from that point is low. Conversely, for points with high local divergences, a structural perturbation here could strongly change the system state.

To draw a physical analogy, consider water flowing into a funnel, and water flowing out of a hose’s nozzle. Since the funnel is a sink and draws in flows of water, it has a negative divergence. If a small structural perturbation is made to the shape of the funnel, say, by slightly reducing the diameter of the funnel, the flow is largely unaffected. However, if one looks at the high-divergence case of the nozzle of a hose through which water flows, a source, a small perturbation (such as perturbing the pipe radius, or the shape of the nozzle a little bit) will lead to dramatic changes in the shape of the flow of water. Notably, for ecosystem dynamics, this variation in divergence happens all along the vector field within the same system, and thus, depending on the instantaneous state of the system on the vector field where the structural perturbation occurs, the effect on the system may be different.

For a vector field f , we may estimate the divergence $\nabla \cdot f$ as the sum of partial derivatives of the vector field’s components in each of the state variables with respect to the state variables, that is:

$$\nabla \cdot f = \frac{\partial f_{x_1}}{\partial x_1} + \frac{\partial f_{x_2}}{\partial x_2} + \cdots + \frac{\partial f_{x_n}}{\partial x_n}$$

In this setup, the unknown (local) vector field function f at a given time maps to the next timepoint’s state value (specifically, in this discrete case, $f(x_i(t)) = x_i(t + 1)$). Therefore, looking back at the form of the Jacobian elements above, we make the general observation that the trace of the local Jacobian for a given vector field is equal to the divergence of

the vector field at that timepoint. Given that the S-map algorithm allows us to compute a local Jacobian at each timestep on the attractor, this fact allows us to compute the local divergence at each time, and see how it varies. Considering the above idea that the local vector field divergence acts as an indicator of structural stability, this is a potential way to track the variation in structural stability of the system over time quantitatively.

Indeed, previous work (Cenci and Saavedra, 2019) has formalized this notion into a nonparametric technique to estimate time-varying structural stability. Formally, the divergence of the vector field may be thought of as the volume contraction rate (VCR) of the attractor—a point with a high rate of increase in volume diverges more, and trajectories here tend to spread out more. Since the divergence is specifically a property of the vector field and not the trajectories on it, this represents structural sensitivity (as opposed to any other metric such as dynamical stability). The VCR has been shown to accurately predict structural stability to perturbations, with states with higher VCR (that is, higher divergence) tending to have a greater extent of deviation from the system’s undisturbed trajectory if it were perturbed at that time. Therefore, the trace of the Jacobian reconstructed at each timestep gives us a quantitative measure of the extent of structural sensitivity of the system over time.

With this framework in place, I briefly make a detour from EDM methodology to justify why tracking the VCR could be effective as a metric of structural stability specifically for microbial communities. I posit that changes in the VCR are indicative of dysbiosis—drastic changes in the baseline relative abundances of microbial taxa, leading to physiological impacts. This is because a change in the value of VCR indicates changing sensitivity to structural perturbations, and therefore, greater/lesser capacity for fluctuations in abundances inducing dysbiosis. A structural perturbation occurring in a state with high VCR is much more likely to reorganize the microbiota into a different state, and therefore changes in the VCR act as a proxy for the chances of dysbiosis over time, especially given that environmental perturbations are constantly occurring. Also, since a structural perturbation deforms the attractor, this also modifies the states achieved on it, which have different intrinsic VCR values. Indeed, the VCR has been shown as a metric of predictability of future dynamical trajectories (Cenci et al., 2020): states with higher VCR tend to be less predictable. High VCR values thus indicate large degrees of fluctuation typical of host systems interacting with perturbative agents and therefore reorganizing their dynamical trajectories, which manifest as regions of lower predictability. A perturbation that reduces VCR also changes

the composition of the ecological network and forces it into a less sensitive dysbiotic state, which may persist and be difficult to recover from. Therefore, changes in the VCR clearly indicate periods of shifting dynamical regimes within the system- in effect, the local structural sensitivity tells us about the propensity for the current community compositional state (in terms of relative abundances) to change if a structural perturbation were to occur at that point.

Given an appropriately long and well-sampled dataset of species abundances in the microbiota, one major aim of ours is hence to track the variation in the VCR over time, as a proxy for structural sensitivity. However, it is not sufficient to simply track the structural sensitivity over time- another important facet is to understand what factors determine its variation over time. This variation in VCR, assuming it is not constant over time, could be modulated by exogenous factors such as food consumption, endogenous ecological community properties such as species abundances or interspecific interactions, or could be completely stochastic. One necessary clarification here is the interpretation of “effect of changes in species abundance on VCR”. This refers to the effect of the natural temporal changes in species abundance as described by trajectories on the attractor, and not perturbations to the state vector of species abundances at a particular time (which would be quantified instead through dynamical stability). This effect would indicate to us how the total abundance of the community as it fluctuates over time plays a role in determining the variation in VCR over time. Similarly, varying interactions over time (say, from competitive to weakly interacting to mutualistic regimes) as the state of the microbiota evolves along the attractor may have a profound effect on structural sensitivity. Since abundance data is used to reconstruct interactions via the Jacobian, which in turn gives us the VCR, these three variables are intimately linked- however, this interrelation may be highly nonlinear and hard to tease out.

In particular, the role of ecological interactions in structuring the observed relative abundances of microbial communities has been a question of long-standing interest. Many studies have shown interspecific interactions as being necessary to generate the observed dynamics and macroecological properties of microbial communities (De Muinck et al., 2017; Cosetta and Wolfe, 2019; Grilli, 2020; Ratzke et al., 2020; Gibson et al., 2021; Ho et al., 2022; Camacho-Mateu et al., 2024). Several of these studies have posited that microbial communities possess largely weak and sparse interactions (as expected of a diverse ecosystem), though their presence plays non-trivial roles in generating the final distributions. However, the overall nature, type, and predicted effect of these interactions inferred through these experiments

has often differed. While some studies (eg. Goldford et al., 2018; Kehe et al., 2021; Culp and Goodman, 2023) have shown that positive interactions such as cross-feeding are common among microbial communities, others have shown that competitive exclusion is highly likely (Foster and Bell, 2012; Coyte et al., 2015; Palmer and Foster, 2022; C.-Y. Chang et al., 2023). A third contrasting view on the role of interactions was shown in a recent study (George and O'Dwyer, 2023), which considered no interactions between species whatsoever, and instead used individual stochastic linear models of microbial growth. This model was able to successfully reproduce several macroscale properties of ecological communities such as the mean abundance distribution of species and the power-law relationship between abundance and its variance commonly called Taylor's Law (L.R. Taylor, 1961).

Though the presence of high numbers of cross-feeding interactions and competitive exclusions are not necessarily mutually exclusive results (for example, a community could have strong cross-feeding mutualisms between specific species and exclude a large pool of other species), characterizing the relative predominance of these effects on determining the state of the community over time could be illustrative. It is also worth noting that these studies have all analyzed these communities using parametric approaches such as the MacArthur Consumer-Resource Model (MacArthur, 1970) or the generalized Lotka-Volterra model, which, as highlighted before, assumed fixed interaction strengths over time, and could lead to unexpected errors in microbial communities. These also frequently do not account for higher-order interactions, which have been shown to play significant roles in structuring microbial communities (Sundarraman et al., 2020; Swain et al., 2022; Morin et al., 2022)-though these, too, have been inferred to be sparse (Arya et al., 2023).

Though I do not directly assess large-scale macroecological trends such as mean abundance distributions in these systems, quantifying the role of changes in ecological interactions on structural sensitivity could provide several valuable insights. Recall that in effect, the structural sensitivity tells us about the propensity for the current observed state of species' abundances to change if a structural perturbation were to occur at that point. Given that structural perturbations are constantly taking place in microbial systems (Nguyen et al., 2021a), the effect of interactions on structural sensitivity effectively tells us the role of a specific interaction structure in maintaining the observed abundance state, or potentially catalyzing or impeding recovery from a disturbed state. Given these potentially interesting community-level factors driving community-level sensitivity, I now introduce causal inference

methods as a potent way to identify linkages between them.

2.3 A brief foray into nonparametric causal inference

As highlighted in the previous section, in order to identify factors determining the sensitivity of microbiota to environmental perturbations, it is important to accurately reconstruct the variation in these properties over time and seek trends. For appropriately sampled abundance data, ecological properties can be reconstructed in a non-parametric framework that accounts for the high extent of nonlinear behavior using tools from empirical dynamic modelling like the S-map algorithm. Similarly, one may source data on exogenous factors like food intake and body mass from sampled metadata collected along with the experimental data. Given these variables, the aim is to look for linkages between them, and identify drivers of sensitivity. However, given the nonlinear nature of the dynamics in these systems, there may be potential periods where the sign and/or strength of correlation fluctuates over time, a phenomenon termed as “mirage correlation” (Sugihara et al., 2012). This can be especially prominent in ecological time-series, where the full dynamics on the attractor can sometimes take a timescale far beyond the sampling time to complete, leading to potentially spurious correlations if a particular sector of the dynamics is unsampled. Moreover, given limited ecological sampling data and a very large set of predictors, finding robust trends can be challenging. Another problem that emerges from nonlinearity is the notion of cascading effects- if a single small change occurs to a particular variable at a particular time, it may have cascading effects on linked variables that propagates through the system and causes drastic changes. Correlative approaches are hence ineffective to link system variables, and to truly estimate the effect of changes in factors on microbiota sensitivity, a causal picture is needed. In nonlinear systems, as relationships change in complex fashions over time, causation need not imply correlation.

Ideally, a dataset optimized for finding effects of system variables on a response variable would have randomized controlled trials and test interventions; clearly, this kind of data is often not realistically available in complex systems. This necessitates an approach grounded in harnessing observational data to make predictions about how changes in one variable could affect the response variable, the framework of **causal inference**.

Causal inference refers to the set of techniques used to identify the independent effect of one variable on another. The basic guiding principle of causal inference is to assess the response of a focal variable if one of the potential cause variable is changed, effectively

simulating an intervention in the system. One classical approach to this is path analysis, which identifies directions of dependencies among a large pool of variables to construct a network of responses. This was further extended to the field of structural equation modeling, which uses multivariate regression on a specified structure to identify causal linkages (Shipley, 2016). Both these methods are based on regressions, though, which, as previously highlighted, may not work for complex nonlinear systems. A more effective way is instead to use non-parametric probabilistic methods based in Pearl’s structural causal models, sometimes called do-calculus (Pearl, 2009). The idea of probabilistic structural causal models is to move beyond focusing on the effects of an intervention and instead see how much a potential cause can affect the probability of observing a particular value in the response variable.

Unlike the previous primer on empirical dynamic modeling, this section is much less exhaustive, and seeks to solely motivate the use of causal inference in nonlinear nonparametric systems, define relevant terminology, and introduce the requisite notions and tools that will be applied to microbial community variables. Interested readers may read further in several other extensive reviews of causal inference (such as Pearl (2009)), with others several specifically focused on ecological systems (such as Saavedra et al., 2022).

On its own, observational data cannot establish causal mechanisms- it can only find associations between variables. This is formalized in an idea called Reichenbach’s Principle (Reichenbach, 1956), which states that if two variables X and Y are statistically related, then there must be a third variable Z that causally influences both. Thus, Z takes the role of a confounding variable in this setup, and may be equal to X or Y itself, indicating that X and Y have a direct causal link between each other. But without any knowledge of Z , or when its effect cannot be blocked, in general, no further information can be determined about X and Y . Another possible complication is when there are colliders: two variables may be common causes of an effect, leading to some measures showing them as statistically related when measuring the effect. In a multivariate system, this can be exacerbated by multiple indirect and direct common causes and effects. To overcome this, a causal hypothesis that proposes possible underlying structure of the system’s variables and their relationship must be generated, which may then be tested.

2.3.1 Probabilistic causal inference

(about a causal inference enthusiast) he
got that DAG in him

@ethanblinck, X (formerly Twitter)

The basic idea of probabilistic causal inference is to generate interventional distributions—that is, to what extent can changes in a likely cause affect the probability of observing a given state of the response variable? This notion will be formally defined as the average causal effect or ACE (Pearl, 2009)—for now, one may informally think of it as if one were to perform an intervention and fix a system variable at a particular value, what is the probability that the response variable will produce a certain output? The ACE is computed for a particular linkage in a given causal graph, which represents the hypotheses to be tested. This causal graph typically takes the form of a directed acyclic graph (DAG). After statistical validation, the ACE can then be thought of as the weight of each arrow in this digraph. Note that the construction of this DAG does not assume a specific form of causal relation between the variables, and only the presence or absence of one. Therefore, this can always be used to describe a non-parametric system. One useful assumption to make, especially in noisy systems, is the Markovian nature of the causal graph, implying that the variables in the graph are affected by mutually independent random variables which are unknown and unrepresented.

First of all, a causal graph must be constructed, based on expert knowledge or causal discovery methods. Causal discovery methods aim to seek hypothetical causal structures from data. A commonly used algorithm for causal discovery is the IC algorithm (Inductive Causation) (Verma and Pearl, 1990), which looks at statistical dependencies between pairs of system variables in a range of possible conditions, then looks for likely causes among them using causal inference criteria. A brief outline of the algorithm is as follows—

- The skeletal structure of the causal graph as undirected edges is inferred using pairwise sequential tests of conditional independence based on conditioning on other system variables.
- These discovered edges are identified but a direction has not been established. To do this, colliders are identified: for each pair of separated variables X and Z with a

common neighbor Y , test if conditioning on Y lead to non-independence of X and Z .

- Finally, establish a direction according to the colliders and other conventions. (a specific convention is proposed in Verma and Pearl, 1990)

Then, ideally, if it is possible, a causal graph proposed as a hypothetical set of linkages must be tested. This can be done by checking for conditional and unconditional dependencies between variables that do not have a linkage between them (said to be d-separated). The independence of the two unconnected variables shows that the causal graph is indeed correct. If there is a linkage between two of the variables, this does not validate the graph as there could be other variables playing a role in this dependence. The structure of the validated causal graph can tell us about likely linked and independent variables- if the d-separations do not hold true, that is, the unconnected variables cannot be shown to be independent, then the causal graph must be redrawn and tested anew. If a likely link is identified by a validated causal graph, then it may be tested further using specific existence-based criteria to add more confidence to the causal effect- these levels of gradation of confidence are called as potential and genuine causes, and are based on the statistical dependence of the variables in different contexts, a set of variables assigned specific values.

Pearl's do-calculus framework provides a rigorous methodology to move from the causal graph and observational data to interventional distributions, by establishing rules for interventions mapping to observations. These rules and their axiomatic formulations are explained in detail in other sources (Pearl, 2009; Bareinboim and Pearl, 2012; Saavedra et al., 2022). I will briefly highlight the intuition and their implications here. An intervention on variable X (represented by a do-operator, producing $do(x)$) may be thought of as externally pinning a X to a fixed value x : if an intervention is done, then the causal graph is modified by removing all edges that point inward towards that variable X . This is because the values taken by X are now determined by the experimental intervention and not other system variables. Essentially, the interventional distribution can be thought of as the observational distribution in the manipulated causal graph. The Markovian property also allows the treatment of interventions as independent and local effects. Given this framework, do-calculus has 3 'rules' that describe conditions under which interventional distributions can be found on a causal graph by deleting/inserting interventions or observations, and exchanging observations and interventions with each other. I state the mathematical formulation of these rules in the appendix, but they are fundamentally based on conditional independence of variables in

causal graphs modified by specific interventions.

Now, using this, it is possible to calculate the average causal effect (ACE) linking two variables X and Y : this is formally defined as the expected increase in the value of Y per unit increase in X . If X and Y are binary valued variables, then the ACE denotes the difference between the probability of a particular outcome for Y if we intervene and set the state of X to either 0 or 1. Thus, in general-

$$ACE(X, Y) = \frac{\partial}{\partial x} \mathbb{E}[Y|do(x)]$$

And for binary variables, we see-

$$ACE(X, Y) = P(y = 1|do(x) = 1) - P(y = 1|do(x) = 0)$$

In effect, the ACE denotes the change in probability that Y happens if X changes, thereby indicating the strength of ‘effect’ between the two causal variables. It is worth noting that the ACE is calculated by considering all the direct and indirect paths through which X and change Y , including paths through mediator variables. This often obscures direct effects: to estimate these, the methodology needs to be modified to instead intervene on the mediators along with X . This leads to the definition of natural direct and indirect effects: the natural direct effect (NDE) between two variables X and Y is the effect of changes in the variable X on the probability of values of Y while keeping all the mediator variable values’ fixed. Similarly, the natural indirect effect (NIE) estimates the effect of the mediator variables on the response variable Y while fixing the other variable X . Mathematically, for binary variables, X and Y with a binary mediator Z as well, these may be expressed as follows-

$$\begin{aligned} NDE &= \sum_{a \in \{0,1\}} P(Z = a|X = 0)[P(Y = 1|X = 1, Z = a) - P(Y = 1|X = 0, Z = a)] \\ NIE &= \sum_{a \in \{0,1\}} P(Y = 1|Z = a, A = 0)[P(Z = a|X = 1) - P(Z = a|X = 0)] \end{aligned}$$

Any continuous-valued dataset could be binarized according to a specific rule: one commonly used one is to take above-median values as 1, and below-median values as 0 (Saavedra et al., 2022). Though this coarse-grains the system’s time-evolution, it permits direct estimation of causal effects in an interpretable manner. Moreover, this approach of binarization is frequently

used in the framework of symbolic dynamics to study nonlinear dynamical trajectories (Smale, 1967). In the next chapter, I describe the construction of a causal graph and binarization of microbial community state variables to predict causal effects according to this framework.

Recently, other methods of nonparametric causal inference techniques for nonlinear systems have also been developed. One class of these approaches is in the domain of inverse modeling (Ives et al., 2003; Almaraz and Oro, 2011). However, this requires a substantial amount of data to be applicable. Another approach that could be potentially effective is a nonlinear causation framework developed specifically in the context of empirical dynamic modeling and reconstructed attractor manifolds, termed as convergent cross-mapping or CCM (Sugihara et al., 2012). CCM is based on the concept of shadow manifold reconstruction defined by Takens' Theorem (see section 2.2.1). The basic idea is that because of transitivity of bijective maps between manifolds, a shadow manifold constructed from lags of one variable is also diffeomorphic to another shadow manifold of a different variable. From this equivalence, causality is established between a pair of variables X and Y by trying to assess prediction skill of current values of Y from time-lags of X (hence the name cross-mapping): if there is indeed some ability to predict the future values of one variable from past values of the other, then they are causally linked. Another additional criteria is the observation of convergence—that is, as more observation vectors are included and the shadow attractors become more dense, does performance tend to increase? While CCM could be potentially highly effective for identifying the causal effect of independent environmental variables on system dynamics, using it to infer linkages between specific ecosystem properties like species interactions could produce potential deviations due to indirect effects of interaction between species, which, as previously highlighted, are extremely common in microbial communities.

In spite of this, it is worth noting that several studies of ecosystem dynamics have used CCM successfully to filter down a fully connected interaction network to only include causal links between species' time-evolving trajectories. Then, the S-map was used to calculate only that subset of interaction strengths, along with additional CCMs to link computed ecosystem metrics to external environmental variables. Even in microbial systems, CCM could also potentially give similar or more fine-grained results to probabilistic causal inference and could thus be worth investigating as a parallel methodology to prune interaction networks and link external variables to ecosystem dynamics. However, in the absence of any prior information about microbial community response to perturbations and interac-

tion strength, it is more reliable to avoid potentially confounding indirect pathways from CCM.

In light of these potential shortcomings of other methods, and the broad general applicability and inference, I will be using the probabilistic causal inference method outlined above to link different factors changing over time in the microbiota to the response variable of structural sensitivity, represented as the VCR.

Chapter 3

Methods: Understanding Microbial Communities

Distortions. Interference. Real data is messy. It's all very, very noisy out there. Very hard to spot the tune. Like a piano in the next room, it's playing your song, but unfortunately it's out of whack, some of the strings are missing, and the pianist is tone deaf and drunk- I mean, the noise! Impossible!

Tom Stoppard, *Arcadia*

In this chapter, following on from the methodologies and model frameworks introduced in the previous chapter, I highlight the application of these techniques to large microbial communities (and associated challenges), and novel extensions I develop that facilitate understanding the occurrence of structural perturbations in microbiota and factors governing their response. All analyses are done in R, version 4.2.1 (R Core Team, 2022), using additional packages of rEDM for EDM functions such as the S-map (Park et al., 2023), deSolve for numerical integration of Lotka-Volterra dynamics (Soetaert et al., 2010), ggplot2 for visualization (Wickham, 2016) and zoo for time series manipulation and filtering (Zeileis

and Grothendieck, 2005).

3.1 Generating Sample Communities

While the approaches highlighted above mitigate several of the problems associated with the modeling of microbial communities, there are still certain associated challenges preventing direct application to microbiota count data. While ultimately, my aim is to understand the behavior of real microbial systems, it can be illustrative to generate a large number and variety of sample communities and analyze them, both as a means of understanding optimal approaches, as well as looking for trends between system variables that can then be validated on real data. One shortcoming of simulated data is that generating biologically realistic metadata on food intake and other health variables can be challenging (and therefore causal discovery is limited)- however, it is still feasible to use simulated data to interrelate fundamental ecological variables such as abundance and interactions with structural sensitivity.

Generating data that is reminiscent of natural nonlinear time-series can be extremely challenging. While one could potentially try EDM-based approaches such as constructing a family of Jacobians whose coefficients vary stochastically over time, then starting with an initial state and repeatedly regressing to generate sequential future states using successive Jacobians, this makes later S-map based inference circular, as the Jacobian is already known at each step. Thus, while developing methods with simulated data, it may be viable to use a simplified parametric linear model with noise, given its ease to simulate reliably and ability to generate a large set of underlying dynamics by modifying internal parameter values. Similar to several other parametric model-based work on microbiota (Joseph et al., 2020; Gibson et al., 2021), I use a stochastic generalized random Lotka-Volterra framework to simulate communities. However, it is worth noting that in the subsequent analyses, I use no properties of the parametric model itself in reconstructing the dynamics of the system. Critically, using the Lotka-Volterra model allows us to make structural perturbations to the system in the form of perturbing parameters such as growth rates and interaction strengths known to be causally related to abundances. In principle, any other model could be used equally successfully.

In order to make the generalized Lotka-Volterra model as similar to a microbiota as possible, and assuming minimal structure of the specific parameters beyond distribution parameters, I simulate large communities according to the following criteria:

- Reminiscent of real microbial communities, there must be a high amount of richness (I use 30 taxa).
- Growth rates are generally high, and sampled from a log-normal distribution (in accordance with the idea that individual abundances across a community are log-normal, as in Preston (1948)). This is making a large assumption that all species have positive growth rates on account of no resource limitation in the microbiota- however, this may not necessarily be true, and setting growth rates to negative is an alternative formulation that could be worth testing.
- Interaction strengths between species are randomly sampled from a Gaussian distribution, similar to random ecosystem models (eg. May, 1972).
- However, microbial communities have been shown to have high sparsity (Gibson et al., 2021; Camacho-Mateu et al., 2024). Therefore, a fraction of interspecific interactions are randomly set to 0 by setting a_{ij} and a_{ji} both to zero. This sparsity may also be emerging from the fact that when sampling dense interaction matrices for large communities, this does not account for spatial structure as a result of which species may not actually interact with each other. Microbial communities in the gut, for example, have been shown to have extensive spatial separation along the lumen, with specific taxa found only in one particular region and therefore not interacting with species local to another region. (Zhang et al., 2014)
- Additionally, to ensure that species abundances do not blow up to unrealistically high values during the simulation, a separate mean and variance are supplied as parameters to generate the diagonal elements of the interaction matrix. The entire interaction matrix is then divided by the square root of diversity to prevent large fluctuations when community size is large (Bunin, 2017).
- An initial transient state of all species having an equal abundance is defined, from which time-evolution towards the attractor rapidly occurs. The system is assumed to possess at least one attractor.
- The system is numerically integrated and solved for the next state using the generalized Lotka-Volterra function, filling in the parameters from the matrix and state abundances.
- At each timestep, (demographic) noise is added to each abundance N_i by randomly generating a noise parameter sampled from a normal distribution with mean 0 and

standard deviation ηN_i , where η is a fixed constant. This encodes stochastic birth and death processes. Any negative values created like this are set back to a small and positive value.

- At the end of each timestep, I also add in dispersal processes similar to what might be seen in the gut microbiota. First, I select a random number of species between 1 and half the maximum diversity (15), and increment their abundances by values chosen from a uniform distribution from 0 to some upper bound (typically 0.3). This represents stochastic immigration of certain species based on food choices made by the individual, which add abundances. Then, to represent the loss of bacteria from the microbiota by cyclic processes like digestion, a fixed fraction of all species is stripped away- stool is an excellent predictor of gut microbial composition, hence, an equal fraction should be represented as well (Vandeputte et al., 2016).
- This process is continued stepwise for a set number of days. After this, a structural perturbation is simulated by perturbing each growth rate and interaction strength by adding a randomly sampled term from a standard normal distribution to them. Another alternative is to perturb a single growth rate and single abundance value.
- The simulation is continued with this new set of parameters for a set number of days- then, the parameters are set back to their original value, denoting recovery.

This generates a long time series with 3 regimes- the initial phase before the perturbation, a perturbed phase where the parameters are changed, and a ‘recovery’ phase where the parameters are set back to their original values. I generate multiple time series with multiple different sets of parameters from this algorithmic framework. Notably, this setup allows us to create communities with different intrinsic properties along various axes- for example, slow-growing communities with high degrees of competition in a very sparse interaction matrix subjected to small-size perturbations, versus fast-growing communities where all interactions are weak, though non-zero, and the perturbation to parameters is large. This diversity allows easy comparison of different quantitative and qualitative results to different regimes.

Since we now have the abundances sampled at continuous time-steps, the dataset is now ready for the application of empirical dynamic modeling. The first step is to make sure that each species sampled has “equal weight” when being compared in the reconstruction- this is done by rescaling each time series of abundance so that they all have mean zero and

variance one. This prevents unnecessary distortion of the state space. Now, in principle, we can apply the S-map to the time series and reconstruct the Jacobians. However, there is one significant limitation to doing this. Recall that the S-map is based on performing a weighted regression, with the weights determined based on the distance of the focal point to others. If the dimension of the state space is too high, Euclidean distances tend to blow up to extremely large values, a phenomenon termed as the “curse of dimensionality” (Verleysen and François, 2005). As a consequence of this, distances inferred are very high, reducing the weight contribution of even nearby points significantly, and making them more similar to faraway points in an exponential kernel. Clearly, this is not a feasible strategy, especially since microbiota tend to contain a very high species richness, and therefore will exist in a very high-dimensional state space if directly reconstructed.

In general, the application of EDM techniques to large ecosystems has been challenging. While some techniques have been developed as potential workarounds (Ye and Sugihara, 2016; C.-W. Chang et al., 2021), these often are sensitive to noise. To overcome this, I propose a simple solution that allows forecasting. This approach involves randomly sampling a large number of small sub-communities from the species pool, reconstructing interactions between them, and then putting them together to analyze system-level responses. I pick the size of each sub-community as 4 to capture sufficient complexity without becoming too large- however, most subsequent results are robust to larger cluster size choices. The basic notion here is to treat all other species as part of the external environment along with other drivers. Therefore, fluctuations in other species also may constitute structural perturbations at the level of this single sub-community; but note that these may not be true structural perturbations at the whole community level, and thus should not be considered directly at face value as periods of external change. This sub-sampling approach has been used effectively to study microbial community assembly and transitions in large species pools as well (Long et al., 2024). Another caveat of this approach is that given the fact that the composition of the environment and the nature of its fluctuation varies in each sub-community (and thus potentially the nature of interaction between two same species in different sub-communities), it would not necessarily be accurate to use these values to fill in values of a larger interaction matrix. However, if sufficiently large number of communities is sampled, this generates a dense distribution of interaction coefficients at each time step for a pair of species, over a range of external environmental conditions. Then, by picking a parameter for this distribution (such as the mean or median), this could produce a coefficient that could be supplied into filling in a larger

Jacobian. This could be an alternative approach to Jacobian reconstruction, though computationally heavy, and has been used in other studies of large microbial communities (eg. Deng et al. (2021)). However, this approach of generating the entire Jacobian notably does not give us insights into variation in response between different subcommunities, which, as I will show, gives critical mechanistic insights into the drivers of response. Furthermore, at high levels of diversity, microbial ecosystems display behaviour indistinguishable from random fluctuation (Cui et al., 2021); therefore, specific responses may be visible only at smaller dimensions.

I note the synergy between this approach of sampling small communities with the technique of multiview embedding (Ye and Sugihara, 2016). However, since clinical interventions in the microbiota are typically done at the level of introduction of a few specific species, we believe that working at the level of single small communities and seeking specific response types characteristic of a particular community structure could provide more targeted insights in terms of recovery from dysbiotic dynamical states. Further, since the samples of bacterial OTUs come from fecal samples, the communities that we may sample are largely ad-hoc and may not actually spatially coexist in the alimentary canal (Zhang et al., 2014); therefore, using multiview embedding to combinatorially generate and average out over the communities may end up including unrealistic communities in the mean prediction; instead, by treating each community individually, realistic communities that have strong interaction effects on each other may be isolated.

To compute Jacobians of these sub-sampled communities using the S-map, a first step is deciding the parameters to use. Since we fix a discrete sampling interval of 1 day to achieve synergy with the form that real ecological data takes, we fix the time step as 1. Then, the next parameter to be determined is the embedding dimension. In order to be able to consistently compare results across communities, I fix the embedding dimension as 4 for all communities, that is, the state space is comprised of raw values of each of the state variables, with no lagged terms. This standardizes the axes and ensures that any two sub-communities' state variables and dimension will be necessarily translatable with each other. Finally, the last parameter to be determined is the locality parameter of the S-map, θ . This is determined by taking a large range of possible values of θ (from 0 to 9) and comparing the root mean squared error in different cases by using the leave-one-out cross-validation method. A different value of θ is selected for each community sampled, as each community may have a different intrinsic amount of nonlinearity and state-dependence.

With all this in place, from a large number of sampled sub-communities of 4 species each from a pool of 30 species, using the S-map algorithm, the time-varying Jacobians are reconstructed. From each of these communities, then, I extract the trace of each of these Jacobians, defining a time-series of the volume contraction rate (VCR). Similarly, to quantify interaction strength between taxa, I define the average interaction strength between taxa as the mean off-diagonal Jacobian value. Though this tends to average out over quite a bit of variability, especially considering that a case of symmetric opposing effects between two species (i.e. $J_{ij} = -J_{ji}$) would be set to zero, it tells us quite a bit about responses at the whole-community level. An increase in average interaction strength would thus indicate that the community as a whole is moving towards more positive interactions, while a decline indicates more negative effects at the whole-community level. Another option is to evaluate the absolute interaction strength variation over time, which looks solely at the coefficient and not the sign, but this can be hard to resolve into actual mechanistic changes to the community. Regardless, I propose this as an alternative metric to also track interactions. Finally, another option is to look into the numbers of different interactions, but this tends to be extremely coarse-grained as it may not reflect the high degree of parametric variation, especially given 4 species encoding a possibility of 6 interaction pairs (and hence capping the maximum possible number of a given interaction type as 6). Thus, to quantify the effects of interaction on sensitivity, I assess how changes in the mean off-diagonal interaction coefficients affect the trace of the Jacobian.

One clarification that could be useful at this juncture is a slightly more detailed explanation of the interpretation of interaction strengths in a Jacobian setting. As highlighted in Section 2.2.3, the Jacobian is a measure of the local aggregate effect of on the population through all pathways (Song and Saavedra, 2021). In light of this, a (+/+) interaction, as seen in the Jacobian, may not be indicative of a pairwise mutualistic effect between the species, but could also potentially be the result of indirect pathways leading to positive net effects on each other without direct interaction. For example, consider a scenario where bacterial taxon *A* releases a metabolite which can be used by taxon *B*, but taxon *B* does not directly confer any benefits to *A*. Thus, the pairwise interaction is commensal (+/0). However, the presence of *B* may compete with another taxon *C* and stymie its growth. If *C* is a competitor or predator of taxon *A*, then the presence of *B* aids the growth of *A*, leading to an indirect benefit for *A*. Therefore, this would be encoded as a mutualism in the Jacobian, without actually taking the form of a pairwise facilitative interaction. Note that while henceforth, I

refer to (+/+) interaction terms in the Jacobian as mutualisms, it is important to remember that this need not be a pairwise mutualism. In some sense, the Jacobian is therefore able to give a more complete picture of the effect of one species on another, while sacrificing the simple interpretability from pairwise interactions. Another benefit of using the Jacobian interpretation of interaction strengths here is that we would expect some sort of relationship between the trace and the interaction strengths a priori- since these encode all indirect pathways as well, the trace elements include not only the self-limitation on the species but also the effect of other species on the focal species through the pathway of its own effect on other species. This also further justifies why correlation-based approaches may not be particularly useful here, since the inter-specific interaction terms are also part of the trace.

Having generated these, I can summarize these as large time series of abundance of each of the sampled communities, their mean (interspecific) interaction strength, and VCR. By plotting each of these over time, I can look into how a community's abundance, interaction strength, and structural sensitivity varies over time, especially during the period of perturbation. However, each of these characterize responses for a single 4-species community: by construction, the external environment and species composition of each of these differs, and the effect of the perturbation differs for each of them (as a random pool of species is randomly perturbed). Additionally, in real data, there may be some other selective external perturbations that affect only some sub-communities and not others. Given this, how do we know if the perturbation has affected the whole community?

The first thing to establish here is a metric of community change when a perturbation emerges. As proposed earlier, the structural sensitivity, quantified as the volume contraction rate (VCR) from the Jacobian, should be able to detect induced structural perturbations by tracking the changing divergence of the vector field when a disturbance occurs, indicating deformations in the attractor. However, depending on the species in the sub-community and random nature of the perturbation, this VCR value should differ strongly across communities. If this is indeed the case, how can we scale up and see whether the VCR significantly changes at the whole community level during the perturbation period? One obvious solution is to simply average the VCR values across each community to get an average structural sensitivity; however, this does not work for the inferred VCRs. This is because the S-map is fundamentally a regression algorithm that does not return the true Jacobian, only one that is directly proportional to the real one. Recall that in order to treat every variable as equally important

and not deform the attractor on account of relative abundances, the S-map requires each abundance time series to be normalized to a standard normal distribution. This results in loss of true interaction strength values, providing instead relative values of slope. This normalization factor is specific to each constituent taxon and appropriate rescalings combining all terms cannot be determined easily and reliably. Therefore, summing up values of VCRs across different communities is not viable, necessitating a different approach. Finally, given a metric such as the VCR that could track structural changes integrated to the community level, we may rephrase our main question into a more computationally tractable and more high-resolution form- **what factors determines the variation in response of each sub-community to perturbation?** That is, in the face of random external perturbations, are there certain common characteristics that lead to certain sub-communities responding with greater departure from their pre-perturbation states, indicating more severe dysbiosis?

3.2 Developing the Similarity Statistic

To quantify whether VCR changes (or lack thereof) occurring in particular sub-communities are reflective of changes in structural sensitivity at the whole-community level, and to compare responses across different communities, I develop a novel z-score based approach that is highly extensible and general as an indicator of perturbation for a wide range of time series and parameters. By scaling up to the whole community level, applying this to the VCR from simulated data simultaneously acts as a test for the effectiveness of the VCR as a metric in logging recorded structural perturbations. Given natural noise and nonlinearity in the system, it is hard to characterize a specific value that is indicative of a pre-perturbation state. Therefore, the basic idea is to ‘learn’ a baseline distribution of values that categorize the pre-perturbation ‘healthy’ attractor, and sequentially compare future values of VCR against this distribution. Then, to scale this up from a time series of sub-communities to the response of the whole community, I track the fraction of sub-communities at a particular time point that deviate by at least a certain amount from this distribution. I describe the method in detail below-

1. Fix an initial window length B to define the baseline pre-perturbation period.
2. Fix a size r_s for the rolling window. An ideal size would be one that generates at least 30 different windows on the baseline period (for example, if a baseline of 50 days is chosen, then the maximum window size that would be appropriate is around 20 days).

The reason for this is because 30 is generally regarded as an estimate for the Central Limit Theorem to hold, indicating that the baseline distribution may be treated as Gaussian (Ross, 2017)

3. Given a time-series of a specific time-varying parameter of a single sub-community (say, the VCR of the community), roll a window of size r_s over the baseline period, and calculate the average value taken by the time series in each window. From this, generate a vector of $B - r_s$ values, and compute the mean, μ_b and standard deviation σ_b of this. These act as our baseline distribution parameters.
4. Then, roll the window of size r_s over the remainder of the time series outside the baseline period, computing the successive sequence of mean parameter values. At each point, compute the z-score of the observed mean parameter value, μ_s , with respect to the distribution of the baseline as $(\mu_s - \mu_b)/\sigma_b$.
5. If the z-score is greater than a cutoff value K , set its value to 1. If it is lesser than $-K$, set it to -1. Else, set it to zero.
6. Repeat this classification for the entire time series to generate a sequence of 1, 0 and -1 over time specifying the z-score of the window compared to the baseline.
7. Repeat this process for all the subcommunities generated.
8. Finally, at each timestep, compute the fraction of sub-communities that have a value of 0, and express this as a time series. I define this value, the fraction of communities with a z-score within a specified cutoff, as the **similarity statistic**.

This process is summarized in Figure 3.1.




Figure 3.1: The Similarity Statistic. Given a family of time series, we define a baseline period, and generate a distribution of values by rolling a window over it. Then, we roll the same window over the outside region, and compute z-scores. Rescaling z-scores according to a cutoff, the final similarity statistic is what fraction of time series have a z-score within the defined cutoff as a function of time.

The reason I use the fraction of z-scores as opposed to averaging raw z-scores is because baseline distributions of low variability can produce very high z-scores, which can skew the mean when comparing across communities. Another option would be to use the rank of the window's value compared to the baseline values, but this can be very coarse. Intuitively, the similarity statistic measures the likelihood across all communities that a particular day's observations could have come from the baseline 'healthy' distribution. If a large fraction of communities deviate significantly from the baseline distribution's bounds on a particular window (generating a low number of 0 scores), it is likely that this is a period of perturbation affecting the system at the whole-community level. Another way to conceptualize this is that the baseline state consists of a section of trajectories on an attractor. The z-score indicates

the possibility that a future value could have come from that same shape of attractor, or whether a perturbation is likely to have deformed the attractor. By assigning a score of 0, this is a means of indicating that there is statistical support for the fact that the observed value is from the same attractor, and thus the system has remained unperturbed. Finally, the fraction of zeroes takes it to a sort of mean-field response by encoding what percentage of communities are not significantly perturbed. This technique is somewhat similar to the EDM-based approach proposed by Säterberg and McCann (2021), which uses prediction accuracy of one regions' reconstructed attractors on another as a metric to predict regime shifts; however, while that compares between two segments of the time-series, this method provides a continuous measurement of similarity over all windows. I propose this as a highly generalized and robust method to detect periods of differing dynamics from a chosen baseline.

While this is a complex and computationally intensive metric to apply, it has many significant advantages. One major advantage is its generality- since it assumes no specific parameter and requires only a family of time series, it can be applied anywhere to identify perturbations in that parameter. The reason that the VCR is of specific interest to us is because changes in the VCR are indicative of structural sensitivity and associated environmental perturbations. However, the same method can be used to track whole-community changes in abundance and interaction strength as well from smaller sampled subcommunities. Similarly, for a different ecosystem response variable such as, say, lake turbidity or biomass, this method can be easily and directly applied. Another major implication is that it allows us to test the exact effective duration of perturbation and recovery periods on the system: given a recorded perturbation, does the whole system attractor on average return instantaneously to its original state after the perturbation, or does it persist in a new and deformed state throughout the perturbation? Further, after the departure of the perturbation and (potential) restoration of the original parameters, does the system take longer time to recover, as a sort of hysteretic response? Note that this may be the case for specific sub-communities, but this may not be true at the whole system level. The similarity statistic of the VCR tells us about the effective duration of the structural change on the system's sensitivity.

Taking this further, since the system picks up statistical changes in state without specifying any properties of the perturbation itself, it could be used to identify structural perturbations in different time series, as a sort of early-warning signal indicated by initial depletion in the similarity statistic value. However, on its own, these may not be direct identifications of

structural perturbations, but rather represent dynamical perturbations (where one or more of the bacterial abundances undergoes a large change). These could also change the VCR (as it is ultimately computed as a nonlinear function of the abundance). Moreover, the dynamical stability may also be calculated from the Jacobian as its maximum eigenvalue (Ushio et al., 2018). Because of this, and the trace being the sum of all eigenvalues, the two measures are well-correlated. The two measures are intimately interlinked (Arnoldi and Haegeman, 2016; Medeiros et al., 2021)- regardless, this provides a pool of potential perturbations to be further probed for causative agents or major compositional changes using other methods.

3.3 Causal Inference on Sub-Sampled Communities

To test the effect of abundance and interaction changes on the variation of the VCR in these sampled communities, I follow the causal inference methodology introduced in Section (2.3.1). Unfortunately, since there is no metadata in these simulated communities, I do not use causal discovery, and instead hypothesize a causal graph based on the 3 main ecosystem variables of interest- the abundance, the mean interaction strength, and the VCR. Specifically, abundance and mean interaction strength are both posited to have a causal effect on the structural sensitivity. In addition, I hypothesize the presence of a causal effect of abundance on interaction strength, something typically not accounted for in ecological models, but could have potential effects. For example, in fish, predator switching behaviour has been observed where the existence of predation depends on the abundance of one prey compared to others (Murdoch et al., 1975). Therefore, as abundance varies, the interaction coefficients may vary as well. However, we do not quantify this effect explicitly, and instead concentrate on the effects on the VCR. Unfortunately, in this causal graph structure, there are no d-separations, as all variables are linked by at least one edge. This prevents statistical validation of the structure. This structure is shown in Figure 3.2.




Figure 3.2: Structure of causal graph. The mean interaction strength acts as a mediator variable. To compute the effect of these terms on the structural sensitivity, we compute the natural direct and indirect effects, which compute the probability of an observed change in the sensitivity when one variable changes, holding the other one fixed.

Here, given the structure of the causal graph, the interaction strength acts as a mediator—thus, to compute the causal effect, I use the Natural Direct and Natural Indirect Effect values respectively defined in the earlier section to compute the effect of abundance change and interaction change on the VCR. To do this, I first detrend each time series by taking successive differences. This mitigates the effect of specific higher/lower values when binarizing, ensuring that the factors considered are not the raw values (which may carry historical contingency from the past) but only the fluctuations. Then, binarize each time series for each subcommunity about their medians. Then compute these effects— in effect, this tells us if there is an above-median or below-median fluctuation in abundance/interaction strength while holding the other value fixed, what is the probability that the VCR will show an above-median fluctuation in value? This produces an NDE and NIE for each sub-community. To analyze these further, I assign these sub-communities into profiles. If both the natural (direct/indirect) effect of abundance and interaction strength on VCR is seen to be positive for a sub-community, then I assign it to the profile (+/+), or PP (for plus/plus). Similarly, I define communities in the profiles PM, MP, and MM. With this, one immediate question that follows is the relative frequency of these different profiles.

These causal profiles therefore represent cases where the ecological community properties have different roles in structuring the system’s sensitivity over time. To take this further, I investigate whether these causal profiles could potentially determine the variation in response to the perturbation among different sub-communities. To do this, I collect the VCRs of the sub-communities associated with each profile as four families of time-series, and apply the similarity statistic onto each of them. Then, I compare the similarity statistics of each of these profiles over time to each other to look for differences in response. This constitutes a complete work flow on the simulated data, starting from generating the data, looking for the existence of perturbations, classifying variation in sampled sub-communities into profiles, and finally comparing these profiles’ response against each other. Given this pipeline, I analyze simulated community data for a wide range of parameters, and compare the responses and results. However, parallel to this, these methods need to be tested if they are applicable and successful on real-world microbial data. This approach is summarized in Figure 3.3.

3.4 Parameters for the Pipeline

Note that in the aforementioned methodology, the parameters of the generalized Lotka-Volterra are only used as a guideline to generate the dataset, and none of the properties are used in evaluating the results, which are all done using non-parametric methods assuming no underlying systemic structure. Regardless, the underlying parameters supplied produce qualitatively different datasets, which need to then be carefully analyzed for differences in behavior. In particular, the main determining parameters for Lotka-Volterra dynamics are the growth rate (r) and the interaction matrix (A). In addition, other parameters of this specific procedure can also be modulated, such as the amount of noise, strength and duration of perturbation, extent of sparsity, and rates of immigration and emigration. Also, given the stochasticity inherent in the system in terms of interaction structure and the random nature of perturbation, it is theoretically possible to have different outcomes and states while specifying the same set of underlying parameters.

Therefore, for each set of parameters that determine the structure of the model, I simulate 10 different communities and look for differences in results. If for a particular set of values, a complete time series featuring an initial phase, a perturbed period, and a subsequent return to the initial phase is not generated after 100 iterations, an error is raised and that set of parameter values is excised from the analysis. Because of the high complexity of the

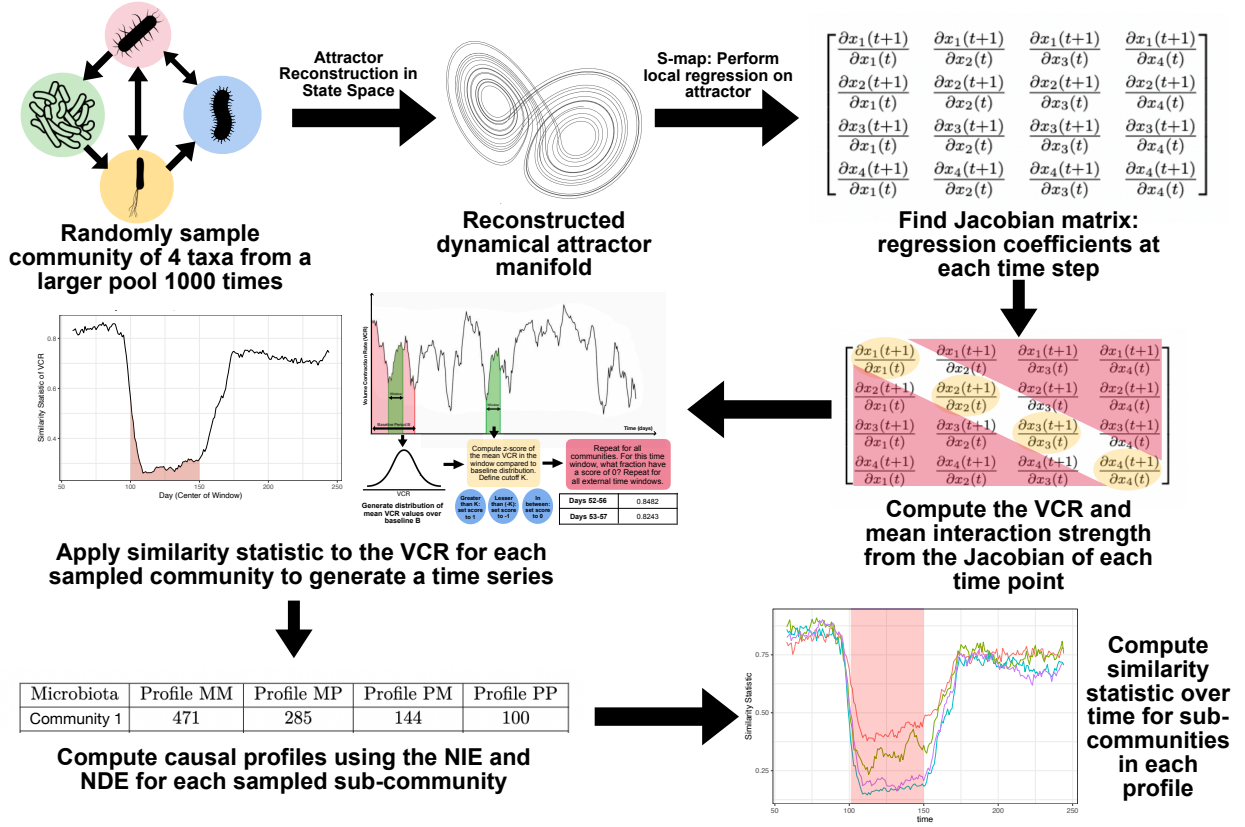


Figure 3.3: Complete workflow. Beginning with raw abundance data, I reconstruct trajectories for subsampled communities, compute ecosystem parameters, and then apply the similarity statistic and causal profiling.

model, I fix several of the parameters during the simulations. The initial phase is set to 100 days, and a perturbed period of 50 days. To achieve synergy with real data which is sampled on a daily basis, the value of the time step is also set to 1: thus, a complete time series contains 251 days of data. The species pool size is set to 30 with all species starting with an equal abundance of 0.5. The standard deviation of growth rates is set to 2 (i.e. growth rates are sampled from a log-normal distribution with standard deviation $\log(2)$), and the mean self-limitation (diagonal) is set to -4 with a standard deviation of 1. The off-diagonal elements of the interaction matrix are set to have a standard deviation of 1 as well. The noise fraction is set to a constant value of 0.2, and the standard deviation for the perturbation distribution is set to 0.3 (the mean being zero). Finally, the maximum immigration is set to 0.3, the sparsity fraction is set to 0.2 (indicating that 20% of pairs in the interaction matrix are randomly set to 0), and the emigration rate is set to 0.75 (indicating one quarter of the population is flushed out at each time point). Given this, in effect, the only parameters that

are left free are μ_r , the mean growth rate whose log specifies the mean of the log-normal growth rate distribution, and μ_A , the mean of the interaction matrix which specifies the center of the Gaussian distribution of interaction strengths. A grid of these values is thus iterated over, and the time series is attempted to be generated in each case. The grid of values of μ_r and μ_A used is provided below in Table 3.1-

Community	μ_r	μ_A
1	3	-4
2	3	-2
3	3	0
4	8	-4
5	8	2
6	8	0
7	13	-4
8	13	-2
9	13	0

Table 3.1: Values used for the stochastic GLV simulation. Each row represents a mean for the distributions of growth rate (μ_r) and interaction strength (μ_A).

Note that this includes no cases where the mean interaction strength is positive- this is because the system is unstable in these settings and no community can be formed (see Section 4.1). If a full community is successfully generated with 30 species, I then sample 100 sub-communities of 4 species each, and apply the analysis onto them. Therefore, ideally, given the 9 different sets of parameters, and 10 communities each per parameter set, this produces a total of 90 different time series, each of which is split into 100 sub-communities of size 4. The reason for choosing 100 is mostly due to computational limitations- for some situations where we deal with a single sample community (for example, in real data), 1000 sub-communities are sampled.

3.5 Sourcing Microbial Data

To understand the effect of external structural perturbations on the composition of real microbiota, a first step is naturally to synthesize sources of data on the time-evolution of different bacterial populations in the human microbiota. The reason I restrict to human datasets is that the experimental period of perturbation is more clearly demarcated as the region where physiological effects are seen (as in disease), and diets and exposure to environmental condition changes tend to be more naturally varying and diverse (as opposed to controlled lab environments and diets). Regardless, there are usable mice-based datasets (Turnbaugh et al., 2009; Gibson et al., 2021) which may potentially be viable. Given the aforementioned pipeline used on simulated data above, I now apply it to real data to understand factors driving variation in response and structural sensitivity of these systems. There are an extensive number of experimental studies that have looked into changes in the human microbiota using metagenomics under a variety of conditions, including dietary changes (David et al., 2014b), onset of disease (Valdes et al., 2018; Hou et al., 2022), aging (Nagpal et al., 2018; Ghosh et al., 2022), use of antibiotics (Dethlefsen et al., 2008), and international travel (Boolchandani et al., 2022; Worby et al., 2023). In light of this, identifying the criteria that defines an appropriate dataset for applying these analyses is crucial.

To begin with, one critical factor is the presence of a major external perturbation during the period sampled. This allows testing whether the VCR is able to identify structural changes in real data, and thereby allowing comparison of responses of different subcommunities to this perturbation. Given the multiple range of physiological and external factors that could modulate the dynamics of the gut microbiota, a wide range of possible events could lead to major structural perturbations and associated reorganizations of gut microbiota. While external environmental perturbations are constantly taking place, many of them on a day-to-day basis may not be significant enough to cause appreciable changes to the microbiota that can be picked up, especially given the intrinsically noisy nature of dynamics in the system. While the similarity statistic could potentially be used to identify periods where the system sensitivity changes as hallmarks of perturbation in unannotated datasets, these changes could also come about as a result of dynamical perturbations or measurement errors, and are difficult to validate. Therefore, while these methods developed above are general enough and could certainly be applied onto unannotated data, the results may not necessarily be indicative of responses to external perturbations specifically.

Another necessary condition for an appropriate microbiota dataset is sufficient length- since the EDM framework is built around using data to reconstruct a non-linear attractor, it requires a sufficiently long sampling timeframe to capture the full range of states and dynamics. Else, the inferred ecosystem properties may be inaccurate, especially since the S-map takes into account neighboring points to make predictions. For bacterial communities, due to fast generation time, this time period may be on the order of a few days to weeks. Similarly, another requisite condition is the sampling timestep- given the pace of bacterial community growth and change, daily samples would be ideal. However, weekly samples could also be potentially useful if the time-series is long enough to capture a wide range of states. For EDM in particular, a critical factor is that the sampling interval size is consistent over time. Though methods to overcome this have been developed into a framework called VS-EDM (Johnson and Munch, 2022), these can often be sensitive to noise. Finally, since the aim is to make ecological inferences, the dataset also needs to track the relative abundances of all microbial community members over time, as opposed to some particular focal species.

With these conditions, the pool of possible studies decreases drastically. Several studies that have explicitly measured external perturbations (eg. Faith et al., 2013; D’Souza et al., 2019; Levy et al., 2020; Worby et al., 2023) are too sparsely sampled for EDM-based inference. Others dealing with medical interventions to the microbiota (Dethlefsen et al., 2008; Fukuyama et al., 2017) do daily sampling for a period before and after the perturbation, but this is for a very limited timeframe. Here also, EDM may not be effective. One extensively sampled dataset (Caporaso et al., 2011) looked into microbial community composition in different organs over many months in different subjects- while this satisfies all the conditions needed in terms of data requirements, there are no logged perturbations, making it impossible to validate structural stability. Another long-term study (David et al., 2014a), however, does record major life events occurring during a year of daily sampling, and therefore fulfils all the requirements for being applicable to our question and methodologies. Unfortunately, this is the only such study found. Though we could potentially use the others, they all bear several caveats and disadvantages. Therefore, while some further analysis is done specifically on these other datasets, I focus on this study as the main test of our methods on real data.

David et al. study the microbiota of two subjects (A & B) as part of a comprehensive longitudinal measurement over the course of a year. As part of this, they track microbial

relative abundance data in both subjects' gut microbiota and Subject A's saliva. This data on time-varying abundances comes from 16s rRNA sequencing of fecal and salivary samples to identify specific bacterial operational taxonomic units (OTUs), and was used to track changes in community composition over time. Notably, both subjects had major life events taking place during the period of the experiment- Subject A travelled abroad for a period of 50 days, during which there were significant dietary changes, as well as periods of illness. For Subject B, there was a period of enteric infection late into the experiment. The sequencing of fecal and salivary samples during these periods revealed significant compositional changes occurring as a result of these life events. Along with this, they recorded 339 health-related metadata variables during the course of the experiment- this includes dietary composition and nutrition, fitness level, mood, and presence of illnesses.

The original datasets have records of count data for 5431 different OTUs (operational taxonomic unit), measured over an experimental window of 389 days for all three microbiota (two gut and one salivary). An OTU here refers to a delineation of distinct bacteria, which may be identified at varying levels of taxonomic resolution. For example, the species may be known for a particular OTU, but only the order for another. Regardless, based on the 16s rRNA sequences, they have been separated into two different clusters. This clustering is done algorithmically on the basis of grouping gene sequences with high similarity together. An immediate problem that arises with the dataset is that not every day has recorded data, and there are many missing data points interspersed with measurements. While interpolation schemes permit filling in missing data points, these could often produce unrealistic results when there are long periods of missing data. The first step is thus to remove long strings of missing data from either end. For Subject A's gut microbiota, the last 29 days all have no recorded counts, so these can be removed, leaving 360 days of data. For Subject A's salivary microbiota, the last 29 days as well as the first 26 days have no data, leaving 334 days of records. Finally, for Subject B's gut microbiota, the last 80 days have no data recorded, leaving 319 days of data. However, if these long stretches of missing data appear in between existing data as opposed to at the beginning or end, it would not be accurate to excise these points and merge the existing data on either end, as the temporal structure of the system is lost, and non-contiguous states are put together. This necessitates some subjective choices to filter down the data further. In particular, in the last 100 days (218-318) of Subject B's gut microbiota records, only 8 days have recorded counts. This extensively sparse data makes EDM-based reconstruction difficult. Therefore, I subset Subject B's gut data to the first 218

days of the experiment. Such large regions of sparse data are comparatively absent in the records of both of A's microbiota; I retain all 360 days in A's gut data and 334 in A's saliva, though it is worth noting that these also contain gaps of up to 13 days (specifically, data is absent from days 259 to 271 for Subject A). A variety of interpolation schemes could be used to fill in these gaps. Given the requirement for abundance data to remain positive, spline interpolation is often inappropriate as it creates negative values. Other viable techniques could include simple linear interpolation or potentially simplex projection for forecasting missing values based on similar states in the past and future of the time series. For now, because of its simplicity, I proceed with linear interpolation; however, in the appendix, I show the same analysis with simplex projection used to fill in missing values (see Appendix A.1).

Now that missing days in the data have been filled in, I now proceed to further processing to the dataset to prepare it for the application of EDM. First, I group all the OTUs at the family level. Several recent studies of microbial communities (Goldford et al., 2018; Tian et al., 2020; Ho et al., 2022) have shown that coarse-graining taxonomy at the family level preserves metabolic relationships and structures, producing similar dynamics to lower taxonomic levels without loss of fine structure. I also replicate the following results at the genus and order level (Appendix A.2). Having done this, I now convert all species counts for a particular day to the relative abundance fraction by dividing by the total count for that day. This is because in bacterial count data, the numbers reported are counts among the bacteria that were counted in that particular sample (as opposed to the absolute abundance). Casting this into relative abundance fractions thus normalizes the counts across different days where the amount of sample counting differs. Subsequent to this, unannotated families, which may be combining the counts from several different orders, are excised. This produces time series of the relative abundances of 151 different bacterial families in all 3 microbiota. However, many of these families have extremely sparse time series, with a recorded abundance of 0 on most days. Though these could potentially be interesting, the high number of zeroes makes them unsuitable for the application of EDM. Therefore, I choose to include a family if and only if it has no more than 13 zeroes in a row during the period of measurement. Thus, families that have zero abundance for a period of 2 weeks or more are excised. This leaves us with a much smaller taxon pool. This is also replicated with smaller and longer cutoff lengths to show robustness (Appendix A.3). Then, each family's time series are normalized to have mean 0 and variance 1, as is the prerequisite for EDM, and 1000 communities of 4 taxa each are sampled from the larger taxon pool (cases for larger subcommunity species

richnesses are shown in Appendix A.4). Similar to the previous simulated dataset, the S-map is applied to each community by selecting an optimal θ by leave-one-out cross-validation, then constructing a sequence of (4×4) Jacobian matrices. Subsequent to this, the similarity statistic is calculated for the set of VCRs, as well as for the set of sub-community abundances and mean interaction strengths. Then, before causal inference is applied to these communities to assign profiles to them based on the sign, I bring in the metadata to apply causal discovery.

3.6 Metadata Processing and Causal Discovery

During the course of the experiment, David et al. collect extensive metadata on the lifestyle of subjects *A* and *B*. This comprises 339 different metrics measured for a total of 387 days, ranging from sleep tracking data, presence of common ailments like a cold or bloating, amount of exercise done, stool properties such as frequency and hardness, mood in the morning and night, oral hygiene activities, amount of macronutrients consumed, and quantity of different types of foods consumed on each day. In addition to days where no data at all was recorded, there are also some days when a particular parameter was unrecorded (for example, on a particular day, sleep quality may not have been logged, but others were). Before causal discovery approaches can be used, these missing values need to be filled in. First, similar to the previous case, I remove unrecorded days from the beginning and end. This leaves 360 days for Subject A’s metadata, the same as the abundance data in the gut (and longer than the 334 days for A’s saliva), but only 176 for Subject B’s metadata, shorter than the 218 days seen for abundance. Where some specific data is missing at the beginning or end, interpolation cannot be done directly as there is no start/end point- for these, the nearest recorded data point is filled in as the rest of the values. Since this is only a few days on either end, the impact on the data is minimal. The remaining missing data points are filled in by linear interpolation, similar to what was done for the abundance data. However, for some metadata variables, linear interpolation needs further correction. This is because many variables (such as Boolean variables like presence of illness, or the quantity of food eaten) need to be either constrained in $\{0, 1\}$ or integers. Thus, for these particular variables (recorded as instance, boolean, integers, food, or diet in the metadata), I round down. Since the metadata for Subject A encompasses both the saliva and gut data, there are only 2 metadata datasets. However, Subject B’s metadata contains some further problems- there are 8 variables for which no data is recorded, and 4 for which only 4 days of data is recorded. These variables are all excised from the analysis. This produces time series of all metadata variables for most of the duration of the experiment.

For performing causal discovery, the Inductive Causation (IC) algorithm described earlier (Verma and Pearl, 1990) is used. This entails checking for statistical independence using the G^2 -test of independence, a variant of the χ^2 -test, using the `pcalg` package in R (Markus Kalisch et al., 2012). Then, the process for applying the IC algorithm is modified from a recent study (Tabi et al., 2023), fixing a significance level of 0.06 for the G^2 -test. One shortcoming of the IC algorithm is that it cannot reliably infer causal links if the number of terms supplied to it at once is too high. Given the extremely high number of metadata variables here, this necessitates a piecewise approach of taking 3 or 4 metadata variables and the 3 ecosystem variables at a time, then finding causal links (3 in the case of Subject B's gut, because of the lesser amount of time-series data). This is further complicated by the fact that since each sub-community's composition differs, there may be causal effects of a certain metadata variable on one community's ecological parameters but not another. The way to overcome this is to analyze each community separately, then infer if a particular causal link affects the whole community by looking at the frequency of significant relationships across all communities. This allows us to look for novel linkages with high levels of effect on the system.

From this, I construct and validate causal graphs, then proceed to compute the NDE and NIE of the abundance and mean interaction strength on VCR, correcting the formulas to account for potential additional causal links picked up by the IC algorithm. Note that the G^2 -test shows that not every one of these links is statistically significant- however, I still retain these in the profile to maximize data availability, noting that making this correction later could strengthen results. Similar to before, I prepare frequency counts for each of these causal profiles, and also compute similarity statistics for each profile. This constitutes the complete pipeline for the real data as well.

Finally, while they carry several limitations, I replicate these analyses on similarly filtered datasets from two other studies- one very short dataset of the effect of antibiotics on the gut (Dethlefsen et al., 2008), which contains abundance data of gut microbiota for three subjects undergoing two independent sessions of antibiotic therapy 8 months apart. No metadata apart from the days of antibiotic treatment are recorded here. Another study (Caporaso et al., 2011) contains extensively sampled daily count data for 2 subjects from skin microbiota on 2 different sites, salivary microbiota, and gut microbiota, but does not log metadata or life

events during this period, making results impossible to validate. These analyses are regarded as secondary and only one or two select results from them are seen in the main text as an auxiliary support to the main dataset from David et al.

Chapter 4

Results and findings

There have been never been individuals.
We are all lichens.

Entangled Life, Merlin Sheldrake

Given the various different analyses done here, the results are presented in a modular fashion, sequentially testing each of the methods on the synthetic and then the real data.

4.1 Synthetic Data: Generating Time Series and Jacobians

First, based on the range of parameter values supplied, in what cases can the community successfully generate a complete time series, featuring a pre-perturbation region, a perturbed period, and a post-perturbed recovery phase? Recall that several of the parameters of the simulation were fixed, and the system was iteratively simulated over a grid of values of the mean growth rate and interaction strength. I observe that time-series are successfully generated in all cases except those where the mean interaction strength is positive (indicating average mutualism). In such cases, the simulation cannot run stably and tends to blow up due to positive feedback loops. To see if it is possible to include these cases, I run further simulations with extremely high emigration rates (90%) to see if high departure rates from

the system could potentially stabilize it. However, these simulations also do not succeed, and several other parameters such as the standard deviation of the interaction matrix, the sparsity, and the noise likely need to be simultaneously modulated to produce stable communities with high numbers of mutualisms. Regardless, the constant gradation of values of mean interaction strength from highly competitive to weakly interacting regimes where species are largely independent over a wide range of growth rates provides enough diversity of dynamics to investigate ranges of responses.

For each parameter value set for which it is possible to produce a complete time series, I generate 10 different 30-species communities. For each of these, I sample 100 sub-communities of size 4 species each, normalize them, and then apply the S-map algorithm with θ selected by leave-one-out cross-validation. While the size 4 is chosen arbitrarily, the simulation can easily be repeated for larger community sizes, with a majority of further results reported for a community size of 4 being qualitatively identical for larger communities (see Appendix Section A.4 for validation of this fact on real data). Then, from each of these, the time series of VCRs, abundance, and mean interaction strengths are produced and stored.

4.2 Synthetic Data: Similarity Statistic

To begin with, I note that there is quite a bit of diversity of response during the perturbation period in different sub-communities and across different communities with the same as well as different parameters. Figure 4.1 depicts the VCR over time for 3 different sub-communities in one of the communities with $\mu_r = 8, \mu_A = 0$.




Figure 4.1: VCR over time for 3 sample communities. Each subcommunity is sampled from the same larger community. Highlighted region denotes the period of the simulated perturbation.

Clearly, there is extensive variation in response of sub-communities to perturbations, war-

ranting further investigation on both whether these are seen at the whole-community scale as well as what determines this variation (beyond identity of the species and the effect of the perturbation on it). The figure shows that while the perturbation causes the VCR of some communities to increase, it decreases in some cases, and for some cases, it remains indistinguishable from the other periods. However, remember that these values cannot be directly compared against each other, only relatively within the same time series. These cannot be directly taken as a measure of effective interaction strength as well, because there could be a significant normalization term that is absent from the Jacobian- in accordance with this, the Jacobian of competitive communities where $\mu_A = -4$ has elements of the order of roughly 0.05. Naturally, the next step is to apply the similarity statistic to these communities. By fixing a baseline period of 50 days, a window size of 14 days, and a cutoff z-score of 2, I apply the similarity statistic to the VCRs of each of the communities.




Figure 4.2: Similarity Statistic for Different Parameter Regimes. Highlighted region denotes the period of the simulated perturbation. Similarity statistic values are calculated by averaging over 10 replicate communities generated for the same parameter values.

Figure 4.2 shows the similarity statistic of the VCR varying over time for sampled communities in 3 different parameter regimes. The scores here are averaged over the 10 communities generated per parameter value. For ease of comparison, here, I fix the growth rate at a middle value ($\mu_r = 8$) while modulating the mean interaction strength from -4, -2 and 0 (the other cases for the other mean growth rate values of 3 and 13 are qualitatively identical and can be seen in Appendix B.3). Clearly, while specific values and the extent of depletion varies across cases, all the time-series show appreciable dips in the value of the similarity statistic during the defined period of perturbation, days 101-150. Specifically, while the community similarity statistic appears to fully recover back to pre-perturbation values in the case where $\mu_A = 0$, there are situations where there appears to be some sort of hysteresis in

the other two competitive communities, where recovery to the original attractor does not occur even after resetting the parameters to their original value. Looking further into this by assessing the similarity statistic individually for each community, we note that this lack of complete recovery is because of certain replicates of the same set of parameters not recovering, while others do, and is not the case that all communities show this partial recovery. This is expected in such a competitive community, given that some perturbations could potentially lead to competitive exclusion of a set of species that was originally present, with constraints to re-establishment achieved via priority effects induced by constant immigration of other species (Grainger et al., 2019). This produces robust states with attractor shapes that are fundamentally different from the original state, indicating persistently different similarity statistics.

In general, however, note that the similarity actively declines in all cases when the perturbation appears- this shows that the ability of the similarity statistic to pick up a perturbation onset is extremely robust and works across a wide range of parameter values and regimes. Recall that while a perturbation was artificially simulated in generating the time series, the S-map reconstruction and similarity statistic have not used this fact anywhere; instead, the statistic has identified it based on the fraction of communities responding with a change in their VCR. Notably, after the parameters are set back to their original values, the similarity statistic recovers to a high value with respect to the baseline in a large number of cases. While this recovery may not always occur depending on the interaction structure of the community or can sometimes take a little time, there is some amount of recovery of the similarity statistic in a majority of cases. Especially in weak interaction regimes, the system attractor appears to robustly return back to the original topology after the parameters are reset to their original values. While this may be trivially true if the system relaxes to a single equilibrium point, the similarity statistic makes no such assumptions and is able to detect this even in cases where the attractor could be more complex like chaos or limit cycles.

Another feature to note is that even in periods of before the perturbation, the similarity statistic does not take on a value of 1. This could be because of two reasons. One is that some sub-communities may be intrinsically highly unstable solely because of the noisy nature of the system and therefore have a high likelihood of falling outside the boundaries defined by the baseline distribution and cutoff. Another reason that the initial state for all species is taken to be an abundance of 0.5, which is likely not on the attractor- thus, there is an initial

transient phase during which the system moves towards the dynamical attractor, which skews the distribution. This can be tested potentially by generating longer time series and excising transients. One variation on the similarity statistic that can be done is to also look at the frequency of cases where the z-scores are higher/lower than the cutoff. While the similarity statistic on its own tells us directly about the presence of perturbations and departure from a particular state, the fraction of higher/lower values can tell us about the direction that the VCR during the perturbed period is taking. Given that the structural perturbations done to the system are random in nature, it can be illustrative to look at their effect on structural sensitivity in different settings. Indeed, in Figure 4.3, I show the frequency of communities that have a z-score more and less than 2 over time for a single representative community with $\mu_r = 8, \mu_A = 0$.




Figure 4.3: VCR values above and below cutoff. Similar to the calculation of the similarity statistic as the proportion of communities within a cutoff, we may compute the frequency above and below the cutoff to track the fate of the average community’s parameter values compared to the baseline over time.

This shows that the fraction of communities that increase in VCR compared to the baseline is very high in the perturbed period, indicating that a majority of communities destabilize strongly from the baseline and enter a state where they become more sensitive to disturbances. The number of low-VCR communities also seems to increase a little bit in the period of perturbation as well, indicating that certain communities may strongly stabilize during this period (though this ‘stabilization’ could simply emerge from large abundance declines). It is, however, worth noting that this specific result is to illustrate this sort of inference, and the specific outcomes and interpretation depends on the structure of the community and

the nature of the perturbation. Another pertinent question about the similarity statistic is whether it could potentially track relative sizes of perturbation- to test this, I fix a mean growth rate of 8 and mean interaction strength of 0, and sample 1000 sub-communities, varying the standard deviation of the perturbation distribution to 3 values- 0.1, 0.5 and 1. Then, to quantify the effect, I calculate the mean of the similarity statistic during the period of perturbation, days 101-150. For standard deviation values of 0.1, 0.5, and 1, I get average statistic values of 0.198, 0.134, and 0.059 respectively. This seems to suggest the similarity statistic is able to successfully also able to capture differences in the strengths of structural perturbations, with larger-sized perturbations pushing the similarity statistic further away. However, this is more of a comparative assessment, and the statistic cannot be directly used to reconstruct a perturbation strength from its value in a particular region without further analysis. Finally, another interesting question about these results is how sensitive it is to different extents of perturbations in terms of how many species get affected. While many external perturbations can change the entire physiological environment, affecting the growth rate of all constituent species, other structural perturbations may have more localized effects, only affecting the interaction strength or growth rate of a few select species. Therefore, to test the effectiveness of the similarity statistic to these types of perturbations, I also simulate cases where only a single parameter becomes perturbed by a single random value sampled from a normal distribution. The results of one such case on a sample community with $\mu_r = 8, \mu_A = 0$, sampling 1000 sub-communities, and with a perturbation distribution standard deviation of 1 are shown in Figure 4.4.

Once again, we see a clear depletion in the similarity, though the amount is far less compared to the all-parameter case, as expected. Note that this result, however, depends on the specific nature and strength of the perturbation and the identity of the species chosen- while a structural change occurs and changes the state of the community, in some cases when the perturbation is small or a particular species with little interaction is chosen, it may not cause a significant enough change in the similarity statistic to be detected- this is just to demonstrate potential capabilities. In any case, the similarity statistic presents a highly robust, sensitive and generalizable method to detect structural perturbations in time series data. Notably, this also shows that the VCR, which was previously developed and used as a predictor of sensitivity to perturbations, can also be used as an indicator of the presence of perturbations, since different states are fundamentally associated with different structural sensitivities, and the probability of departure from those states is governed by their local structural sensitivity. Additionally, the similarity statistic is a general metric that




Figure 4.4: Similarity Statistic with a Single Structural Perturbation. Here, I compute the similarity statistic of a perturbed Lotka-Volterra community, but perturb only a single species' growth rate and a single pairwise interaction by a certain amount. The similarity statistic successfully picks up the perturbation.

can be applied to any family of time series data, including the mean interaction strengths and abundances. Figure 4.5 depicts the similarity statistic of abundance and mean interaction strength in a sample community with $\mu_r = 8, \mu_A = 0$, indicating similar results. However, in the specific context of assessing sensitivity to external perturbations, the VCR is the best indicator of community-level changes, as a metric of structural stability. Abundances can be drastically changed by dynamical perturbations, while mean interaction strengths could pick up perturbations to the interaction matrix but may be less successful for changes in growth rate.

Before proceeding, it is important to highlight some potential caveats with the use of the similarity statistic to identify the deviation from a baseline state over time. One is that the definition of the baseline is a critical feature here that must be chosen carefully, since this is used as an indicator of 'health', and determines the associated distribution. If a good baseline is not available, reliable predictions of periods of change from the baseline will not be possible. However, if pre-perturbation baselines are not available, the method may still be applied in reverse, by using data from during the perturbation as a baseline, and assessing the depletion from the baseline state over time as an indicator of recovery of the whole community. Another potential pitfall is the presence of dominating levels of noise, which




Figure 4.5: Similarity Statistic for Abundance and Mean Interaction Strength. Similar to VCR, the similarity statistic is applied to these other parameters as well, and is shown to successfully highlight the perturbation period in these as well.

could prevent the perturbed community states and VCRs from being delineated from each other reliably, producing fluctuating similarity statistic values with no clear demarcation.

In general, the similarity statistic presents a potent way to reconstruct perturbations in a large, noisy interacting community with minimal structural assumptions by looking at the change in the state of its sensitivity over time in smaller sub-communities. This metric is robust over a wide range of parameter values and types of random perturbation, and highly sensitive even to small perturbations.

4.3 Synthetic Data: Causal Profiling

Next, I detrend the time series of VCR, mean interaction strength and abundance for each sub-community in each larger community. Then, binarizing about their median values, I compute the Natural Direct Effect for the effect of abundance on the VCR and the Natural Indirect Effect of interactions on VCR. Then, based on the signs of these causal effects, I assign each sub-community into profiles PP, PM, MP or MM (the first letter denoting the effect of abundance on VCR). Then, since each community contains 100 sub-communities, I calculate the fraction of different profiles in each community. Then, since I generated 10 communities for each set of parameters, I average the fraction of these profiles across those larger communities to get a final proportion associated with that set of parameters. I repeat this for all the parameter values and summarize the results in Table 4.1.

Mean Growth Rate	Mean Interaction Strength	MM	MP	PM	PP
3	-4	28.2	22.8	15.4	33.6
3	-2	34.8	28.2	12.7	24.3
3	0	45.7	35.8	8.2	10.3
8	-4	25.1	22.9	13.3	38.7
8	-2	27.1	26.9	14.1	31.9
8	0	34.4	30.4	15.8	19.4
13	-4	22.2	20.4	18	39.4
13	-2	26.8	26.9	14.8	31.5
13	0	37.2	29.4	14.6	18.8

Table 4.1: Causal Profile Frequency. These frequencies of each causal profile are computed by finding the fraction of the 100 subcommunities in each replicate with a particular profile, then averaging this over the 10 replicates with a particular parameter value. As growth rate is fixed and interaction strength is varied, there appears to be a transition from PP-dominated systems to MM/MP-domination.

An immediate observation is that depending on the nature of the community parameters chosen, the proportion of different profiles is radically different. First, in all cases, communities of type PM appear to be unusually rare- this could be due to the specific nature of the fact that we work with largely competitive communities, and an artifact of the Lotka-Volterra model structure. This is likely a result of the specific parameter values chosen- indeed, this trend does not persist if some other parameters such as sparsity and standard deviation of interactions/growth rate are modified- this could be emerging due to the differing distributions of the diagonal and off-diagonal elements. I do not analyze this further at this point, and instead concentrate on the other profiles- despite the smaller number, these also show trends somewhat consistent with what I will highlight now. To understand the variation in profile frequency in different parameter cases, I concentrate only on the case when $\mu_r = 8$ and vary interaction strength- the other cases of growth rate follow similar

trends. We note that when competition is high with mean -4, there is a surge in the number of communities of type PP, and comparatively fewer of type MM and MP. However, this dominance tends to decrease as competition strength is weakened to -2, with MM and MP type communities becoming more common. Finally, at the weak interaction case, the pattern inverts- MM and MP become the most common types of communities, while PM and PP tend to be rarer. This is a sort of continuous transition as the interaction strength is tuned upwards. However, every community contains an appreciable proportion of profiles of each type, highlighting the fact that the effects of interaction and abundance changes on the VCR are different across different sub-communities, implying the lack of a single universal and generalizable directional ecological effect on these communities' dynamics. This continuous transition, however, suggests that different communities have fundamentally different broad ecological effects, which may mediate their response to perturbations as well. This merits further investigation.

In general, profiles of type PM and MM are interesting from an ecological standpoint. While both positive and negative arguments could be made for the effect of total species abundance on community sensitivity, especially depending on the identities of the species whose abundances are increasing, classical ecological theory, especially based on Lotka-Volterra pairwise models, generally posits that an increase in competition tends to stabilize communities, and a decline in competition towards mutualism tends to drastically destabilize the community, an effect we saw when simulating the time series as well (May, 1972; Allesina and Tang, 2012; Coyte et al., 2015). However, when considering aggregate effects from the Jacobian, it appears that the effect of increasing mutualisms and weakening competition is less direct, and can vary in its effect on stability, being stabilizing in the case of some communities (PM/MM) and destabilizing in some (PP/MP).

This could be occurring due to a variety of reasons, chief among them being that the classical formulation of destabilization is based specifically on the case of dynamical and **not** structural stability. Indeed, for a dynamically unstable system such as the mutualistic Lotka-Volterra communities I attempted to simulate, a dynamical perturbation such as the addition of noise could lead to severe instability. Previous work (Saavedra et al., 2017; Cenci and Saavedra, 2018a; Song et al., 2020) has shown a fundamental link between the structural stability of a community and the size of the feasibility domain, the vectors of intrinsic growth rates leading to positive equilibria (Logofet, 1993; Rohr et al., 2014). A

larger feasibility domain is linked to higher ecological resilience in some sense, as the range of structural perturbations that do not lead to loss of the feasible community state is larger; however, there is also the factor that in a larger feasibility domain, typically, retention of the feasible community state against a wide range of structural perturbation will prevent large-scale deformations caused by subsequent exclusion of species if moving outside the domain. In the context of the feasibility domain, a decline in competition and increase in mutualism has been shown to increase the size of the domain (Rohr et al., 2014; Saavedra et al., 2016); in particular, communities with predatory interactions have been shown to have the most stabilizing effect against structural changes by increasing the size of the domain (Song and Saavedra, 2018), indicating that a perturbation setting one term of a competitive Jacobian interaction to a positive value could have an appreciable effect on structural stability. Therefore, in this setting, since we track structural stability, a variable role of interaction strength on the structural sensitivity is in line with existing theory- the variation in sign could be emerging depending on the specific interaction types of the constituent species modulating the size of the feasibility domain, and the specific perturbation experienced by them to their parameters. Also, recall that in a Jacobian setting, the definition of a ‘mutualism’ is based on an aggregate effect and not pairwise interactions, which could also skew existing theoretical predictions based on pairwise models when complex combinations of positive and negative effects interplay in determining final interaction coefficients. Parallely, several other approaches have shown positive effects of positive interactions on several ecosystem parameters such as preserving species abundances, diversity, and conferring resilience to perturbations (Hale et al., 2020; Qian and Akçay, 2020; Stone, 2020), especially in the context of re-approaching original abundances after a depletion. Microbial communities have also been shown to contain an appreciable amount of mutualistic interactions (Machado et al., 2021).

The important takeaway here (apart from the above observation) is twofold. One is that the causal profile varies across sub-communities, and all different profiles are represented with significant frequencies in sub-communities. The second is that as communities’ interaction strengths move from a competitive regime to a regime where most interactions are weak and equally likely to be mutualistic and competitive, the composition of different causal profiles is modified. In competitive regimes, PP and PM type communities dominate, while the opposite is seen in weak interaction cases. The exact reason why this is the case is unclear, and needs further investigation. These frequencies, however, seem to be characteristic of a

particular regime- this appears to be reasonably indicative.

4.3.1 Sensitivity of Causal Profiles

Having partitioned each sub-community into different causal profiles, I now apply the similarity statistic onto collections of only the time series from each of these profiles separately in order to assess differences in sensitivity for each. Fixing a mean growth rate of $\mu_r = 8$ and varying mean interaction strengths as -4, -2 and 0, I plot the average sensitivity over time for each profile in Figure 4.6- these values are generated by averaging over the 10 communities with the same set of parameters. The cases for other values of the growth rate, producing qualitatively similar results, are shown in Appendix B.3.




Figure 4.6: Similarity statistic across profiles. This shows the similarity statistic over time for each causal profile in 3 cases with different interaction strengths. Clearly, competitive communities show a disparity in response, with PP/PM responding less. This is not so in weak interactions, where all respond largely equally.

Here, we note immediate differences across different cases and profiles. In the competitive communities in Figure 4.6(a) ($\mu_A = -4$ or -2), the profiles corresponding to MM and MP show much more depleted similarity statistic values compared to the other profiles. This indicates that they respond much more strongly to the perturbation, entering a period of higher dissimilarity to the original baseline. Indeed, by averaging the value of the similarity statistic during the perturbed region of Days 101-150 for the $\mu_A = -4$ case, we may see that the average MM and MP similarity statistic value is 0.328 and 0.284, significantly lower than the 0.639 and 0.505 seen for PP and PM. This has several implications. First, it suggests that certain profiles are less stable than others, explaining the variation seen in response. This suggests that the causal profile appears to be a potential determinant of variation in similarity

statistic during the perturbation. Furthermore, looking back at the causal profiles, it appears that the more volatile MM and MP type profiles are less frequently occurring among the generated sub-communities. Also, since the major difference between responses seem to be separated largely by the first letter (MM/MP versus PM/PP), this suggests that at least in this competitive setting, the effect of abundance on the VCR is the primary determinant of dynamics in these systems, as opposed to the effect of interactions. However, since the dynamics of the different profiles even within these two partitions are not exactly the same, this does suggest that interactions do play an important role in structuring the finer-level responses of the system.

Looking at the other case of different interaction strength, where $\mu_A = 0$, we see that there is no particularly more or less sensitive profile, and all of them seem to respond fairly similarly. Since these plots are averaged over 10 replicates, this presents two possibilities- either all profiles respond similarly in all replicates, or different profiles respond differently in distinct replicates, but on average, do not show a stereotypical response. By averaging similarity statistic values over these 10 communities in the perturbed region, we may show different profile sensitivities of 0.162, 0.069, 0.094, and 0.09 for PP, PM, MP and MM profiles respectively. While PP appears to be a little more stable, this disparity is far less compared to what is seen in other cases. To understand this homogeneity further, I compute the average similarity statistic value in the perturbed region for each profile separately for each of the 10 constituent communities and tabulate their values below in Table 4.2.

From this, we see that the second case appears to be true- in different replicates of the system, different communities appear to have different response of similarity statistic to perturbation. Recall that these replicates have different interaction matrices and growth rate vectors, as well as different random perturbations acting on them, and hence show different responses. Given that the interactions are weak, it is possible that this emerges because there is no strong network structure governing a stereotypical response to a random perturbation- therefore, the response of the system may depend specifically on the identity of the perturbed species and the specific direction of the perturbation experienced by those species. Recall that despite random perturbations with the same distribution, competitive communities showed a consistent identity of most sensitive species, which also held true within individual replicates, possibly because the perturbation was small enough to not drastically change the types of interactions in the community. It is worth noting that while PP and PM type

Community	PP	PM	MP	MM
1	0.0806	0.0122	0.0288	0.004
2	0.085	0.038	0.1517	0.206
3	0.185	0.189	0.079	0.028
4	0.179	0.053	0.113	0.081
5	0.117	0.037	0.137	0.096
6	0.414	0.152	0.133	0.201
7	0.084	0.046	0.044	0.118
8	0.281	0	0.08	0.09
9	0.207	0.088	0.153	0.007
10	0.091	0.014	0.019	0.063

Table 4.2: Similarity statistics in the perturbed region for each replicate in the weak interaction regime. Averaging the score over the whole region, we find that in different replicates, different profiles respond with more/less sensitivity. Differences in values within a column emerge because of the randomly sampled strength of the perturbation- some stronger, some weaker.

communities are rarer in weak interaction regimes, they are not always the least stable communities and it depends on the nature of the perturbation- thus, this relation may not hold true when interactions are weak. However, it is possible that a consistent trend would emerge if interactions are made significantly positive- a complete test of this transition would entail testing whether mutualistic communities have more sensitive PP/PM type profiles as well as lower frequencies of occurrence. However, since it is not possible to generate these without compromising other parameters severely, it is necessary to stop simulations at weak interaction regimes. However, one potential way to stabilize mutualisms is to add in factors such as negative growth rates (currently sampled from a log-normal distribution) or Holling Functional Response curves (C. S. Holling, 1959), which saturate based on density of interacting partner as opposed to growing linearly indefinitely.

In summary, in competitive communities, certain sub-communities seem to be more sensitive to perturbations, with much more deviation in similarity statistic of the baseline state during the perturbed period. Thus, the causal profile appears to determine a large part of the variation in response inherent in different sub-communities provided the community is dominated by competitive interactions on average. However, in weak interaction regimes, the identity of what the most sensitive sub-communities are appears to change in different replicates. When interactions are strongly competitive, the rarer profiles (MP/MM) are also the most sensitive. However, when interactions are weak, despite the PP profile being rarer, the sensitivity of profiles appears to depend on the specific nature of the perturbation and species identity.

4.4 In Vivo Microbiota Data

I replicate similar results on the real data, but include additional sections on the composition of the data and the causal discovery steps.

4.4.1 Family Composition and Structure

By filtering out the rare species with more than 2 weeks of no abundance, I see that Subject A’s gut microbiota contains 30 common families, Subject B’s gut microbiota has 18, and Subject A’s saliva contains 38. As expected from the higher baseline species richness, salivary microbiota tend to contain more common families as well (Maki et al., 2021). The included families in the analysis are listed in Section B.1 in the appendix. One interesting thing to note here is that there is quite a bit of difference between the families included in the salivary and gut microbiota of Subject A, despite being sampled from the same subject, but all families seen in B’s gut microbiota (save for one, Verrumicrobiaceae), are found in A’s gut microbiota as well, suggesting a possible “core” gut microbiota composition, without necessarily having the same dynamics (Sharon et al., 2022). Given this, I now normalize this relative abundance data, sample 1000 communities of 4 species each, and apply the S-map with an optimal θ to each of the 3 datasets. The results are replicated on larger community sizes in Appendix A.4.

Similar to the case of the stochastic GLV, I randomly sample 3 sets of 1000 sub-communities from Subject A’s gut, Subject B’s gut, and Subject A’s saliva respectively, and look for variation in response and periods of changed dynamics by looking at the time series of their VCRs over time. These sample time series are shown in Figure 4.7.




Figure 4.7: Sample VCRs in real microbiota. This shows the extensive variation in VCRs of different sub-communities over time. The colored regions denote the annotated period of perturbation. There is a clear variety of VCR responses occurring.

Clearly, there is quite a bit of variation- however, several communities seem to show characteristically different VCR values during the periods of recorded perturbation, highlighted in the plots. Similar plots are shown for the abundance and mean interaction strengths of the same set of sampled communities in Appendix B.2. A natural next question is to see if the perturbation is significant at the whole-community level by applying the similarity statistic onto the community VCRs. I use a baseline period of 50 days, a window size of 14 days, and a cutoff z-score of 2 for the gut microbiota. However, since the pre-perturbation period for the salivary microbiota is shorter due to the lack of sampling at the beginning of the experiment, I shorten the baseline period size to 40 days.

4.4.2 Similarity Statistic on Real Data

Figure 4.8 depicts the similarity statistic of the VCRs of each of the microbiota over time.




Figure 4.8: Similarity statistics for each of the 3 microbiota. The colored regions are the annotated perturbations. A baseline window of 40/50 days, a window size of 14, and a cutoff score of 2 are used. The periods of perturbation are associated with dips in the value of the similarity statistic.

First of all, from this, we note that periods of recorded perturbations are also associated with major depletion in the similarity statistic of the VCR- this proves that in real data as well, the similarity statistic is able to pick up periods of external perturbation! This also reveals that the real data is also highly noisy, with fluctuating values of similarity statistic (though few as low as the major perturbation recorded). These may be indicative of smaller, briefer disturbances arising from a range of possible sources. Another important observation is that in all cases, though the system starts at a fairly high amount of similarity to the baseline (especially in the case of the gut microbiota- the saliva baseline being shorter necessitates less inclusion of baseline periods), the recovery process takes a slightly longer time to occur. This could be indicative of the milieu of internal governing physiological parameters slowly relaxing towards the original pre-perturbation state long after the main perturbation has ended. Alternatively, it could reflect the fact that since the microbiota is in a different state, it could be more sensitive to the natural pool of environmental perturbations in this state, preventing immediate recovery. If this is indeed the case, then it could be that the VCR tends to take on higher values during the period of illness. Figure 4.9 shows the frequency of z-scores greater than 2 for each sub-community, similar to Figure 4.3. This indicates that the nature of the perturbation's effect on the VCR is different across different microbiota. In Subject A's gut, while the amount of high-VCR communities increases slightly, but the

amount of low-VCR communities drastically increases. This seems to suggest that a majority of communities display a large *decrease* in VCR during the period of perturbation, indicating some kind of stabilization of sorts into a novel dysbiotic state. However, looking at the case of Subject A's saliva and Subject B's gut, the opposite is true- it appears that the number of communities with a higher-than-cutoff VCR increases during the period of the disturbance, suggesting a deformation into an unstable, highly sensitive state. Comparing A's gut and B's gut, this shows that the nature of disturbances can often be fundamentally different on the VCR on the same organ system microbiota (especially given the distinct agents causing them- one an infection, the other being travel abroad). Comparing A's gut and saliva is more interesting- despite the perturbation being the same, its effect on the VCR of the two microbiota appears to be opposite. This warrants future studies potentially linking gut and salivary dynamics.




Figure 4.9: Above and below cutoff frequencies for each microbiota. The first row of plots depict above-cutoff frequencies, and below-cutoff is the second row. With a cutoff of 2, this shows that the perturbation appears to have different effects- it increases the VCR on average in A's saliva and B's gut, but decreases it in A's gut.

Another fact that is worthwhile to note in Figure 4.8 is the response of different microbiota's similarity statistics. While Subject A and B's gut microbiota display significant dips during the period of perturbation from a highly similar initial state, despite vastly different

natures of the external perturbation that caused this, this is less so for the salivary microbiota, which begins from a lower similarity statistic, then increases only much more gradually back to a higher value. It also appears to possess a fairly significant unannotated period of perturbation some time around Day 260. This is further explored later along with possible explanations in Section 4.4.4 with the similarity statistics split by causal profile, but the main message for now is that the salivary microbiota appears to respond strangely to the recorded perturbation as compared to the gut microbiota.

In spite of the high degree of noise and nonlinearity associated with the response, the similarity statistic applied to the VCR successfully does pick up periods of perturbation, also highlighting potential unrecorded ones, and showing that recovery can often take longer than the period of the perturbation itself. There is also the possibility of hysteretic responses that we saw in the earlier simulated cases, where the system never truly recovers and instead persists in a dysbiotic state for a longer period even after the perturbation disappears- this behaviour, though possible in microbiota (Khazaei et al., 2020; Mangal et al., 2023), appears to not be observed in this particular set of data. This could potentially be occurring in B’s gut, but the dataset after the illness is not long enough to test duration of complete recovery- however, it is noteworthy that the extent of depletion and the time needed for recovery is quite long, a result also stated in the original study.

4.4.3 Causal Discovery and Profile Assignment

With the similarity statistic validated as a potent and robust means of detecting changes in community composition, I now proceed to describe methods of causal discovery applied to the data, with the aim of linking variables in the metadata associated with physiological choices and food intake to changes in community-level properties such as total abundance, mean interaction strength and VCR. As described earlier, this is achieved by sequentially taking groups of 3/4 metadata variables, detrending and binarizing, and then applying the IC algorithm to them with a significance value α of 0.06. I define a metadata variable as having a significant cause if at least 10% of the subcommunities (that is, 100 of them) have their ecosystem parameters affected by fluctuations in the metadata variable. Due to Subject B’s metadata being only 176 days, the ecosystem variables are correspondingly cut down from 218 to 176 days as well. However, this does include the period of the perturbation (days 150-160) as well as the period immediately after, which are of potential interest to us.

Having done this, I observe an unexpected result. None of the metadata variables appear to have a significant effect on any of the community-level metrics, with the maximum numbers being seen being body weight affecting total community abundance in 24 out of 1000 cases in Subject B's gut microbiota, with an overwhelming majority of metadata variables (above 90%) having no perceptible effect whatsoever on any of the 1000 groups of the 3 ecosystem variables, and the majority of the remainder affecting the ecosystem in only 1 or 2 sub-communities (which may be emerging solely by statistical chance). This appears to suggest that no metadata variable can play a significant role in governing the observed dynamics of the community, a result that is in contrast to several studies (eg. David et al., 2014b; Conlon and Bird, 2014; Strasser et al., 2021) that show significant roles of physiological factors on bacterial abundances at the very least. This lack of structure, however, does not indicate that these factors do not affect the microbiota at all, and could emerge from several other sources-

- One highly likely possibility is that the metadata variable's value (for example, the food eaten) indeed causes fluctuations in the community level variables. However, these may not be significant enough to cause a recorded change in the value of the response variable about the median since the response is also binarized. It is notable that most of the previous studies of food on microbiota have applied the treatment for a sustained period as well- the singular nature of these changes (choices made on a particular day) could potentially fail to cause enough directed fluctuation to be recorded as a change in the binary ecosystem properties. This is the most likely reason.
- Another possibility is that the metadata has significant causal roles only during particular times. One shortcoming of this probabilistic causal approach is that it requires a whole time series and assumes that causal profiles do not vary with time. It could be possible to show these relationships in contiguous periods of no recorded perturbation only. In accordance with this idea, David et al. (2014a) show based on correlative approaches on the abundances of bacteria in the original dataset that very few OTUs (210 out of 5431 in A's gut, and 16 of 5431 in A's saliva, representing 8 and 6 different families respectively) are significantly affected by food intake, and this effect is only seen during periods of comparative stability.
- It is also possible that this particular set of variables does not have explanatory power for these ecosystem properties, and there are other unrecorded physiological variables (or sums/rescalings of the measured variables) that do have causal roles.

- Another possibility is that these systems are dominated by endogenous dynamics, and therefore no causal role can be ascribed. However, prior studies linking diet to microbiota seem to suggest some role.
- Finally, another possibility is that there are complex linkages between the metadata variables that necessitates their being analyzed together in order to infer the causal graph structure (for example, if one acts as a collider of sorts). This is testable- by selecting all possible permutations of 4 metadata variables from the larger pool, then running the IC algorithm on these successively for each of the 1000 sub-communities, then testing for significant causal effects. However, this too produces no significant results.

One way to overcome this issue and identify causal variables would be to develop and apply causal inference methodologies based on continuous fluctuation of variables while accounting for mediation and other complex structure (without binarizing). However, given the difficulty of doing this in a robust and general setting, I leave this as an exciting and promising future area of research. It is notable that the IC algorithm does successfully pick up likely causes between the ecosystem variables themselves in a large number of causes, indicating that a majority of sub-communities have significant effects of the abundance and interactions on VCR, as confirmed by the G^2 -test, as well as a few highly trivial links between some metadata variables (such as, for example, the total amount of cheese eaten and the amount of a specific type of cheese eaten). Regardless, this produces, similar to the GLV setting, a causal graph of abundance, interactions, and VCR, with the structure specified as before- the abundance affects the VCR and interactions, and the interactions affect the VCR. It is unfortunate that this causal graph cannot be validated by adding another metadata-associated node that would add d-separations to the system- however, based on our knowledge of ecosystem dynamics, it is reasonable to make these assumptions as a causal graph. I compute the natural direct and indirect effects in each case for the communities, and calculate proportions of each profile.

The results for each microbiota are summarized below in Table 4.3. From this, we note that in all 3 microbiota, profiles of type MM appear to be the most common, followed by MP, PM and then PP. Despite the wide range of dynamics, ecological structure, and perturbation types characteristic of each of the three microbiota, with different similarity statistic patterns, this pattern appears to be conserved. Looking back at the previously generated Lotka-Volterra systems, we see that these values and their gradation is very similar to the characteristics

Microbiota	Profile MM	Profile MP	Profile PM	Profile PP
A's gut	442	263	184	111
A's saliva	362	264	209	165
B's gut	355	253	213	179

Table 4.3: Frequencies of each profile in different microbiotas, out of 1000 sub-communities. Letters (M/P) are assigned to each profile on the basis of the sign of the natural causal effect of the abundance/interaction on the VCR- thus, profile MP has a negative sign of abundance on VCR, and a positive sign of interaction on VCR.

of a community with weak interactions. Indeed, many recent studies (Gibson et al., 2021; Camacho-Mateu et al., 2024) have also shown similar results, that microbial communities tend to be sparse with weak interactions, but taking on both positive as well as negative values (Machado et al., 2021).

4.4.4 Similarity Statistic Across Profiles

In accordance with the results shown for the Lotka-Volterra case of weak interactions, in the absence of any prior data on the structural effect of perturbations and thus treating them as random, I would expect mixed results of sensitivity of different profiles, especially given the radically different natures of communities' organizations and the perturbations acting on them. Recall that we saw earlier that in the regime of weak interaction, depending on the specific nature of the perturbation, a variety of outcomes are possible.

Given the profile assigned to each sub-community, I apply the similarity statistic to each set of VCRs from the profiles for each microbiota, using a baseline of 50 (or 40) days, window size of 14 days, and cutoff score of 5 (to accentuate differences). The results can be seen in figure 4.10.

I go stepwise through these results, starting with A's gut microbiota. Here, we see that while all 4 profiles dip in their similarity statistic during the period of perturbation, the 4 profiles only separate significantly from each other in the period after Day 100, during which profiles of type PP and PM appear to show much greater sensitivity to the perturbation




Figure 4.10: Similarity statistic in the 3 microbiota split by profiles. This uses a baseline size of 40/50, a window size of 14, and a cutoff score of 5. This reveals the consistent response across all 3 microbiota, despite the differing species compositions and nature and effect of perturbations, that communities in the profiles PP and PM appear to respond more strongly.

compared to MP and MM. Proceeding onto the salivary microbiota's response, we see that the depletion during the recorded disturbance period is comparatively much lower compared to the other two microbiota, showing no major dip and remaining fairly high in similarity throughout. The different profiles respond fairly similarly and homogeneously during this period, with only a minor dip from the baseline- while there appears to be a slight bias towards profiles PM and PP, this is not consistent across other reasonable parameter value choices (see Appendix A). However, outside the region of recorded perturbation, salivary responses tend to be more varied, with profile responses starting to diverge after day 250. The main region of 'interesting' dynamics seen in the salivary microbiota is instead around days 250-300, where an unannotated perturbation appears to have been picked up by the

similarity statistic. Here, the similarity drops drastically in most profiles, though the response of different profiles here is unequal. It is visible that the PP and PM type communities tend to respond much strongly here as well. I revisit this strange perturbation shortly. In B's gut microbiota too, we see a response similar to Subject A's gut, where profiles PP and PM appear to respond more strongly to the perturbation, which is quite brief compared to the dip in A's gut (as expected). The similarity statistic remains quite high in the interim pre-perturbation period, however, as expected.

While it might be possible to highlight profiles PM and PP here as being more sensitive here, this claim needs further investigation and validation. Specifically, given the idea from prior studies that microbial interactions are largely weak and sparse, the low sample size of 2 microbiota with an appreciable response to a recorded perturbation (3 if including the unannotated perturbation in saliva), and given the Lotka-Volterra simulated communities showing that weak interaction regimes can in principle have any community responding more strongly depending on the nature of the perturbation, this is not enough information to conclude this as a consistent pattern in microbiota yet. It is possible that by chance, these communities have the same response given a likelihood of many possible responses. To investigate this further, I replicate these same analyses on a variety of alternative parameter choices. As highlighted earlier, these include grouping at the genus/order level, picking different rare species inclusion cutoffs, different interpolation schemes, and different subcommunity sizes. These plots, shown in Appendix A, show that this result is extremely noisy, with some cases showing differing results across the 3 microbiota (even across different replicates of the same system parameter set, though this is not shown in the Appendix). While it is possible that these inconsistencies may be emerging due to other parameter choices failing to capture certain subtle aspects of the system by coarse-graining it or adding noise, it is difficult to conclude robustly from this alone that certain communities are more/less stable in real ecosystems, especially given the synergy with weak interaction regimes shown by the causal profile frequencies. Further investigation with refined causal profiles, more realistically parameterized simulations, and trials are needed for concluding this, though crucially, this does provide a potential set of candidate communities for further analysis. It is worth noting that these profiles also tend to be rarer in terms of proportion among the sub-communities across all 4 microbiota, indicating a possible ecological phenomenon of the natural structure of the microbiota organizing themselves to have a larger number of more stable communities if true, similar to how the rarer MM/MP type communities were less frequently occurring in

competitive systems.

Before proceeding, I finally confront the unannotated perturbation and the apparent lack of strong response to the known disturbance from the saliva. First of all, the salivary microbiota does not show much of a response across all profiles during the period of recorded perturbation, with the value of the similarity statistic remaining high throughout. This could be because of the fact that when one samples from the salivary microbiota by sequencing salivary samples, the saliva contains species taken from many different niches in the oral cavity (Mark Welch et al., 2020; Wilbert et al., 2020). This leads to mixing together of different communities in different states in the sample, possibly averaging out the dynamics. Another alternative possibility is simply that during the perturbation of moving abroad, the salivary microbiome simply remained majorly unperturbed, possibly due to the lower residence time of food in it. Now, I turn to the unannotated perturbation where the salivary microbiota’s similarity statistic does show a significant trend. One possibility, of course, is that there is a genuine perturbation during this period. However, the metadata is missing during this period, making validation impossible. Given that it seems to follow similar sensitivity patterns seen in the other recorded perturbations from the microbiota, this is still however a genuine possibility. Another possibility is that it is a dynamical perturbation that is changing the state and therefore the VCR, and is not necessarily a structural perturbation. Finally, another possibility is seen by looking at the pre-filtering raw data, noting that in the region of days 250-300 of the experiment, 16 days out of 50 are missing (including 13 from days 259-271). Therefore, it is also possible that these dynamics are coming about as a strange artifactual result of the linear interpolation. To test this, I apply a nonlinear interpolation algorithm based on simplex projection to fill in the data and replicate these results in Appendix A.1.

The main takeaway from this section is that while the kinds of profiles that respond with higher sensitivity (PP/PM) are seen in the 3 microbiota, this result is extremely noisy and can be emerging by chance. This requires further investigation. If this is indeed the case, the reason for this is not clear, and could come from the specific nature of the perturbation or from undetermined internal structure of the community.

4.5 Other Datasets

While this scale of inference is not possible on other datasets due to several aforementioned reasons, certain aspects of the analysis can be adapted onto datasets that satisfy some of the requirements. In particular, I highlight two other studies- one from a study that sought to assess whether the baseline composition of the human microbiota fluctuates over time and compare its composition across different organ systems (Caporaso et al., 2011). This dataset could not be used because it lacked records of perturbations or metadata- thus, the fluctuations seen could be arising from any factors, including from missing data or dynamical perturbation, and apart from identifying potential regions of interest, these do not tell us anything useful to validate the VCR or different profile sensitivities on. However, the dataset could be useful to assign the causal profiles and compare the frequency, as this does not assume any structural perturbations need to occur in the system. Another dataset is that explicitly includes an external perturbation is a study on the response of gut microbiota to antibiotic treatment (Dethlefsen et al., 2008). The daily count data for this experiment includes a week of pre-perturbation data, 5 days of antibiotic therapy, then a week of recovery data, making up a total of 19 days. However, because the dataset is so short, the causal profiling could be inaccurate as it takes a very small slice of the attractor and infers probabilities from that, which is both very constrained and very coarse. However, this could be used as an additional test of the similarity statistic at the whole-community level, without performing any causal profiling.

4.5.1 Moving Pictures of the Human Microbiota

This dataset is extremely rich and well-sampled. It includes data for 4 different microbiota for 2 different subjects (F and M)- skin microbiota from the left and right palm, salivary microbiota, and gut microbiota. Both subjects were sampled daily- Subject M was sampled for a period of approximately 444 days, while Subject F was sampled for a period of 186 days. Here, our aim is solely to analyze causal profile frequencies, and not carry out assessments of similarity over time. Similar to what was done for previous microbial data, the OTU count data is grouped at the family level, and families with sufficient rarity (no abundance for at least 2 weeks) are excised. The number of common taxa is listed alongside the causal profile frequencies below in Table 4.4.

The main focus here is the causal profile annotation, which is done by assuming the

Sample	MM	MP	PM	PP	Richness
F Left Palm	343	209	207	241	136
F Right Palm	473	255	156	116	102
F Saliva	397	260	186	157	25
F Gut	455	267	165	113	21
M Left Palm	531	213	131	125	57
M Right Palm	522	205	125	151	60
M Saliva	460	289	148	103	29
M Gut	447	189	226	138	25

Table 4.4: Causal profile frequencies from each microbiota sampled indicates a consistent trend of MM being the most common, then typically followed by MP, PM and PP.

same causal graph structure, computing the natural direct and indirect effects, and assigning profiles PP/MP/PM/MM to each sub-community. First, one thing that could be interesting to note for future studies is the high species richness in the skin microbiota. Given the ease of sampling the skin microbiota compared to the gut, long-term time series of the skin microbiota structure could be extremely insightful. Second, and more significantly, the causal profiles of a majority of the communities appears to follow a highly conserved and consistent pattern, that is, MM profiles are the most common, followed by MP, then PM and PP. This is identical to the pattern that we see in all 3 microbiota, suggesting that this pattern could be a general systems-level property of microbiota across individuals and across different parts of the body. However, due to the difficulty in ascribing the identity of potential perturbations to an actual structural effect, no further annotation or analysis of sensitivity is done.

4.5.2 Antibiotic Recovery

This data comes from a study of the effect of antibiotic perturbation on the composition of the gut microbiota (Dethlefsen et al., 2008). By giving 3 subjects a 5-day course of antibiotic Ciproflaxin, the effect on the gut microbiota was studied. This was done in 2 trials, 8 months

apart. Given the large duration between the two studies, I treat the two as independent. For both trials, samples were taken before the treatment to establish a baseline- however, these were taken months in advance and therefore represent only a single snapshot of the dynamics. Daily baseline data is available for a period of 1 week before the treatment- this can be used as the baseline state. In total, the daily data, as is required of EDM, is available for a total of 19 days- 7 before the treatment, 5 during the treatment, and 7 after. The short nature of the time series makes defining a baseline here is also extremely tough, as the distribution cannot be assumed to be Gaussian with such a limited number of time points, making the z-score based approach slightly shaky. Any causal inference also becomes challenging (as the probabilistic nature of the natural causal effects estimated is made extremely coarse by the small length of the time series). Therefore, I only test if the similarity statistic is able to detect the perturbations occurring here as well. The data is split by subject and each trial is treated independently, producing 19 days of data each. Given the extremely short nature of the time series, a baseline cutoff of less than 7 days of constant 0 abundances is used to filter the data. This leads to the inclusion of 14 and 16 taxa in Trial 1 and Trial 2 for Subject 1, 17 and 15 for Subject 2, and 10 and 12 for Subject 3. Picking a baseline size of 5 days and a window of 2 days, I plot the similarity statistic of six trials (2 for each patient) below in Figure 4.11.

Though results are very coarse and suffer from insufficiently sampled baseline values, in a majority of cases, there appears to be a decline in the similarity statistic occurring somewhere during the highlighted period of perturbation, followed by a subsequent period of comparative increase after the perturbation recedes. This timing of the minimum is different across different people and treatments, and can often, as noted before, be noisy- longer recovery phases may also be needed to fully understand the dynamics. While not extremely convincing on its own, in light of the success of the similarity statistic on the other longer dataset, this tells us that in general, the statistic is largely successful at identifying perturbations associated with compositional and structural changes.




Figure 4.11: Similarity statistic for antibiotic treatment. Despite the insufficient baseline of 5 days, a window size of 2, and a cutoff score of 2, almost all treatments appear to show at least some decline in similarity statistic and subsequent recovery after the perturbation.

Chapter 5

Discussion, summary and outlook

To make a prairie, it takes one clover,
and a bee, and reverie.

Emily Dickinson

First and foremost, I hope this thesis has given you a far deeper and more nuanced appreciation of both the nature and importance of microbiota, as well as the slew of incredible tools being developed to sift through the modeling mess that is natural ecosystems and extract general, robust properties. My mother likes to use the phrase ‘slipping a needle through a banana’ to describe something happening smoothly and subtly, which I think is an appropriate description of these methods on ecological time series. In this section, I synthesize and contextualize these results in a more general setting, and their implications for future studies of microbial communities. In addition, I extensively detail shortcomings of these methods, both as a cautionary signal for prospective users of these methods in future studies, as well as potential directions for future studies to strengthen, validate and implement these results towards a deeper level of understanding of microbial community dynamics. I first briefly summarize the main results and their motivation developed in the previous sections.

5.1 Synthesizing results, briefly

The complexity inherent in microbial systems has made and continues to make any modeling approaches difficult to apply- this necessitates the development of several novel techniques as ‘compromises’ to overcome limitations of many approaches. Therefore, this work unifies approaches from non-parametric dynamical reconstruction along with causal inference with the aim of trying to glean as much about the system’s dynamics as possible while assuming minimal amounts about the specific structure that produces those dynamics. At this juncture, it may be helpful to recall the aim we started out with- **what factors determine the sensitivity of microbiota to environmental perturbations?** From here, I have qualified environmental perturbations as those that perturb the internal vector field structure of the dynamics, leading to the selection of structural stability as the most appropriate metric to track the sensitivity of the system. While structural stability can be reconstructed by using the S-map and then returning the VCR (volume contraction rate) as the trace of each Jacobian, a weighted regression algorithm from empirical dynamic modeling, it cannot be done in large communities easily. While the S-map can be applied to smaller sub-communities, it is difficult to scale up whether changes observed at the single sub-community level actually have effects at the whole-community level. To do this, I develop the similarity statistic as a robust, generally applicable and highly sensitive method to identify and quantify the effects of external perturbations. Though similar methods have been applied to understanding resilience and recovery in ecosystems (Ingrisch and Bahn, 2018), they often assume that a system relaxes back to a fixed point instead of a distribution of values, and may not be suitable for large ecosystems where the responses of different sets of taxa are radically different. Moreover, many of these methods are optimized for system states, and not more descriptive properties like their sensitivity.

This method is extensively validated on both simulated (using a stochastic generalized Lotka-Volterra framework with an artificially induced structural perturbation) and real longitudinal microbial data, with the conclusion that the similarity statistic applied to the VCR is able to successfully, robustly and consistently detect periods of external perturbation as a reduction in its value. Furthermore, as the perturbation recedes, the value can frequently be seen to recover back to a higher-similarity state, albeit slowly in real systems. This is also validated on a short antibiotic treatment dataset. With the belief that the VCR acts as a metric of perturbations, and high degrees of variation observed in its values over time

for different communities, the question of factors affecting sensitivity therefore comes to the forefront. Specifically, could the effect of interactions and abundance on the VCR explain variations in strength of response seen among different sub-communities? This is done by using probabilistic causal inference, which binarizes the data, and then assuming a causal graph structure of abundance and interactions both affecting VCR and abundance affecting interaction, computes the natural direct and indirect effects in each sub-community. Then, based on the sign of the effect, these are assigned to different profiles.

I calculate the frequency of each of these profiles- in competitive simulated communities, the profiles with negative effects of abundance on VCR are rarer; however, in simulated communities with weak interactions, the trend inverts, with the profiles (PP/PM) having a positive effect of abundance on VCR becoming rarer. Though they cannot be simulated easily, I expect mutualistic communities to have a similar frequency to the weak-interaction regime. Real data from 3 main (David et al., 2014a) and 8 supplementary microbiota (Caporaso et al., 2011) shows similar results to the weak interaction case, where communities of PP/PM are seen to be rarer. Finally, I apply the similarity statistic separately onto different causal profiles- this shows that across the real microbiota, while PP/PM appear to be more stable for the given parameters of the data, this may not be a general property- this appears to not be statistically robust, depending on the choice of values taken. In the simulated systems, competitive communities showed a consistent trend of communities of type MM and MP being more unstable, while in weak interaction regimes, the identity of the most specific community appeared to depend upon the nature of the random perturbation itself.

While both Subject A's and B's gut microbiota respond strongly to the recorded perturbation, as tracked by the whole-community similarity statistic, this is less so the case for the saliva, which shows only a fairly brief decline in the similarity during the perturbed period, and with lesser separation between the profile responses. This could be due to the saliva being an average of many niches, or potentially also because of the residence time in the system. Since the salivary microbiota tends to experience a much more direct contact with food, and is constantly refreshed by comparatively lower residence times of novel bacteria, the structural effects guiding perturbations in the oral cavity may be distinct from the gut, where there is much longer exposure to food, and much more spatial extent. Previous studies (Zaura et al., 2015) have also shown the salivary microbiome to be more resilient to antibiotic therapies. Mouth infections have been shown to affect oral microbiota composition (Belstrøm, 2020),

but in the case of invasion of novel microbial cohorts which tend to affect functioning lower in the alimentary canal (Nguyen et al., 2021b), the salivary microbiota may remain potentially undisturbed.

5.2 Implications

What this work shows can be organized into two orthogonal axes- one in terms of what this means from an ecological standpoint, where we confront questions about how sensitivity to environmental perturbations changes with time in large communities, and secondly, from a clinical-adjacent standpoint, concerned with how microbiota can be made to be more stable against perturbations or recover from dysbiotic states faster.

5.2.1 Ecological Implications

One major highlight is the development of the similarity statistic as a non-parametric and generally tractable tool to assess the presence, effect and duration of perturbations on large ecosystems. Recent work has shown that multivariate early warning signals, based on detecting critical transitions, have limited ability to predict ecosystem changes (O’Brien et al., 2023a)- it is possible that this attractor distribution based approach may be able to overcome some of the issues that are associated with this, as it does not detect critical transitions but rather motion away from the state distribution of a baseline attractor. In fact, the S-map would not work for critical transitions, as it requires the presence of points on the same attractor to weigh the regression appropriately. A fascinating follow-up could be to see if the performance of changing similarity statistic as an early-warning signal for transitions could work on other datasets where EWS-based methods have had limited success (possibly because the system displays more of an attractor deformation type of response rather than a critical transition between distinct attractors). Moving on, the probabilistically driven causal profiling gives us a very generic way to assign the effect of factors governing a particular ecosystem variable. Though this is hardly an unexpected result, I validate that different subcommunities do indeed have different structures and effects of abundance and interactions on their stability. Total abundance is shown to have a role in structuring the sub-community’s sensitivity, but effects vary. This can be conceptualized by thinking about the intrinsic dynamics of different taxa. Certain taxa’s dynamics may be more intrinsically unstable, and therefore an increase in their abundance may cascade and destabilize the whole community. Conversely, an increase in the abundance of stabilizing taxa could tend to increase structural

stability by constraining perturbations in other species. More significantly, the role of mean interaction strength is also shown to have a role on the stability of a particular compositional state of the microbiota in a majority of communities, despite the mean off-diagonal element as a metric requiring averaging over large variation in interspecific interaction coefficients. This is in contrast to a recent study (George and O'Dwyer, 2023), which showed that several fundamental microbial ecosystem properties can be reconstructed by use of a stochastic linear model with no interactions between species. It is, however, possible that while abundance mostly determines the broad-scale macroecological properties of the system, the interactions play a role in fine-tuning it, and ensuring the maintenance of that structure against environmental stochasticity and variation (since we do show its role in specifically modulating structural stability). In support of this, though the sensitivity statistic across profiles appears to show grouping of responses on the basis of the abundance effect sign, the interactions do play a role in smaller-scale variation between the profiles. Another significant inference from the interaction's causal effect on stability is the conflicting role of increasing interaction strength on structural stability. Classically, increasing interaction strength is regarded to cause unstable feedback loops, reducing the dynamical stability of the system. However, positive interactions could benefit structural stability as well (Rohr et al., 2014) by increasing the size of the feasibility domain. Given these opposing effects, it is natural to expect both cases co-existing in different subcommunities- while the size of the feasibility domain may increase, it is also possible to have a scenario where the community is not able to remain in said feasibility domain due to changes in its orientation or local structure (Desallais et al., 2024). Another case that could explain this is that in some communities, there may not be an explicit increase in mutualisms per se, but an increase in the strength of the positive term in predation effects (+/-) which can be highly stabilizing (Allesina and Tang, 2012; Song and Saavedra, 2018). This warrants further exploration to delineate the exact ecosystem-level changes in the interaction matrix occurring that lead to specific outcomes.

Another interesting ecological perspective on this comes from looking at the frequencies of different causal profiles with different parameter values and in different settings. From the generalized Lotka-Volterra simulations with different interaction strengths, we note that in competitive communities where mean interaction strength is highly negative, the system tends to have more numbers of communities where increases in abundance tend to destabilize the community (PP/PM). While the mechanistic reason for this is unclear, it could be related to the fact that in a competitive community, decreases in growth rate of a particular species

due to structural perturbation could push up the abundances of a large number of species, increasing the overall population-level extent of competition and thereby causing structural destabilization. However, as the interaction strength is tuned up towards zero, there is a continuous transition to a state where MM and MP become the dominant communities and PP and PM become rarer, which appears to be a consistent hallmark of weak interaction regimes. It is interesting that the case of equal frequencies of each profile does not occur in the case where the mean interaction value is 0 but requires instead low levels of competition. The reason for this is unclear, and could be a consequence of the appearance of positive interactions in a significant amount causing unfavourable feedback loops in weak interaction regimes. Though it is not possible to generate communities with average positive interaction strength given our current model framework, I suspect that as mean interaction strength is further tuned up, the results will continue to show a dominance of MM and MP type communities. However, there will likely still be some cases with the PP/PM type of interaction, as is seen in the case when competition is increased as well that MM/MP type communities persist.

This is not a particularly surprising result on its own, suggesting solely that different interaction regimes have different effects of interaction and abundance on sensitivity (though the framing of this as effects on structural stability is rather non-trivial, and needs to be carefully worked out). However, the continuous transition and the fact that it balances out at an intermediate value are certainly future questions to be resolved. Importantly, this also establishes an expectation set of values for indicating what sort of interactions may be dominating in a community, assuming that Lotka-Volterra dynamics are indeed a good approximation of its behaviour. Real microbiota appear to be dominated by MM/MP type interactions, indicating that these are likely not driven predominantly by competitive interactions, an idea that is coming to the forefront more and more with more microbial interaction studies (Coyte and Rakoff-Nahoum, 2019; Figueiredo and Kramer, 2020; Blasche et al., 2021).

Another important ecological insight comes from applying the similarity statistic separately to each causal profile and comparing the sensitivities during periods of perturbation. In competitive communities, it appears that the rarer profiles in terms of frequencies appear to be the ones that respond with more sensitivity. While a cursory analysis shows this may be true in real microbiota, this result is extremely noisy and is not robust to changes in system sampling and filtering parameters. However, if it is true, this could be suggesting that

real microbiota tend to structure themselves to possess a majority of more stable microbiota substructures (and therefore more stable at the whole community level), while retaining some high-sensitivity cases, either as a necessary consequence of the network construction, or in order to act as a potential pool of standing high-sensitivity internal substructure that could easily move between states if a perturbation were to occur. However, the weak interaction regime appears to show no consistent identity of more or less sensitive communities, instead being determined possibly by the specific network structure and the type of perturbations. Assuming the PP/PM profiles to genuinely be more sensitive in real microbiota, given the alignment of causal profile frequencies with weak regimes, this inconsistency could be emerging due to two potential reasons-

- While the fact that any profile could be more sensitive to a random perturbation is seen in the case of weak-interaction communities in the GLV setting, this need not be the true structure of real data. Instead, the hypothesis that the underlying microbial community structure is similar to weak interaction Lotka-Volterra simulations is made solely on the basis of the similarity in the frequencies of causal profiles, and previous parametric model based studies. However, since this method is based on EDM and assumes no parametric structure, the underlying dynamics could be radically different from any Lotka-Volterra framework at all, and simply happen to produce similar results. Even if the Lotka-Volterra model is a good approximation of the community dynamics, it could be possible that the true microbiota dynamics tend to be more mutualistic while preserving similar causal profile frequencies. In such a case, similar to how the competitive community has higher sensitivity of the rarer MM and MP type of profiles, the mutualistic community may also have higher sensitivity to rarer PP and PM type of profiles as an inversion of effect of sorts.
- Another possible factor is the nature of perturbations themselves. In the Lotka-Volterra simulations, we take random Gaussian perturbations, and consider that in the weak regime, the most sensitive profiles appear to be determined by the specific nature of the perturbation. However, in a real ecosystem, the perturbations may be directed and preferentially increase or decrease the growth rate of a large number of species (possibly determined by their functional roles). Therefore, a stereotypical type of perturbation occurring in real ecosystems during all these diverse types of disturbances could also lead to a stereotypical response in terms of the sensitivity of different profiles. This effect could be complex- despite the nature of the perturbation in A's gut microbiota

seemingly different from that of the other two (in terms of the increase/decrease in VCR, as shown in Figure 4.9), with VCR declining on average in A's gut versus increasing in the other two, the results on sensitivity appear to be the same.

Formalizing a simulatable scenario that consistently produces similar trends and responses to real microbiota is therefore a future goal- consumer-resource models could be an alternative starting point that could provide a lot of insights, although cross-feeding interactions in the gut could complicate this. However, dynamical models explaining microbiota resource uptake in the presence of cross-feeding have been developed and used successfully to recreate ecological dynamics (Fant et al., 2021). Identifying whether certain communities are indeed consistently more or less stable in real microbiota, or whether this identity is dependent on specific facets of the perturbation remains a future area of focus.

Finally, I theorize why the higher sensitivity of PP/PM type communities may be the case, if it is not genuinely a true trend and is instead a statistical artifact. One option is that during periods of perturbations, taxa that are typically very rare during normal periods have a massive increment in (relative) abundance, while the decline in common taxa is less in terms of percentage increase. This means that since these species show a much greater abundance value in these cases, the general response of sub-communities across all profiles containing these species during perturbation periods would be to increase in abundance over their median values. Thus, for PP or PM-type communities that contain these sorts of highly-increasing taxa, the community is highly destabilized. MM and MP type profiles containing these are highly stabilized; however, since these highly increasing taxa are not particularly common and much of the core microbiome decreases by some extent, due to the negative effect of abundance increases on VCR, these communities destabilize. This could also explain why the trend is not robust to picking larger sub-community sizes such as 7 or 10, as the increment to abundance induced a particular taxon tends to be cancelled out by a large number of other taxa declining. However, this requires further validation and investigation, especially whether PP/PM type communities are enriched in taxa that have this sort of increasing pattern, and when other taxa are able to weather this effect.

5.2.2 Implications for Clinical Interventions and Microbial Therapeutics

While these results and their implications are in line with our understanding of the human microbiota as an ecological system, and posits a method to identify perturbations and what

could potentially variation in sensitivity to them, quite a bit more work may be required before these methods can be applied as diagnostic or therapeutic tools for recovery or stabilization of the gut microbiota. Control of microbiome dynamics has been a long-standing goal in medicine, with approaches borrowed from network modeling and control theory to identify breakpoints and nodes in metabolic pathways (Angulo et al., 2019; Y.-Y. Liu, 2023). In addition to understanding the specific effect of perturbation types, it remains to be confirmed whether real microbiota having higher numbers of PP/PM type profiles compared to others are more sensitive to reaching dysbiotic states; however, without further understanding of what exactly constitutes these communities in terms of members, functional roles, or local interaction structure, it becomes difficult to assign these numbers without prior long-term sampling, especially given extensive variation in microbiota between different subjects. This is because the causal profiles are ascribed based on past data. Also, since causal profiles are assumed to be fixed over time, these comparisons may only be done at the level of comparing two microbiota to each other as opposed to tracking periods when the fraction PP-type communities increase in a particular subject's gut microbiota. This could potentially be overcome by splitting the time series into different sufficiently long segments and assigning causal profiles to the same communities in each segment. Though disjoint partitions would require a tremendous amount of data to be effective, this could be scaled up into some sort of rolling-window approach to detect where causal profile frequencies tend to fluctuate and change. This requires careful application of the do-calculus axioms to not violate any causal inference properties. Regardless, this method on its own also could have some therapeutic benefits. By assigning profiles to sub-communities and identifying the dynamics of these select bacterial taxa as biomarkers, one may be able to track some amount of behaviour. For example, if a particular community (or set of communities) that has been identified as volatile (regardless of profile) is shown to have an increase in abundance, this may be indicative of incoming large-scale destabilization in that entire microbiota's dynamics.

Another aspect limiting the application to therapeutic outcomes is the lack of any species-level quantification. This is omission by design rather than lack of consideration- there is a great deal of diversity of taxa not just at the level of presence but also abundance between individuals, and identifying promoters of sensitivity consistent across microbiota at the single-taxa level could be both extremely difficult from a methodological standpoint, as well as ecologically non-tractable, given emergent properties in microbiota at the community level (Goldford et al., 2018; C.-Y. Chang et al., 2023). This study therefore looks for unified

approaches at the community level, since species-level metrics may be personalized and not generalizable. However, given this framework established thus far, this opens up avenues for future trait or species based annotations.

Another difficulty is the requirement of a robust baseline sampled before the perturbation actually occurs in order to model its effects. Given that clinical interventions typically occur only upon the onset of perturbation, and sourcing prior data in the absence of any ailment on a daily basis is logistically challenging, this presents one significant limitation. However, this is less of a problem if the sampling can be done from the day of the perturbation starting, ending when the perturbation ends, and using this as a baseline. This inverts the problem, and the similarity statistic now tracks the extent of recovery of different profiles and the whole community. In such a setting, given that MP and MM type communities could potentially be less sensitive, one interventional form of therapeutics upon the onset of a perturbation could be specific measures that enrich or deplete these specific communities' abundances. Other future directions could be to relate different community profiles with clusters of functional roles and pathogenicity to isolate the mechanistic basis of differential profile volatility.

All may not be futile at the species level, though- a recent study (Medeiros et al., 2023) highlighted a technique of ranking species by their sensitivity to perturbation under non-equilibrium conditions using EDM: relating individual sensitivity contributions from this to causal profiles or functional roles could provide mechanistic insights into the sensitivity of particular taxa. Some studies have also looked into integrating metabolic factors into EDM (Munch et al., 2023) as a means of generating more biophysically relevant predictions: applying a similar amalgamation of physiological factors with EDM may also lead to potentially more interpretable results.

5.3 Limitations

As with any method or result, it is important to question the limitations and caveats with its application. As touched upon in the previous section, one shortcoming of the similarity statistic is that defining a sufficiently long baseline that is indicative of a pre-perturbation state is a significant factor. This requires both a good amount of sampling as well as belief that this is truly indicative of the system's baseline- for a system continuously subjected to large-scale external disturbances, this may be extremely difficult. Moreover, given no

further information about the system beyond the abundances of different taxa to identify perturbations, establishing this baseline and whether a recorded perturbation is genuinely an indicator of a structural change becomes significantly more difficult.

Another shortcoming is the extremely limited sample size of valid microbial data that satisfies the necessary conditions for use- in effect, we have only 3 microbiota to study. While these do show a variety of responses, addition of other longitudinal microbial data could potentially show further and novel variations. Specifically, given the logistical and ethical difficulties associated with producing such dense, long-term datasets in human gut microbiota, replicating these experiments with model systems such as fruit flies and mice could be a valid further direction. Such datasets have been generated in some recent studies (Turnbaugh et al., 2009; Gibson et al., 2021), with recorded structural perturbations in terms of diet and antibiotic treatment as well. These, however, do carry the problem (as well as advantage, in some sense) of regulated laboratory conditions, stamping out some of the noisy and heterogeneous continuous disturbances occurring to the microbiota by human lifestyles. It is potentially possible that these three microbiota show a particular trend, but similar to the case of weak interactions in the GLV simulations, in other datasets, other outcomes are also likely to be possible, and this consistency is occurring by chance.

Next, I move onto possible methodological limitations, beginning with EDM-based approaches. The Jacobian can be sensitive to parameter values and embedding dimension choices in some regions of the attractor, especially where trajectories can be densely bundled together. Small changes to the data such as changing the interpolation technique to fill in missing values, changing the value of θ selected, or adding in alternative kernels or regularization terms (Cenci et al., 2019) may change the results observed in some communities. Some understanding of this could come from quantifying error terms in the Jacobian using methods outlined in previous papers (Cenci and Saavedra, 2018b). It may, in general, be worthwhile to add regularization terms to the data, to ensure more accurate Jacobian coefficients. However, when applying this to 1000 sub-communities over a long time-series, this can become highly computationally intensive (especially if one with a cross-validated parameter such as elastic-net, H. Zou and Hastie (2005)), and raises the question whether the same regularization scheme is appropriate for every sub-community. Another angle presenting difficulties is that the sampling approach of 4 communities on account of the S-map’s limitation to work in high dimensions coupled with the fact that it returns a Jacobian only directly proportional

to the true one limits easy interpretation of community-level trends. This approach is also fairly data-intensive, requiring long time series. However, by adding in EDM methods such as dewdrop regression (Hsieh et al., 2008), this could potentially be extended to be applicable to more datasets. Another weakness of the S-map is that it implicitly assumes the presence of the system remaining on a single attractor over time- this is because when weighing states to fit the Jacobian, the presence of a critical transition between attractors can lead to nonsensical results from combining states from the two attractors if explicit delineation of states on each attractor cannot be done. Finally, the S-map also assumes stationarity of the sampled time-series, which may not always necessarily be the case, for example, if some physiological parameter governing the system gradually changes as the person ages. However, recent extensions to EDM (Gee et al., 2023) have proposed methods of overcoming this.

The causal inference methodologies also carry several critical defects. One particularly noteworthy one is that the causal effects are assumed to remain fixed over time- in rapidly fluctuating natural ecosystems, this may not be true. In Section 5.2.2, I propose a technique to potentially overcome this by using rolling windows and tracking causal profiles in each, but this requires a lot of data and may not be extensively reliable. In the current regime, however, signatures of perturbation onset or recovery associated with changes in local causal profile within the same community may be eclipsed. Another limitation, and one that could be causing the failure of the causal discovery model, is the binarization of the data. Since this binarizes about the median, fluctuations that are barely above the median are treated identically to major spikes, and small increments or decrements to the raw value caused by minor environmental changes like food intake may be cancelled out. Developing a bespoke causal approach based on derivatives that admits continuous values could facilitate better resolution of causal results in this system. Successful causal discovery would also be extremely beneficial as it allows the validation of the structure of the causal graph. Adding lags in the metadata to try to explain causal effects could also be insightful, as the effects of lifestyle changes may take time to actually affect the ecological structure of the microbiota. Finally, one subtle potential fault when assigning the causal profiles is that while each causal effect certainly has a sign, it may not necessarily be significant. This can be validated by using a G^2 -test: it may provide some more resolution into the dynamics of each profile if these non-significantly signed communities are excised.

Finally, a limitation with the sampling methodology is that these 4-taxon communities

are formed ad-hoc from a larger taxon pool and may not necessarily exist in reality, given extant spatial separation in the gut- however, this is a problem at the level of the whole ecosystem as well. This presents a problem for any sampling-based microbiota study that uses samples of fecal/salivary matter without actually using tissue samples from the relevant organ section (which is naturally impossible in a majority of human gut studies). However, I do not consider this as a major potential bias, since many physiological changes, especially to external perturbations occur in a concerted manner affecting a whole bacterial metapopulation, and if two particular taxa are indeed spatially separated and regulated by distinct ecosystem processes, their dynamics will be independent and thereby the interactions inferred by the Jacobian will be minimal.

5.4 Future Directions

This thesis has only been a preliminary step into understanding the dynamics of perturbation response and recovery in complex nonlinear ecosystems. In this section, I highlight potential further analyses that could elucidate a deeper understanding of the properties of these systems. These are suggested at multiple different levels of the analysis, and while not all of them may be feasible, or produce meaningful and complementary results, I trace them out here in any case as a logbook of potential ideas.

First, I think about the microbiota itself as a system, and what further questions may be useful to ask. For one, it could be extremely fruitful to scale this down to the species level and look for distinctive trends. One such analysis could be whether the same taxa in each microbiota have similar dynamics, and if not, why so. Another approach could be to compute species-level sensitivity rankings, similar to the approach taken in a recent study using EDM (Medeiros et al., 2023), and see how this changes over time as the perturbation begins and ends. Another recent study (Medeiros and Saavedra, 2023) takes this further and decomposes the VCR into terms of time-varying individual species variability terms and species correlated responses, the linkage between changes in the abundances of different species, which could also tell us the role of specific interactions in structuring the system. It could also be interesting to see if the responses of different taxa to perturbations, possibly quantified using the similarity statistic, are phylogenetically correlated (similar to Sireci et al. (2023)). An alternative approach could also potentially to perform some sort of causal discovery of individual taxa abundances on other variables (including other taxa abundances)- this could identify taxa that respond early and recover early and potentially cause others to

deplete/recover, thereby acting as indicators of sorts. On another note, it could be insightful to compare the fractions of taxa in different profiles, and see if there are consistent patterns of occurrence (or co-occurrence between taxa) that are consistent predictors and indicators of a particular causal profile.

Yet another potentially interesting direction to lead this analysis would be to apply EDM methods and see if there are dynamical linkages between salivary microbiota changes and gut microbiota changes, considering the movement of food and saliva into the gut. While the original analysis by David et al., 2014a did not show any connections at the raw abundance level, there may be EDM-based approaches that successfully predict some aspect of gut microbiota processes with changes in the saliva microbiota as a prior warning of sorts. This could revolutionize biological monitoring, given how much easier it is to produce salivary samples compared to fecal samples. Taking this further, building a working baseline for compositions of coexisting communities at different positions in the alimentary canal could enable building spatially explicit metapopulation-based models of perturbation, and assessing how perturbations may propagate down to more or less critical or sensitive regions. Finally, it is worth noting that most of these bacterial OTUs have not been successfully cultured in a laboratory environment, a factor that is suspected due to their association with very specific cross-feeding partners and environmental conditions (Turnbaugh et al., 2007; Tanaka and Benno, 2015; Lagier et al., 2016). These studies that look into interaction strength variation while explicitly looking into environmental variation could provide a framework for identifying potential conditions and partners which would allow the growth of these taxa, permitting further metabolic, functional and genomic annotation. For example, the phylum *Saccharibacteria* (or TM7), resident in the oral microbiota, and recorded as one of the more common taxa, has been only recently cultured in a lab, and found to display fascinating epibiont behavior, dependent entirely on attaching to other host bacteria for survival (Bor et al., 2019). Rare species can also have significant effects on the ecological dynamics (Säterberg et al., 2019), especially if they take the form of keystones (Wang et al., 2023)- including the previously excluded species and working around their sparsity could provide higher-resolution windows into dynamical changes.

Next, we turn to the EDM based approach and the stochastic simulations. For one, it may be interesting to look into how variation in other simulation parameters affect the results. Another issue that has been brought up several times earlier is to add negative growth rates

and to replace the linear functional response with a sigmoidal Holling's functional response in the standard Lotka-Volterra model, aimed at stabilizing positive interaction settings. This allows us to compute sensitivity and causal profiles in mutualistic communities as well, and compare with weak-interaction and competitive cases. Another potential angle would be to modify the EDM methodology slightly to perform a CCM (convergent cross-mapping) between the taxa abundances *a priori* to construct a more pruned and robust interaction network before applying the S-map. This approach is more similar to what some other EDM studies of species networks have carried out (Ushio et al., 2018; Zhao et al., 2023; Merz et al., 2023). We do not use this approach, since the interactions we compute are based on aggregate effects, and thus, causal interaction mapping could potentially lead to spurious results by mixing together indirect effects. This is potentially less of an issue in communities with stronger trophic structure; however, given overlapping metabolic pathways and cross-feeding behavior observed in the gut microbiota (Culp and Goodman, 2023), we avoid this approach. Regardless, it could be illuminating to apply these methods to a CCM-reconstructed interaction network and seeing if any differences in dynamics are observed. In accordance with this, it could also be interesting to compare these results with the output of the MDR S-map (Chang et al., 2021), a method that uses multiview embedding (to generate multiple possible permutations of lower-dimensional state-space reconstructions, as per Ye and Sugihara (2016)) along with the regularized S-map (Cenci et al., 2019) to model high-dimensional ecosystem responses. We have avoided this because of potential concerns with noise, but this could also show similar results.

Another important direction forward would be naturally to validate if the trend seen in the real microbiota is true, that profiles PP/PM could potentially be more sensitive. This could be done either by development of mechanistically motivated models, experimental validation, or more sophisticated and careful statistical analysis on these and other datasets. One way to increase the pool of viable datasets is to move beyond focusing specifically on human microbiota. It could be interesting to look into this in other systems such as the ocean, other organismal and bodily microbiotas, as well as macroecological scales such as vegetative dynamics. Recent work on community transitions (Long et al., 2024) has shown a distinction in vaginal and ocean microbiota compared to salivary and gut microbiota, possibly based on the seasonality of the dynamics. Another potential experimentally-driven angle would be to look at cocultures of these taxa and compute their traditional Lotka-Volterra interaction strength (this can be done by growing one alone, and then in the presence of the other, and

evaluating ratio of final outcome). This gives us the raw pairwise interaction, which can then be subtracted from the Jacobian-determined interaction strength to assess the role of higher-order interactions in this system, and if they facilitate the display of responses seen. Previous work (Gibbs et al., 2022) has shown that higher-order interactions can promote coexistence in diverse communities. From this, we may revisit the attractor, and potentially fit polynomials in lieu of a hyperplane to test if these could act as potential coefficient matrices for higher-order interactions too.

To sum up, our work provides a set of possible metrics and results that can be compared across microbiota to look for commonalities, aimed towards building a general framework for catalyzing ecological recovery in microbiota. This work highlights the key role of interactions in shaping microbial community sensitivity, describes a method of identifying perturbations, and suggests a method of partitioning gut bacteria into profiles on the basis of the role of fluctuations in their community-level parameters of abundance and interaction strength on sensitivity as a possible determinant of their variation in response. Given a sufficiently rich time series, this tractable framework could be used effectively towards designing interventions for a range of large ecosystems undergoing external perturbation.

Appendix A

Robustness Checks

“Nurse”, he said, “it’s an appendix!”

Madeline, Ludwig Bemelmans

In producing the above results, several schemes, parameter values and methods have been assumed in order to carry out this pipeline. While these choices are largely made based on logical assumptions, it is important to test that other changes and definitions of these parameters do not lead to drastically different results. This section of the appendix is mostly concerned with this. Specifically, I confront the following questions-

- To fill in missing data points, I use linear interpolation. Another alternative way to fill in these missing values would be to use some kind of nonlinear interpolation algorithm. While spline produces frequently negative results, unrealistic for real ecosystems, I adapt the framework of simplex projection to fill in missing values, and confirm that results remain qualitatively similar.
- I grouped results at the family results as a means of dimensionality reduction while preserving metabolic function. What happens if this grouping is done at the genus or order level instead?
- To remove rare families, I set a cutoff of no more than 14 days of no zero abundances

as a threshold. What happens if I make it stricter (7 days) or looser (21 days), leading to inclusion of different subsets of the taxon pool?

- Finally, this analysis has been done taking sub-community cluster sizes of 4. What happens if we increase the size of these? Are results unchanged?

Our major three results are the similarity statistic, the causal profile frequencies, and the differing sensitivity of each causal profile. Thus, for each of the above cases, I demonstrate these three. While there are additional cases combining the two that could be tried out (for example, grouping at the order level and using a cluster size of 7)- however, naturally, due to the prohibitively large number of combinations that this would produce, I do not carry these out.

A.1 Simplex Projection to fill in missing values

The idea of simplex projection is described in Section 2.2.2- basically, the idea is to predict the fate of a focal point by looking at the fate of a given set of neighboring points, and taking the average. Here, I use simplex projection in a slightly unconventional application, to interpolate missing values in the microbiota dataset of the 2 subjects in the study. This is useful because it interpolates nonlinearly, so any trends that could potentially be emerging from the linear interpolation of missing data are removed. Specifically, since every sub-community is embedded in 4 dimensions, for each embedded point right before a missing value, I find the 5 nearest neighbours by Euclidean distance from the focal point, and take their average as the next state. This is sequentially repeated for each missing point. This is a slightly unnatural application of the technique, considering I use future values to fill in past values, but this is still logically consistent since it is within the same time series, and takes only the distance into account. This fills in the data and allows S-map reconstruction the same way as before. Then, by cutting out rarer families according to the 2-week cutoff previously used, this leaves us with 38 taxa in A's gut, 44 in A's saliva, and 21 in B's gut. Note that this is more than the linear interpolation case, because the simplex can potentially predict a non-zero state between two zero values if the distance metric predicts this, which is not possible in linear interpolation- thus, a few more taxa are retained. Then, reconstructing the VCRs, abundances and interaction strengths, the similarity statistic and causal profiles are computed. Figures A.1 and A.2 display the similarity statistic over time for each of the microbiota at the whole community level and then split by profiles, and Table A.1 lists the

frequency of each profile.




Figure A.1: Similarity Statistics for 3 microbiota with simplex interpolation. Note the response during the perturbation being roughly similar, but other trends disappearing- for instance, the large dip observed in A's saliva around Day 275. This shows that the success of the similarity statistic to pick up known perturbations, in any case, is not affected by interpolation schemes.

Clearly, qualitative results to the whole community remain largely unchanged- the similarity statistic still effectively picks up the periods of perturbation. However, notably, the major perturbations that were seen around Day 275 for Subject A's salivary microbiota have disappeared, leaving behind a constant but low similarity statistic as opposed to the significant dip seen earlier. This is also true of the smaller perturbation seen in Subject B's gut. This suggests that these could potentially be the result of artifacts introduced by linear interpolation, and not major genuine perturbations. However, the simplex projection appears to generate a few new perturbation periods, such as immediately after the perturbation period, as well as a large dip seen some time after Day 200. Given this, it is hard to identify and assign specific identities to unannotated perturbations in A's saliva. The profile frequencies also remain preserved across the previous case of linear interpolation, with MM profiles dominating, and PP and PM generally being rarer.

Looking at the results of Figure A.2, we see that qualitative results remain somewhat unchanged but become extremely noisy- in Subject A's and B's gut, the similarity declines strongly, preferentially affecting PP and PM type communities, but the profile MP also appears to respond highly in B's gut. In Subject A's saliva, the trend around Day 275 disappears, and the profile PP/PM, contrary to expectation, appears to become more *stable* during the period of perturbation. However, simplex projection appears to pick up a curious response immediately after the perturbation ends, wherein PP/PM respond more strongly. This is not seen in the original linearly interpolated data, and is especially strange because

Microbiota	Profile MM	Profile MP	Profile PM	Profile PP
A's gut	475	245	148	132
A's saliva	441	271	182	106
B's gut	357	228	206	209

Table A.1: Frequencies of each profile in different microbiotas. Note the consistent trend across all 3 microbiotas, exactly qualitatively similar as before- MM is the most common, then MP, then PM, then PP.




Figure A.2: Profile-based sensitivity for the microbiota. Note that trends remain roughly unchanged, but the higher sensitivity of Subject A's saliva in the perturbed region seems to switch somewhat. There also appear to be new major perturbations immediately after the perturbation and somewhere in the middle.

this region (as well as the region around Day 200 with another spike) does not have any major missing periods of data. Therefore, it may be challenging to infer anything meaningful or specific about these novel perturbations without further investigation. However, it is important to note that while the profile-based sensitivity is extremely noisy, the results on

causal profile frequency and whole-community similarity remain highly consistent with the original case, proving that these are not affected by the values filled in by linear interpolation.

A.2 Other Taxonomic Levels

In this section, I group the data at the genus and order levels. Going further to species or class level may be either too coarse or too specific to make any generalizable conclusions. The same cutoff period of 2 weeks for rare species is used in all cases. This leaves 17 orders and 51 genera in A's gut microbiota, 24 orders and 61 genera in A's salivary microbiota, and 10 orders and 28 genera in B's gut microbiota. Once again, the S-map is applied and the pipeline of similarity statistic and causal profile calculation is applied. Results are shown below in Figure A.3 with both whole-community and profile-level similarity statistic, and Table A.2 for the case of the grouping at the order level. Results for the family level are seen in Figures A.4 (both whole and profiles) and Table A.3 beneath that.




Figure A.3: Order-level grouping. While the whole-community level similarity statistic remains highly robust and effect, the profiles become highly noisy, possibly because of the coarse-graining of different orders.

We see from these two analyses that major trends remain largely unchanged- the causal profiles and similarity statistic remain successful and consistent. However, the profile similarity statistic becomes far noisier- looking at the order level, in A's gut, profile type PM remain more sensitive, while PP is indistinguishable from the others. A's saliva during the perturbation

Microbiota	Profile MM	Profile MP	Profile PM	Profile PP
A's gut	402	273	179	146
A's saliva	347	249	240	164
B's gut	370	290	221	119

Table A.2: Frequencies of each profile in different microbiotas, grouped at the order level. The trend remains consistent with the original annotation.




Figure A.4: Genus level. Here also, we note the success of the whole-community similarity statistic, but observe that while the profiles generally seem to respond, the response is much more noisy and mixed here as well.

period is incredibly messy, with all profiles responding together and no consistent behavior—however, the strange perturbation around Day 275 continues to show higher PP/PM sensitivity. B's gut shows profiles PP, PM and MP all responding very similarly and strongly to the perturbation. This is also the case for the genus-level grouping, which shows similar trends to the order level in terms of these responses becoming extremely noisy and inconsistent. While these could likely be a consequence of coarse-graining/addition of noise arising from this potentially not being the optimal taxonomic level, this shows that the causal profiles in

Microbiota	Profile MM	Profile MP	Profile PM	Profile PP
A's gut	549	256	119	76
A's saliva	389	279	183	149
B's gut	314	209	242	235

Table A.3: Causal profile frequencies at the genus level- note the same consistent trend across all 3 microbiota also seen in the family-level and order-level data.

general may not be the most accurate explainer of variation in sensitivity response during perturbations, though there could potentially be some sort of insight from this.

A.3 Differing Rare Species Cutoffs

Once again, 2 alternative cases are assessed here, pertaining to the sparsity cutoffs for inclusion of a particular family in the final datasets. Since the original cutoff was no more than 2 weeks of no abundance, I use one cutoff of only one week of zero abundance permissible (stricter) and one of three weeks of zero abundance permissible (looser). The stricter one-week cutoff retains 26, 35, and 14 taxa in A's gut, A's saliva and B's gut respectively. The looser 3-week cutoff retains 37, 45 and 20 taxa in the same order. Note that these are not too different from the 2-week cutoff- this is because a vast majority of taxa are extremely rare, being recorded only during a few days and being zero otherwise. Figure A.5 shows the similarity statistic both with and without the causal profiling for the 1-week cutoff, with the frequencies recorded in Table A.4. The three-week cutoff can be seen in Figure A.6 and Table A.5.

Seeing the plots, we do not see much qualitative difference between these and the original results- the similarity statistic still successfully logs the perturbation, the causal profile frequencies remain consistently ordered. The causal profile sensitivity appears to show PP/PM as more sensitive here also, though A's saliva during the perturbation period is highly inconsistent. However, the inclusion/exclusion of these rarer species may not have much of an appreciable effect on the average subcommunity's dynamics, so this is to be expected.

Similarity Statistics: One Week Threshold for Rare Species




Figure A.5: Similarity statistic plots with a 1-week cutoff. Note minimal qualitative change from the 2-week case- responses remain largely the same, except the saliva's perturbation period appearing to be much noisier, and the profiles being much more similar in B's gut, though PP/PM remain more sensitive.

Microbiota	Profile MM	Profile MP	Profile PM	Profile PP
A's gut	468	255	179	98
A's saliva	395	275	174	156
B's gut	337	229	206	228

Table A.4: Frequencies of each profile in different microbiotas: one week cutoff. The trend of MM being the most common and PP/PM being rarer in general continues here as well.

A.4 Larger Cluster Sizes

Finally, all this analysis has chosen sub-communities with a fixed size of 4 taxa. What happens if we choose larger values? I do not pick smaller values, as then the community structure becomes too constrained and not diverse enough to show rich dynamics. Picking clusters too large will show effectively random dynamics (Cui et al., 2021). Below, I show

Similarity Statistics: Three Week Threshold for Rare Species




Figure A.6: Similarity statistic plots with a 3-week cutoff. Subject A's saliva appears to show no consistent response here neither in the perturbed region nor clearly outside, though Subject A's gut and B's gut both display consistent trends.

Microbiota	Profile MM	Profile MP	Profile PM	Profile PP
A's gut	487	239	169	105
A's saliva	363	260	231	146
B's gut	342	280	195	183

Table A.5: Frequencies of each profile in different microbiotas: three week cutoff. The trend of MM being dominant with PP typically being the rarest persists.

results for cluster sizes of 7 (Figure A.7 and Table A.6) and 10 species (Figure A.8 and Table A.7). Note that when taking larger community pools, a smaller portion of the total combinations get represented by 1000 samples (as the ways of picking 10 species out a pool of 30 is far higher than the ways of picking 4- 1096 times more ways!) Therefore, there may potentially be minor biases emerging here from the extremely small pool of total communities sampled, but it is worth noting that 1000 groups of 10 species is likely to sample interactions between any 2 species multiple times, providing a complete picture.




Figure A.7: Results with a subcommunity size of 7 taxa. Note that while the similarity statistic on the whole community continues to operate with high fidelity, the results on the profile similarity statistic appear to become inconsistent.

Microbiota	Profile MM	Profile MP	Profile PM	Profile PP
A's gut	442	286	175	97
A's saliva	334	259	226	181
B's gut	350	250	167	233

Table A.6: Frequencies of each profile in different microbiotas with a cluster size of 7. The trend remains conserved.

First, we note that the causal profile frequencies and similarity statistic on the whole community remain highly similar to the original case, indicating that these are genuine trends that hold in any case. Therefore, it would be safe to infer that the similarity statistic can pick up whole community perturbations, and the causal profile frequencies of microbiota follow a fixed pattern. However, the causal profile does not appear to be a consistent predictor of




Figure A.8: Results with a subcommunity size of 10 taxa. Similar to the case of 7 taxa, the similarity statistic on the whole community succeeds, but the causal profile sensitivity is not consistent with what is seen before.

Microbiota	Profile MM	Profile MP	Profile PM	Profile PP
A's gut	442	286	175	97
A's saliva	334	259	226	181
B's gut	350	250	167	233

Table A.7: Frequencies of each profile in different microbiotas, sampling 10 taxa each into a subcommunity.

variation in sensitivity of response, as evidenced here. In principle, selecting larger taxon pools should not have much of an effect on the observed dynamics, unless a prohibitively large size is taken (in which case there may be random behavior, as previously mentioned). However, 7 and 10 should logically be fairly stable given the richness of roughly 30 taxa seen otherwise. In the cluster-7 case, we see that A's gut continues to have PM slightly more

sensitive, but PP stabilizes and is no longer responding distinctively. A's saliva appears to have the opposite during the period of perturbation- PP responds strongly, but PM appears to be indistinguishable from MP/MM profiles. While the PP/PM communities do respond much more strongly during the later unannotated perturbation, given the results from simplex projection in Appendix A.1 completely removing this dip, I do not treat this as a potential validation, and it is entirely possible that this trend is an artifact of missing data. Finally, the trend also completely disappear in Subject B's gut, which shows an appreciable dip, but profile PP seems to stabilize while MM seems to be responding more strongly along with PM. In the case of 10 species, all profiles in A's gut appear to be responding identically, while only PP is more unstable in A's saliva. B's gut, similar to the 7-cluster case, shows a more unstable PM and MM. This suggests that the trends that we see taking 4 species at a time are in general, not consistent, given that sampling larger sub-communities should not have a major effect on the system. It is possible that the trends generally tend to be true but non-robust, with even a small increase in cluster size leading to loss of structure as the average community becomes more random, but this likely indicates that there may be other additional factors alongside the causal profile that truly determine differential sensitivity.

In light of all of the above analyses, I find it difficult to conclude that profiles PP and PM have higher sensitivities in real microbiota- while this certainly provides an interesting avenue for future justification and validation, it is difficult to conclude this as a consistent result given the current state of data.

Appendix B

Additional parameter value plots

To ruminate means to consider to chew this life into something more digestible. A cow might spend half its living just ruminating.

Sam Sax

In this section, I include plots of many of the main results for other different values- the main text contains many plots for specific parameter choices to enhance interpretation. The remainder of the plots validating that these results are true across different settings can be seen here.

B.1 Families included in each microbiota

A's gut	A's saliva	B's gut
Actinomycetaceae	Actinomycetaceae	Actinomycetaceae
Alcaligenaceae	Aerococcaceae	Bacteroidaceae
Bacteroidaceae	Bacillaceae	Carnobacteriaceae

Bifidobacteriaceae	Bifidobacteriaceae	Catabacteriaceae
Campylobacteraceae	Burkholderiaceae	Clostridiaceae
Carnobacteriaceae	Campylobacteraceae	Clostridiales Family XIII*
Catabacteriaceae	Cardiobacteriaceae	Comamonadaceae
Clostridiaceae	Carnobacteriaceae	Coriobacteriaceae
Clostridiales Family XI*	Clostridiales Family XI*	Enterobacteriaceae
Clostridiales Family XIII*	Clostridiales Family XIII*	Erysipelotrichaceae
Comamonadaceae	Comamonadaceae	Lachnospiraceae
Coriobacteriaceae	Coriobacteriaceae	Leuconostocaceae
Desulfovibrionaceae	Corynebacteriaceae	Moraxellaceae
Enterobacteriaceae	Enterococcaceae	Rikenellaceae
Enterococcaceae	Erysipelotrichaceae	Ruminococcaceae
Erysipelotrichaceae	Eubacteriaceae	Streptococcaceae
Eubacteriaceae	F16	Veillonellaceae
Gemellaceae	Flavobacteriaceae	Verrucomicrobiaceae
Lachnospiraceae	Fusobacteriaceae	
Leuconostocaceae	Gemellaceae	
Moraxellaceae	Hydrogenophilaceae	
Pasteurellaceae	Intrasporangiaceae	
Porphyromonadaceae	Lachnospiraceae	
Prevotellaceae	Methylophilaceae	

Rikenellaceae	Micrococcaceae	
Ruminococcaceae	Neisseriaceae	
Staphylococcaceae	Oxalobacteraceae	
Streptococcaceae	Pasteurellaceae	
Turicibacteraceae	Peptococcaceae	
Veillonellaceae	Peptostreptococcaceae	
	Planococcaceae	
	Porphyromonadaceae	
	Prevotellaceae	
	Propionibacteriaceae	
	Spirochaetaceae	
	Staphylococcaceae	
	Streptococcaceae	
	Veillonellaceae	

*Incertae Sedis

B.2 Abundance and Mean Interaction Strength Variation

In Figure 4.7, I showed the VCR variation over time for some sample communities. Here, I briefly show the variation in abundance and mean interaction strength over time for these same communities to show further records of response diversity, and how complex and nonlinearly related these dynamics can be. Figure B.1 depicts these plots for abundance, and B.2 for mean interaction strength. This shows clear diversity in response across different

sub-communities, with some communities' parameter of interest increasing, decreasing or remaining unaffected (at least, at first glance) to the period of perturbation. While this is easy to imagine for abundance, this justifies why isolating drivers of this variation in the VCR is a complex issue.




Figure B.1: Community abundances. This is the same set of communities in Figure 4.7- this indicates that each of them have different abundances over time as well, changing in no characteristic manner, especially during the highlighted perturbation period.




Figure B.2: Community mean interaction strengths. Similar to the previous plot, for the same set of communities in Figure 4.7, this shows the diversity of responses in mean interaction strength with time.

B.3 Simulated Time-series Plots: Other Parameter Values

This section displays the similarity statistics of the simulated Lotka-Volterra communities in the other growth rate regimes. Figure B.3 shows the similarity statistic for the $\mu_r = 3$ and $\mu_r = 13$ cases. Similarity to the $\mu_r = 8$ case seen in Figure 4.2, we see an appreciable amount of recovery in the $\mu_A = 0$ case, and some cases of hindered recovery when μ_A is competitive. Results are clearly identical, showing that growth rate, while having an effect, does not determine as much as the interaction strength here. Similarly, I replicate these results by splitting these into causal profiles and applying the similarity statistic in Figure B.4. Here too, we see identical results to the $\mu_r = 8$ case- that is, in competitive settings, MM and MP type communities respond more strongly, while results in the weak-interaction regime,

no particular profile is dominant. It has already been shown in Table 4.1 that interaction strength changes do induce a sort of continuous transition of profile frequencies across growth rates- therefore, it is not unexpected that they produce similar results.




Figure B.3: Similarity statistics over time for simulated communities with mean growth rates of 3 (first row) and 13 (second row). Highlighted region indicates period of perturbation- note the conserved dip in similarity statistic upon onset of perturbation, and possible incomplete recovery in competitive communities, as expected and seen in the case of mean growth rate 8 in Figure 4.2

B.4 Distribution of Natural Causal Effects

Here, in Figure B.5, I provide a histogram of the values of natural direct and indirect effects of each community. This shows that most effects, while centered about 0, show quite a bit of spread. This spread appears to be higher in the case of abundance, while interaction effects are consistently slightly closer to 0. However, there is an appreciable range of values, both within a particular system as well as across systems. Interaction effect histograms look largely the same across all 3 microbiota, though- this can be refined with a G^2 -test, or further analysis of the types of interaction changes occurring in the system.




Figure B.4: Similarity statistics by profile for simulated communities with mean growth rates of 3 (first row) and 13 (second row). Highlighted region indicates period of perturbation. We see here that in competitive communities, MM and MP, as expected, respond more, while responses are mixed in weak interaction regimes, similar to the other case of $\mu_r = 8$ in Figure 4.6.




Figure B.5: Histograms of NIE and NDE. The NDE (top row) shows a wide distribution of values, while the NIE is more centered about 0. Regardless, this shows a range of responses in the community, as well as drawing out a commonality in NIE histograms.

References

- Abarbanel, H.D.I. and Kennel, M.B., 1993. Local False Nearest Neighbors and Dynamical Dimensions from Observed Chaotic Data. *Physical Review E* 47.5, 3057–3068. ISSN: 1063-651X, 1095-3787. DOI: [10.1103/PhysRevE.47.3057](https://doi.org/10.1103/PhysRevE.47.3057).
- Allesina, S. and Tang, S., 2012. Stability Criteria for Complex Ecosystems. *Nature* 483.7388, 205–208. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature10832](https://doi.org/10.1038/nature10832).
- Almaraz, P. and Oro, D., 2011. Size-Mediated Non-Trophic Interactions and Stochastic Predation Drive Assembly and Dynamics in a Seabird Community. *Ecology* 92.10, 1948–1958. ISSN: 0012-9658. DOI: [10.1890/11-0181.1](https://doi.org/10.1890/11-0181.1).
- Amor, D.R., Ratzke, C., and Gore, J., 2020. Transient Invaders Can Induce Shifts between Alternative Stable States of Microbial Communities. *Science Advances* 6.8, eaay8676. ISSN: 2375-2548. DOI: [10.1126/sciadv.aay8676](https://doi.org/10.1126/sciadv.aay8676).
- Andronov, A.A. and Pontryagin, L.S., 1937. Rough systems. In *Dokl. Akad. Nauk SSSR* 14.5, 247–250.
- Angulo, M.T., Moog, C.H., and Liu, Y.-Y., 2019. A Theoretical Framework for Controlling Complex Microbial Communities. *Nature Communications* 10.1, 1045. ISSN: 2041-1723. DOI: [10.1038/s41467-019-08890-y](https://doi.org/10.1038/s41467-019-08890-y).
- Ansari, A.F., Acharya, N.I.S., Kumaran, S., Ravindra, K., Reddy, Y.B.S., Dixit, N.M., and Raut, J., 2019. 110th Anniversary: High-Order Interactions Can Eclipse Pairwise Interactions in Shaping the Structure of Microbial Communities. *Industrial & Engineering Chemistry Research* 58.51, 23508–23518. ISSN: 0888-5885, 1520-5045. DOI: [10.1021/acs.iecr.9b03190](https://doi.org/10.1021/acs.iecr.9b03190).
- Arnold, V.I., 1986. *Catastrophe Theory* (Berlin, Heidelberg: Springer Berlin Heidelberg). ISBN: 978-3-540-16199-8 978-3-642-96937-9. DOI: [10.1007/978-3-642-96937-9](https://doi.org/10.1007/978-3-642-96937-9).
- Arnoldi, J.-F. and Haegeman, B., 2016. Unifying Dynamical and Structural Stability of Equilibria. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 472.2193, 20150874. ISSN: 1364-5021, 1471-2946. DOI: [10.1098/rspa.2015.0874](https://doi.org/10.1098/rspa.2015.0874).
- Arya, S., George, A.B., and O'Dwyer, J.P., 2023. Sparsity of Higher-Order Landscape Interactions Enables Learning and Prediction for Microbiomes. *Proceedings of the National Academy of Sciences* 120.48, e2307313120. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2307313120](https://doi.org/10.1073/pnas.2307313120).

- Bar-On, Y.M., Phillips, R., and Milo, R., 2018. The Biomass Distribution on Earth. *Proceedings of the National Academy of Sciences* 115.25, 6506–6511. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1711842115](https://doi.org/10.1073/pnas.1711842115).
- Bareinboim, E. and Pearl, J., 2012. Causal Inference by Surrogate Experiments: Z-Identifiability. DOI: [10.48550/ARXIV.1210.4842](https://doi.org/10.48550/ARXIV.1210.4842).
- Baruah, G., Ozgul, A., and Clements, C.F., 2022. Community Structure Determines the Predictability of Population Collapse. *Journal of Animal Ecology* 91.9, 1880–1891. ISSN: 0021-8790, 1365-2656. DOI: [10.1111/1365-2656.13769](https://doi.org/10.1111/1365-2656.13769).
- Bäumler, A.J. and Sperandio, V., 2016. Interactions between the Microbiota and Pathogenic Bacteria in the Gut. *Nature* 535.7610, 85–93. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature18849](https://doi.org/10.1038/nature18849).
- Belstrøm, D., 2020. The Salivary Microbiota in Health and Disease. *Journal of Oral Microbiology* 12.1, 1723975. ISSN: 2000-2297. DOI: [10.1080/20002297.2020.1723975](https://doi.org/10.1080/20002297.2020.1723975).
- Bender, E.A., Case, T.J., and Gilpin, M.E., 1984. Perturbation Experiments in Community Ecology: Theory and Practice. *Ecology* 65.1, 1–13. ISSN: 0012-9658, 1939-9170. DOI: [10.2307/1939452](https://doi.org/10.2307/1939452).
- Benincà, E., Ballantine, B., Ellner, S.P., and Huisman, J., 2015. Species Fluctuations Sustained by a Cyclic Succession at the Edge of Chaos. *Proceedings of the National Academy of Sciences* 112.20, 6389–6394. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1421968112](https://doi.org/10.1073/pnas.1421968112).
- Benincà, E., Huisman, J., Heerkloss, R., Jöhnk, K.D., Branco, P., Van Nes, E.H., Scheffer, M., and Ellner, S.P., 2008. Chaos in a Long-Term Experiment with a Plankton Community. *Nature* 451.7180, 822–825. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature06512](https://doi.org/10.1038/nature06512).
- Biggs, R., Carpenter, S.R., and Brock, W.A., 2009. Turning Back from the Brink: Detecting an Impending Regime Shift in Time to Avert It. *Proceedings of the National Academy of Sciences* 106.3, 826–831. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.0811729106](https://doi.org/10.1073/pnas.0811729106).
- Bjørnstad, O.N. and Grenfell, B.T., 2001. Noisy Clockwork: Time Series Analysis of Population Fluctuations in Animals. *Science* 293.5530, 638–643. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1062226](https://doi.org/10.1126/science.1062226).
- Blasche, S., Kim, Y., Mars, R.A.T., Machado, D., Maansson, M., Kafkia, E., Milanese, A., Zeller, G., Teusink, B., Nielsen, J., Benes, V., Neves, R., Sauer, U., and Patil, K.R., 2021. Metabolic Cooperation and Spatiotemporal Niche Partitioning in a Kefir Microbial Community. *Nature Microbiology* 6.2, 196–208. ISSN: 2058-5276. DOI: [10.1038/s41564-020-00816-5](https://doi.org/10.1038/s41564-020-00816-5).
- Boolchandani, M., Blake, K.S., Tilley, D.H., Cabada, M.M., Schwartz, D.J., Patel, S., Morales, M.L., Meza, R., Soto, G., Isidean, S.D., Porter, C.K., Simons, M.P., and Dantas, G., 2022. Impact of International Travel and Diarrhea on Gut Microbiome and Resistome Dynamics. *Nature Communications* 13.1, 7485. ISSN: 2041-1723. DOI: [10.1038/s41467-022-34862-w](https://doi.org/10.1038/s41467-022-34862-w).

- Bor, B., Bedree, J., Shi, W., McLean, J., and He, X., 2019. Saccharibacteria (TM7) in the Human Oral Microbiome. *Journal of Dental Research* 98.5, 500–509. ISSN: 0022-0345, 1544-0591. DOI: [10.1177/0022034519831671](https://doi.org/10.1177/0022034519831671).
- Brown, K., DeCoffe, D., Molcan, E., and Gibson, D.L., 2012. Diet-Induced Dysbiosis of the Intestinal Microbiota and the Effects on Immunity and Disease. *Nutrients* 4.8, 1095–1119. ISSN: 2072-6643. DOI: [10.3390/nu4081095](https://doi.org/10.3390/nu4081095).
- Bunin, G., 2017. Ecological communities with Lotka-Volterra dynamics. en. *Phys. Rev. E* 95.4.
- Camacho-Mateu, J., Lampo, A., Sireci, M., Muñoz, M.A., and Cuesta, J.A., 2024. Sparse Species Interactions Reproduce Abundance Correlation Patterns in Microbial Communities. *Proceedings of the National Academy of Sciences* 121.5, e2309575121. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2309575121](https://doi.org/10.1073/pnas.2309575121).
- Caporaso, J.G., Kuczynski, J., Stombaugh, J., Bittinger, K., Bushman, F.D., Costello, E.K., Fierer, N., Peña, A.G., Goodrich, J.K., Gordon, J.I., Huttley, G.A., Kelley, S.T., Knights, D., Koenig, J.E., Ley, R.E., Lozupone, C.A., McDonald, D., Muegge, B.D., Pirrung, M., Reeder, J., Sevinsky, J.R., Turnbaugh, P.J., Walters, W.A., Widmann, J., Yatsunenkov, T., Zaneveld, J., and Knight, R., 2010. QIIME Allows Analysis of High-Throughput Community Sequencing Data. *Nature Methods* 7.5, 335–336. ISSN: 1548-7091, 1548-7105. DOI: [10.1038/nmeth.f.303](https://doi.org/10.1038/nmeth.f.303).
- Caporaso, J.G., Lauber, C.L., Costello, E.K., Berg-Lyons, D., Gonzalez, A., Stombaugh, J., Knights, D., Gajer, P., Ravel, J., Fierer, N., Gordon, J.I., and Knight, R., 2011. Moving Pictures of the Human Microbiome. *Genome Biology* 12.5, R50. ISSN: 1465-6906. DOI: [10.1186/gb-2011-12-5-r50](https://doi.org/10.1186/gb-2011-12-5-r50).
- Carding, S., Verbeke, K., Vipond, D.T., Corfe, B.M., and Owen, L.J., 2015. Dysbiosis of the Gut Microbiota in Disease. *Microbial Ecology in Health & Disease* 26.0. ISSN: 1651-2235. DOI: [10.3402/mehd.v26.26191](https://doi.org/10.3402/mehd.v26.26191).
- Carpenter, S.R. and Brock, W.A., 2006. Rising Variance: A Leading Indicator of Ecological Transition. *Ecology Letters* 9.3, 311–318. ISSN: 1461-023X, 1461-0248. DOI: [10.1111/j.1461-0248.2005.00877.x](https://doi.org/10.1111/j.1461-0248.2005.00877.x).
- Cenci, S., Medeiros, L.P., Sugihara, G., and Saavedra, S., 2020. Assessing the predictability of nonlinear dynamics under smooth parameter changes. en. *J. R. Soc. Interface* 17.162, 20190627.
- Cenci, S. and Saavedra, S., 2018a. Structural Stability of Nonlinear Population Dynamics. *Physical Review E* 97.1, 012401. ISSN: 2470-0045, 2470-0053. DOI: [10.1103/PhysRevE.97.012401](https://doi.org/10.1103/PhysRevE.97.012401).
- 2018b. Uncertainty Quantification of the Effects of Biotic Interactions on Community Dynamics from Nonlinear Time-Series Data. *Journal of The Royal Society Interface* 15.147, 20180695. ISSN: 1742-5689, 1742-5662. DOI: [10.1098/rsif.2018.0695](https://doi.org/10.1098/rsif.2018.0695).
- 2019. Non-Parametric Estimation of the Structural Stability of Non-Equilibrium Community Dynamics. *Nature Ecology & Evolution* 3.6, 912–918. ISSN: 2397-334X. DOI: [10.1038/s41559-019-0879-1](https://doi.org/10.1038/s41559-019-0879-1).

- Cenci, S., Sugihara, G., and Saavedra, S., 2019. Regularized S-map for Inference and Forecasting with Noisy Ecological Time Series. *Methods in Ecology and Evolution* 10.5. Robert B. O'Hara, ed., 650–660. ISSN: 2041-210X, 2041-210X. DOI: [10.1111/2041-210X.13150](https://doi.org/10.1111/2041-210X.13150).
- Chang, C.-Y., Bajić, D., Vila, J.C.C., Estrela, S., and Sanchez, A., 2023. Emergent Coexistence in Multi-species Microbial Communities. *Science* 381.6655, 343–348. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.adg0727](https://doi.org/10.1126/science.adg0727).
- Chang, C.-W., Miki, T., Ushio, M., Ke, P.-J., Lu, H.-P., Shiah, F.-K., and Hsieh, C.-h., 2021. Reconstructing Large Interaction Networks from Empirical Time Series Data. *Ecology Letters* 24.12. Serguei Saavedra, ed., 2763–2774. ISSN: 1461-023X, 1461-0248. DOI: [10.1111/ele.13897](https://doi.org/10.1111/ele.13897).
- Chang, C.-W., Ushio, M., and Hsieh, C.-h., 2017. Empirical Dynamic Modeling for Beginners. *Ecological Research* 32.6, 785–796. ISSN: 0912-3814, 1440-1703. DOI: [10.1007/s11284-017-1469-9](https://doi.org/10.1007/s11284-017-1469-9).
- Chng, K.R., Ghosh, T.S., Tan, Y.H., Nandi, T., Lee, I.R., Ng, A.H.Q., Li, C., Ravikrishnan, A., Lim, K.M., Lye, D., Barkham, T., Raman, K., Chen, S.L., Chai, L., Young, B., Gan, Y.-H., and Nagarajan, N., 2020. Metagenome-Wide Association Analysis Identifies Microbial Determinants of Post-Antibiotic Ecological Recovery in the Gut. *Nature Ecology & Evolution* 4.9, 1256–1267. ISSN: 2397-334X. DOI: [10.1038/s41559-020-1236-0](https://doi.org/10.1038/s41559-020-1236-0).
- Clark, T.J. and Luis, A.D., 2019. Nonlinear Population Dynamics Are Ubiquitous in Animals. *Nature Ecology & Evolution* 4.1, 75–81. ISSN: 2397-334X. DOI: [10.1038/s41559-019-1052-6](https://doi.org/10.1038/s41559-019-1052-6).
- Conlon, M. and Bird, A., 2014. The Impact of Diet and Lifestyle on Gut Microbiota and Human Health. *Nutrients* 7.1, 17–44. ISSN: 2072-6643. DOI: [10.3390/nu7010017](https://doi.org/10.3390/nu7010017).
- Cosetta, C.M. and Wolfe, B.E., 2019. Causes and Consequences of Biotic Interactions within Microbiomes. *Current Opinion in Microbiology* 50, 35–41. ISSN: 13695274. DOI: [10.1016/j.mib.2019.09.004](https://doi.org/10.1016/j.mib.2019.09.004).
- Coyte, K.Z. and Rakoff-Nahoum, S., 2019. Understanding Competition and Cooperation within the Mammalian Gut Microbiome. *Current Biology* 29.11, R538–R544. ISSN: 09609822. DOI: [10.1016/j.cub.2019.04.017](https://doi.org/10.1016/j.cub.2019.04.017).
- Coyte, K.Z., Schluter, J., and Foster, K.R., 2015. The Ecology of the Microbiome: Networks, Competition, and Stability. *Science* 350.6261, 663–666. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.aad2602](https://doi.org/10.1126/science.aad2602).
- Cui, W., Marsland, R., and Mehta, P., 2021. Diverse Communities Behave like Typical Random Ecosystems. *Physical Review E* 104.3, 034416. ISSN: 2470-0045, 2470-0053. DOI: [10.1103/PhysRevE.104.034416](https://doi.org/10.1103/PhysRevE.104.034416).
- Culp, E.J. and Goodman, A.L., 2023. Cross-Feeding in the Gut Microbiome: Ecology and Mechanisms. *Cell Host & Microbe* 31.4, 485–499. ISSN: 19313128. DOI: [10.1016/j.chom.2023.03.016](https://doi.org/10.1016/j.chom.2023.03.016).
- D'Souza, A.W., Potter, R.F., Wallace, M., Shupe, A., Patel, S., Sun, X., Gul, D., Kwon, J.H., Andleeb, S., Burnham, C.-A.D., and Dantas, G., 2019. Spatiotemporal Dynamics of Multidrug Resistant Bacteria on Intensive Care Unit Surfaces. *Nature Communications* 10.1, 4569. ISSN: 2041-1723. DOI: [10.1038/s41467-019-12563-1](https://doi.org/10.1038/s41467-019-12563-1).

- Dai, L., Korolev, K.S., and Gore, J., 2015. Relation between Stability and Resilience Determines the Performance of Early Warning Signals under Different Environmental Drivers. *Proceedings of the National Academy of Sciences* 112.32, 10056–10061. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1418415112](https://doi.org/10.1073/pnas.1418415112).
- Dakos, V., Glaser, S.M., Hsieh, C.-h., and Sugihara, G., 2017. Elevated Nonlinearity as an Indicator of Shifts in the Dynamics of Populations under Stress. *Journal of The Royal Society Interface* 14.128, 20160845. ISSN: 1742-5689, 1742-5662. DOI: [10.1098/rsif.2016.0845](https://doi.org/10.1098/rsif.2016.0845).
- Dakos, V., Van Nes, E.H., D’Odorico, P., and Scheffer, M., 2012. Robustness of Variance and Autocorrelation as Indicators of Critical Slowing Down. *Ecology* 93.2, 264–271. ISSN: 0012-9658, 1939-9170. DOI: [10.1890/11-0889.1](https://doi.org/10.1890/11-0889.1).
- David, L.A., Materna, A.C., Friedman, J., Campos-Baptista, M.I., Blackburn, M.C., Perrotta, A., Erdman, S.E., and Alm, E.J., 2014a. Host Lifestyle Affects Human Microbiota on Daily Timescales. *Genome Biology* 15.7, R89. ISSN: 1465-6906. DOI: [10.1186/gb-2014-15-7-r89](https://doi.org/10.1186/gb-2014-15-7-r89).
- David, L.A., Maurice, C.F., Carmody, R.N., Gootenberg, D.B., Button, J.E., Wolfe, B.E., Ling, A.V., Devlin, A.S., Varma, Y., Fischbach, M.A., Biddinger, S.B., Dutton, R.J., and Turnbaugh, P.J., 2014b. Diet Rapidly and Reproducibly Alters the Human Gut Microbiome. *Nature* 505.7484, 559–563. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature12820](https://doi.org/10.1038/nature12820).
- De Filippo, C., Cavalieri, D., Di Paola, M., Ramazzotti, M., Poulet, J.B., Massart, S., Collini, S., Pieraccini, G., and Lionetti, P., 2010. Impact of Diet in Shaping Gut Microbiota Revealed by a Comparative Study in Children from Europe and Rural Africa. *Proceedings of the National Academy of Sciences* 107.33, 14691–14696. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1005963107](https://doi.org/10.1073/pnas.1005963107).
- De Muinck, E.J., Lundin, K.E.A., and Trosvik, P., 2017. Linking Spatial Structure and Community-Level Biotic Interactions through Cooccurrence and Time Series Modeling of the Human Intestinal Microbiota. *mSystems* 2.5. Morgan G. I. Langille, ed., e00086–17. ISSN: 2379-5077. DOI: [10.1128/mSystems.00086-17](https://doi.org/10.1128/mSystems.00086-17).
- Dedrick, S., Warrier, V., Lemon, K.P., and Momeni, B., 2023. When Does a Lotka-Volterra Model Represent Microbial Interactions? Insights from *in Vitro* Nasal Bacterial Communities. *mSystems* 8.3. Miguel Lurgi, ed., e00757–22. ISSN: 2379-5077. DOI: [10.1128/msystems.00757-22](https://doi.org/10.1128/msystems.00757-22).
- Dempsey, E. and Corr, S.C., 2022. Lactobacillus Spp. for Gastrointestinal Health: Current and Future Perspectives. *Frontiers in Immunology* 13, 840245. ISSN: 1664-3224. DOI: [10.3389/fimmu.2022.840245](https://doi.org/10.3389/fimmu.2022.840245).
- Deng, J., Angulo, M.T., and Saavedra, S., 2021. Generalizing Game-Changing Species across Microbial Communities. *ISME Communications* 1.1, 22. ISSN: 2730-6151. DOI: [10.1038/s43705-021-00022-2](https://doi.org/10.1038/s43705-021-00022-2).
- Desallais, M., Loreau, M., and Arnoldi, J.-F., 2024. *On the Robustness of Species Coexistence to Environmental Perturbations*. Preprint. Ecology. DOI: [10.1101/2024.01.21.575550](https://doi.org/10.1101/2024.01.21.575550).
- Descheemaeker, L. and De Buyl, S., 2020. Stochastic Logistic Models Reproduce Experimental Time Series of Microbial Communities. *eLife* 9, e55650. ISSN: 2050-084X. DOI: [10.7554/eLife.55650](https://doi.org/10.7554/eLife.55650).

- Dethlefsen, L., Huse, S., Sogin, M.L., and Relman, D.A., 2008. The Pervasive Effects of an Antibiotic on the Human Gut Microbiota, as Revealed by Deep 16S rRNA Sequencing. *PLoS Biology* 6.11. Jonathan A Eisen, ed., e280. ISSN: 1545-7885. DOI: [10.1371/journal.pbio.0060280](https://doi.org/10.1371/journal.pbio.0060280).
- Dey, S. and Joshi, A., 2013. Effects of Constant Immigration on the Dynamics and Persistence of Stable and Unstable Drosophila Populations. *Scientific Reports* 3.1, 1405. ISSN: 2045-2322. DOI: [10.1038/srep01405](https://doi.org/10.1038/srep01405).
- Deyle, E.R., Fogarty, M., Hsieh, C.-h., Kaufman, L., MacCall, A.D., Munch, S.B., Perretti, C.T., Ye, H., and Sugihara, G., 2013. Predicting Climate Effects on Pacific Sardine. *Proceedings of the National Academy of Sciences* 110.16, 6430–6435. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1215506110](https://doi.org/10.1073/pnas.1215506110).
- Deyle, E.R., May, R.M., Munch, S.B., and Sugihara, G., 2016. Tracking and Forecasting Ecosystem Interactions in Real Time. *Proceedings of the Royal Society B: Biological Sciences* 283.1822, 20152258. ISSN: 0962-8452, 1471-2954. DOI: [10.1098/rspb.2015.2258](https://doi.org/10.1098/rspb.2015.2258).
- Deyle, E.R. and Sugihara, G., 2011. Generalized Theorems for Nonlinear State Space Reconstruction. *PLoS ONE* 6.3. Matej Oresic, ed., e18295. ISSN: 1932-6203. DOI: [10.1371/journal.pone.0018295](https://doi.org/10.1371/journal.pone.0018295).
- Domínguez-García, V., Dakos, V., and Kéfi, S., 2019. Unveiling Dimensions of Stability in Complex Ecological Networks. *Proceedings of the National Academy of Sciences* 116.51, 25714–25720. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1904470116](https://doi.org/10.1073/pnas.1904470116).
- Eckburg, P.B., Bik, E.M., Bernstein, C.N., Purdom, E., Dethlefsen, L., Sargent, M., Gill, S.R., Nelson, K.E., and Relman, D.A., 2005. Diversity of the Human Intestinal Microbial Flora. *Science* 308.5728, 1635–1638. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1110591](https://doi.org/10.1126/science.1110591).
- Faith, J.J., Guruge, J.L., Charbonneau, M., Subramanian, S., Seedorf, H., Goodman, A.L., Clemente, J.C., Knight, R., Heath, A.C., Leibel, R.L., Rosenbaum, M., and Gordon, J.I., 2013. The Long-Term Stability of the Human Gut Microbiota. *Science* 341.6141, 1237439. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1237439](https://doi.org/10.1126/science.1237439).
- Falkowski, P.G., Fenchel, T., and Delong, E.F., 2008. The Microbial Engines That Drive Earth's Biogeochemical Cycles. *Science* 320.5879, 1034–1039. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1153213](https://doi.org/10.1126/science.1153213).
- Fant, L., Macocco, I., and Grilli, J., 2021. *Eco-Evolutionary Dynamics Lead to Functionally Robust and Redundant Communities*. Preprint. Ecology. DOI: [10.1101/2021.04.02.438173](https://doi.org/10.1101/2021.04.02.438173).
- Faust, K., Lahti, L., Gonze, D., De Vos, W.M., and Raes, J., 2015. Metagenomics Meets Time Series Analysis: Unraveling Microbial Community Dynamics. *Current Opinion in Microbiology* 25, 56–66. ISSN: 13695274. DOI: [10.1016/j.mib.2015.04.004](https://doi.org/10.1016/j.mib.2015.04.004).
- Figueiredo, A.R.T. and Kramer, J., 2020. Cooperation and Conflict Within the Microbiota and Their Effects On Animal Hosts. *Frontiers in Ecology and Evolution* 8, 132. ISSN: 2296-701X. DOI: [10.3389/fevo.2020.00132](https://doi.org/10.3389/fevo.2020.00132).

- Foster, K.R. and Bell, T., 2012. Competition, Not Cooperation, Dominates Interactions among Culturable Microbial Species. *Current Biology* 22.19, 1845–1850. ISSN: 09609822. DOI: [10.1016/j.cub.2012.08.005](https://doi.org/10.1016/j.cub.2012.08.005).
- Fujita, H., Ushio, M., Suzuki, K., Abe, M.S., Yamamichi, M., Iwayama, K., Canarini, A., Hayashi, I., Fukushima, K., Fukuda, S., Kiers, E.T., and Toju, H., 2023. Alternative Stable States, Nonlinear Behavior, and Predictability of Microbiome Dynamics. *Microbiome* 11.1, 63. ISSN: 2049-2618. DOI: [10.1186/s40168-023-01474-5](https://doi.org/10.1186/s40168-023-01474-5).
- Fukuyama, J., Rumker, L., Sankaran, K., Jeganathan, P., Dethlefsen, L., Relman, D.A., and Holmes, S.P., 2017. Multidomain Analyses of a Longitudinal Human Microbiome Intestinal Cleanout Perturbation Experiment. *PLOS Computational Biology* 13.8. Elhanan Borenstein, ed., e1005706. ISSN: 1553-7358. DOI: [10.1371/journal.pcbi.1005706](https://doi.org/10.1371/journal.pcbi.1005706).
- Gause, G.F., 1934. EXPERIMENTAL ANALYSIS OF VITO VOLTERRA'S MATHEMATICAL THEORY OF THE STRUGGLE FOR EXISTENCE. *Science* 79.2036, 16–17. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.79.2036.16-a](https://doi.org/10.1126/science.79.2036.16-a).
- Gee, K., Munch, S., and Rogers, T., 2023. *Quantifying Nonstationarity in Ecological Time Series*. Preprint. Preprints. DOI: [10.22541/au.168900665.57938650/v1](https://doi.org/10.22541/au.168900665.57938650/v1).
- George, A.B. and O'Dwyer, J., 2023. Universal Abundance Fluctuations across Microbial Communities, Tropical Forests, and Urban Populations. *Proceedings of the National Academy of Sciences* 120.44, e2215832120. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2215832120](https://doi.org/10.1073/pnas.2215832120).
- Ghosh, T.S., Shanahan, F., and O'Toole, P.W., 2022. The Gut Microbiome as a Modulator of Healthy Ageing. *Nature Reviews Gastroenterology & Hepatology* 19.9, 565–584. ISSN: 1759-5045, 1759-5053. DOI: [10.1038/s41575-022-00605-x](https://doi.org/10.1038/s41575-022-00605-x).
- Gibbs, T., Levin, S.A., and Levine, J.M., 2022. Coexistence in Diverse Communities with Higher-Order Interactions. *Proceedings of the National Academy of Sciences* 119.43, e2205063119. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2205063119](https://doi.org/10.1073/pnas.2205063119).
- Gibson, T.E., Kim, Y., Acharya, S., Kaplan, D.E., DiBenedetto, N., Lavin, R., Berger, B., Allegretti, J.R., Bry, L., and Gerber, G.K., 2021. *Microbial Dynamics Inference at Ecosystem-Scale*. Preprint. Systems Biology. DOI: [10.1101/2021.12.14.469105](https://doi.org/10.1101/2021.12.14.469105).
- Gilbert, J.A., Quinn, R.A., Debelius, J., Xu, Z.Z., Morton, J., Garg, N., Jansson, J.K., Dorrestein, P.C., and Knight, R., 2016. Microbiome-Wide Association Studies Link Dynamic Microbial Consortia to Disease. *Nature* 535.7610, 94–103. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature18850](https://doi.org/10.1038/nature18850).
- Glaser, S.M., Fogarty, M.J., Liu, H., Altman, I., Hsieh, C.-H., Kaufman, L., MacCall, A.D., Rosenberg, A.A., Ye, H., and Sugihara, G., 2014. Complex Dynamics May Limit Prediction in Marine Fisheries. *Fish and Fisheries* 15.4, 616–633. ISSN: 1467-2960, 1467-2979. DOI: [10.1111/faf.12037](https://doi.org/10.1111/faf.12037).
- Gleason, H.A., 1926. The Individualistic Concept of the Plant Association. *Bulletin of the Torrey Botanical Club* 53.1, 7. ISSN: 00409618. DOI: [10.2307/2479933](https://doi.org/10.2307/2479933). JSTOR: [2479933](https://www.jstor.org/stable/2479933).

- Goldford, J.E., Lu, N., Bajić, D., Estrela, S., Tikhonov, M., Sanchez-Gorostiaga, A., Segrè, D., Mehta, P., and Sanchez, A., 2018. Emergent Simplicity in Microbial Community Assembly. *Science* 361.6401, 469–474. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.aat1168](https://doi.org/10.1126/science.aat1168).
- Grainger, T.N., Letten, A.D., Gilbert, B., and Fukami, T., 2019. Applying Modern Coexistence Theory to Priority Effects. *Proceedings of the National Academy of Sciences* 116.13, 6205–6210. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1803122116](https://doi.org/10.1073/pnas.1803122116).
- Grilli, J., 2020. Macroecological Laws Describe Variation and Diversity in Microbial Communities. *Nature Communications* 11.1, 4743. ISSN: 2041-1723. DOI: [10.1038/s41467-020-18529-y](https://doi.org/10.1038/s41467-020-18529-y).
- Grimm, V. and Wissel, C., 1997. Babel, or the Ecological Stability Discussions: An Inventory and Analysis of Terminology and a Guide for Avoiding Confusion. *Oecologia* 109.3, 323–334. ISSN: 0029-8549, 1432-1939. DOI: [10.1007/s004420050090](https://doi.org/10.1007/s004420050090).
- Grobman, D.M., 1959. Homeomorphisms of systems of differential equations, *Doklady Acad. Nauk SSR* 128, 880–881.
- Grziwotz, F., Chang, C.-W., Dakos, V., Van Nes, E.H., Schwarzländer, M., Kamps, O., Heßler, M., Tokuda, I.T., Telschow, A., and Hsieh, C.-h., 2023. Anticipating the Occurrence and Type of Critical Transitions. *Science Advances* 9.1, eabq4558. ISSN: 2375-2548. DOI: [10.1126/sciadv.abq4558](https://doi.org/10.1126/sciadv.abq4558).
- Guarner, F. and Malagelada, J.-R., 2003. Gut Flora in Health and Disease. *The Lancet* 361.9356, 512–519. ISSN: 01406736. DOI: [10.1016/S0140-6736\(03\)12489-0](https://doi.org/10.1016/S0140-6736(03)12489-0).
- Guk, J., Bridier-Nahmias, A., Magnan, M., Grall, N., Duval, X., Clermont, O., Ruppé, E., d’Humières, C., Tenailon, O., Denamur, E., Mentré, F., Guedj, J., Burdet, C., and for the CEREMI study group, 2022. Modeling the Bacterial Dynamics in the Gut Microbiota Following an Antibiotic-induced Perturbation. *CPT: Pharmacometrics & Systems Pharmacology* 11.7, 906–918. ISSN: 2163-8306, 2163-8306. DOI: [10.1002/psp4.12806](https://doi.org/10.1002/psp4.12806).
- Gutman, Y., 2015. Takens’ Embedding Theorem with a Continuous Observable. DOI: [10.48550/ARXIV.1510.05843](https://doi.org/10.48550/ARXIV.1510.05843).
- Guttal, V. and Jayaprakash, C., 2008. Changing Skewness: An Early Warning Signal of Regime Shifts in Ecosystems. *Ecology Letters* 11.5, 450–460. ISSN: 1461-023X, 1461-0248. DOI: [10.1111/j.1461-0248.2008.01160.x](https://doi.org/10.1111/j.1461-0248.2008.01160.x).
- Hale, K.R.S., Valdovinos, F.S., and Martinez, N.D., 2020. Mutualism Increases Diversity, Stability, and Function of Multiplex Networks That Integrate Pollinators into Food Webs. *Nature Communications* 11.1, 2182. ISSN: 2041-1723. DOI: [10.1038/s41467-020-15688-w](https://doi.org/10.1038/s41467-020-15688-w).
- Hartman, P., 1960. On local homeomorphism of Euclidean spaces, *Bol. Bol. Sc. Mat. Mexicana* 2, 220–241.
- Hartstra, A.V., Bouter, K.E., Bäckhed, F., and Nieuwdorp, M., 2015. Insights Into the Role of the Microbiome in Obesity and Type 2 Diabetes. *Diabetes Care* 38.1, 159–165. ISSN: 0149-5992, 1935-5548. DOI: [10.2337/dc14-0769](https://doi.org/10.2337/dc14-0769).

- Hassell, M.P., Lawton, J.H., and May, R.M., 1976. Patterns of Dynamical Behaviour in Single-Species Populations. *The Journal of Animal Ecology* 45.2, 471. ISSN: 00218790. DOI: [10.2307/3886](https://doi.org/10.2307/3886). JSTOR: 3886.
- Hastie, T., Friedman, J., and Tibshirani, R., 2001. *The Elements of Statistical Learning* Springer Series in Statistics. (New York, NY: Springer New York). ISBN: 978-1-4899-0519-2 978-0-387-21606-5. DOI: [10.1007/978-0-387-21606-5](https://doi.org/10.1007/978-0-387-21606-5).
- Hastings, A., Hom, C.L., Ellner, S., Turchin, P., and Godfray, H.C.J., 1993. Chaos in Ecology: Is Mother Nature a Strange Attractor? *Annual Review of Ecology and Systematics* 24.1, 1–33. ISSN: 0066-4162. DOI: [10.1146/annurev.es.24.110193.000245](https://doi.org/10.1146/annurev.es.24.110193.000245).
- He, Y., Wu, W., Zheng, H.-M., Li, P., McDonald, D., Sheng, H.-F., Chen, M.-X., Chen, Z.-H., Ji, G.-Y., Zheng, Z.-D.-X., Mujagond, P., Chen, X.-J., Rong, Z.-H., Chen, P., Lyu, L.-Y., Wang, X., Wu, C.-B., Yu, N., Xu, Y.-J., Yin, J., Raes, J., Knight, R., Ma, W.-J., and Zhou, H.-W., 2018. Regional Variation Limits Applications of Healthy Gut Microbiome Reference Ranges and Disease Models. *Nature Medicine* 24.10, 1532–1535. ISSN: 1078-8956, 1546-170X. DOI: [10.1038/s41591-018-0164-x](https://doi.org/10.1038/s41591-018-0164-x).
- Hinrichsen, D. and Pritchard, A., 1986. Stability Radii of Linear Systems. *Systems & Control Letters* 7.1, 1–10. ISSN: 01676911. DOI: [10.1016/0167-6911\(86\)90094-0](https://doi.org/10.1016/0167-6911(86)90094-0).
- Ho, P.-Y., Good, B.H., and Huang, K.C., 2022. Competition for Fluctuating Resources Reproduces Statistics of Species Abundance over Time across Wide-Ranging Microbiotas. *eLife* 11, e75168. ISSN: 2050-084X. DOI: [10.7554/eLife.75168](https://doi.org/10.7554/eLife.75168).
- Holling, C S, 1973. Resilience and Stability of Ecological Systems. *Annual Review of Ecology and Systematics* 4.1, 1–23. ISSN: 0066-4162. DOI: [10.1146/annurev.es.04.110173.000245](https://doi.org/10.1146/annurev.es.04.110173.000245).
- 1959. Some Characteristics of Simple Types of Predation and Parasitism. *The Canadian Entomologist* 91.7, 385–398. ISSN: 0008-347X, 1918-3240. DOI: [10.4039/Ent91385-7](https://doi.org/10.4039/Ent91385-7).
- Hou, K., Wu, Z.-X., Chen, X.-Y., Wang, J.-Q., Zhang, D., Xiao, C., Zhu, D., Koya, J.B., Wei, L., Li, J., and Chen, Z.-S., 2022. Microbiota in Health and Diseases. *Signal Transduction and Targeted Therapy* 7.1, 135. ISSN: 2059-3635. DOI: [10.1038/s41392-022-00974-4](https://doi.org/10.1038/s41392-022-00974-4).
- Hsieh, C.-h., Anderson, C., and Sugihara, G., 2008. Extending Nonlinear Analysis to Short Ecological Time Series. *The American Naturalist* 171.1, 71–80. ISSN: 0003-0147, 1537-5323. DOI: [10.1086/524202](https://doi.org/10.1086/524202).
- Hsieh, C.-h., Glaser, S.M., Lucas, A.J., and Sugihara, G., 2005. Distinguishing Random Environmental Fluctuations from Ecological Catastrophes for the North Pacific Ocean. *Nature* 435.7040, 336–340. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature03553](https://doi.org/10.1038/nature03553).
- Huang, Y.-J., Chang, C.-W., and Hsieh, C.-H., 2024. Detecting shifts in nonlinear dynamics using Empirical Dynamic Modeling with Nested-Library Analysis. en. *PLoS Comput. Biol.* 20.1, e1011759.
- Ingrisch, J. and Bahn, M., 2018. Towards a Comparable Quantification of Resilience. *Trends in Ecology & Evolution* 33.4, 251–259. ISSN: 01695347. DOI: [10.1016/j.tree.2018.01.013](https://doi.org/10.1016/j.tree.2018.01.013).

- Ives, A.R. and Carpenter, S.R., 2007. Stability and Diversity of Ecosystems. *Science* 317.5834, 58–62. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1133258](https://doi.org/10.1126/science.1133258).
- Ives, A.R., Dennis, B., Cottingham, K.L., and Carpenter, S.R., 2003. ESTIMATING COMMUNITY STABILITY AND ECOLOGICAL INTERACTIONS FROM TIME-SERIES DATA. *Ecological Monographs* 73.2, 301–330. ISSN: 0012-9615. DOI: [10.1890/0012-9615\(2003\)073\[0301:ECSAEI\]2.0.CO;2](https://doi.org/10.1890/0012-9615(2003)073[0301:ECSAEI]2.0.CO;2).
- Janda, J.M. and Abbott, S.L., 2007. 16S rRNA Gene Sequencing for Bacterial Identification in the Diagnostic Laboratory: Pluses, Perils, and Pitfalls. *Journal of Clinical Microbiology* 45.9, 2761–2764. ISSN: 0095-1137, 1098-660X. DOI: [10.1128/JCM.01228-07](https://doi.org/10.1128/JCM.01228-07).
- Johnson, B. and Munch, S.B., 2022. An Empirical Dynamic Modeling Framework for Missing or Irregular Samples. *Ecological Modelling* 468, 109948. ISSN: 03043800. DOI: [10.1016/j.ecolmodel.2022.109948](https://doi.org/10.1016/j.ecolmodel.2022.109948).
- Jones, E.W., Carlson, J.M., Sivak, D.A., and Ludington, W.B., 2022. Stochastic Microbiome Assembly Depends on Context. *Proceedings of the National Academy of Sciences* 119.7, e2115877119. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2115877119](https://doi.org/10.1073/pnas.2115877119).
- Joseph, T.A., Shenhav, L., Xavier, J.B., Halperin, E., and Pe'er, I., 2020. Compositional Lotka-Volterra Describes Microbial Dynamics in the Simplex. *PLOS Computational Biology* 16.5. Vasilis Dakos, ed., e1007917. ISSN: 1553-7358. DOI: [10.1371/journal.pcbi.1007917](https://doi.org/10.1371/journal.pcbi.1007917).
- Ke, S., Hartmann, J., Ressler, K.J., Liu, Y.-Y., and Koenen, K.C., 2023. The Emerging Role of the Gut Microbiome in Posttraumatic Stress Disorder. *Brain, Behavior, and Immunity* 114, 360–370. ISSN: 08891591. DOI: [10.1016/j.bbi.2023.09.005](https://doi.org/10.1016/j.bbi.2023.09.005).
- Kehe, J., Ortiz, A., Kulesa, A., Gore, J., Blainey, P.C., and Friedman, J., 2021. Positive Interactions Are Common among Culturable Bacteria. *Science Advances* 7.45, eabi7159. ISSN: 2375-2548. DOI: [10.1126/sciadv.abi7159](https://doi.org/10.1126/sciadv.abi7159).
- Kennel, M.B., Brown, R., and Abarbanel, H.D.I., 1992. Determining Embedding Dimension for Phase-Space Reconstruction Using a Geometrical Construction. *Physical Review A* 45.6, 3403–3411. ISSN: 1050-2947, 1094-1622. DOI: [10.1103/PhysRevA.45.3403](https://doi.org/10.1103/PhysRevA.45.3403).
- Khazaei, T., Williams, R.L., Bogatyrev, S.R., Doyle, J.C., Henry, C.S., and Ismagilov, R.F., 2020. Metabolic Multistability and Hysteresis in a Model Aerobe-Anaerobe Microbiome Community. *Science Advances* 6.33, eaba0353. ISSN: 2375-2548. DOI: [10.1126/sciadv.aba0353](https://doi.org/10.1126/sciadv.aba0353).
- Kugiumtzis, D., Lingjærde, O., and Christophersen, N., 1998. Regularized Local Linear Prediction of Chaotic Time Series. *Physica D: Nonlinear Phenomena* 112.3-4, 344–360. ISSN: 01672789. DOI: [10.1016/S0167-2789\(97\)00171-1](https://doi.org/10.1016/S0167-2789(97)00171-1).
- Lagier, J.-C., Khelaifia, S., Alou, M.T., Ndongo, S., Dione, N., Hugon, P., Caputo, A., Cadoret, F., Traore, S.I., Seck, E.H., Dubourg, G., Durand, G., Mourembou, G., Guilhot, E., Togo, A., Bellali, S., Bachar, D., Cas-sir, N., Bittar, F., Delerce, J., Mailhe, M., Ricaboni, D., Bilen, M., Dangui Nieko, N.P.M., Dia Badiane, N.M., Valles, C., Mouelhi, D., Diop, K., Million, M., Musso, D., Abrahão, J., Azhar, E.I., Bibi, F., Yasir,

- M., Diallo, A., Sokhna, C., Djossou, F., Vitton, V., Robert, C., Rolain, J.M., La Scola, B., Fournier, P.-E., Levasseur, A., and Raoult, D., 2016. Culture of Previously Uncultured Members of the Human Gut Microbiota by Culturomics. *Nature Microbiology* 1.12, 16203. ISSN: 2058-5276. DOI: [10.1038/nmicrobiol.2016.203](https://doi.org/10.1038/nmicrobiol.2016.203).
- Laska, M.S. and Wootton, J.T., 1998. THEORETICAL CONCEPTS AND EMPIRICAL APPROACHES TO MEASURING INTERACTION STRENGTH. *Ecology* 79.2, 461–476. ISSN: 0012-9658. DOI: [10.1890/0012-9658\(1998\)079\[0461:TCAEAT\]2.0.CO;2](https://doi.org/10.1890/0012-9658(1998)079[0461:TCAEAT]2.0.CO;2).
- Levin, S.A., 1970. Community Equilibria and Stability, and an Extension of the Competitive Exclusion Principle. *The American Naturalist* 104.939, 413–423. ISSN: 0003-0147, 1537-5323. DOI: [10.1086/282676](https://doi.org/10.1086/282676).
- Levins, R., 1969. Some Demographic and Genetic Consequences of Environmental Heterogeneity for Biological Control. *Bulletin of the Entomological Society of America* 15.3, 237–240. ISSN: 0013-8754. DOI: [10.1093/besa/15.3.237](https://doi.org/10.1093/besa/15.3.237).
- Levy, R., Magis, A.T., Earls, J.C., Manor, O., Wilmanski, T., Lovejoy, J., Gibbons, S.M., Omenn, G.S., Hood, L., and Price, N.D., 2020. Longitudinal Analysis Reveals Transition Barriers between Dominant Ecological States in the Gut Microbiome. *Proceedings of the National Academy of Sciences* 117.24, 13839–13845. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1922498117](https://doi.org/10.1073/pnas.1922498117).
- Lewontin, R.C., 1969. The Meaning of Stability. *Brookhaven Symposia in Biology* 22, 13–24. ISSN: 0068-2799.
- Liu, O.R. and Gaines, S.D., 2022. Environmental Context Dependency in Species Interactions. *Proceedings of the National Academy of Sciences* 119.36, e2118539119. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2118539119](https://doi.org/10.1073/pnas.2118539119).
- Liu, Y.-Y., 2023. Controlling the Human Microbiome. *Cell Systems* 14.2, 135–159. ISSN: 24054712. DOI: [10.1016/j.cels.2022.12.010](https://doi.org/10.1016/j.cels.2022.12.010).
- Locey, K.J. and Lennon, J.T., 2016. Scaling Laws Predict Global Microbial Diversity. *Proceedings of the National Academy of Sciences* 113.21, 5970–5975. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1521291113](https://doi.org/10.1073/pnas.1521291113).
- Logofet, D.O., 1993. *Matrices and Graphs: Stability Problems in Mathematical Ecology* (Boca Raton: CRC press). ISBN: 978-0-8493-4246-2.
- Long, C., Deng, J., Nguyen, J., Liu, Y.-Y., Alm, E.J., Solé, R., and Saavedra, S., 2024. Structured Community Transitions Explain the Switching Capacity of Microbial Systems. *Proceedings of the National Academy of Sciences* 121.6, e2312521121. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2312521121](https://doi.org/10.1073/pnas.2312521121).
- Lorenz, E.N., 1963. Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences* 20.2, 130–141. ISSN: 0022-4928, 1520-0469. DOI: [10.1175/1520-0469\(1963\)020<0130:DNF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2).
- Lotka, A.J., 1925. *Elements of physical biology* (Williams & Wilkins).

- Lozupone, C.A., Stombaugh, J.I., Gordon, J.I., Jansson, J.K., and Knight, R., 2012. Diversity, Stability and Resilience of the Human Gut Microbiota. *Nature* 489.7415, 220–230. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature11550](https://doi.org/10.1038/nature11550).
- Ludington, W.B., 2022. Higher-Order Microbiome Interactions and How to Find Them. *Trends in Microbiology* 30.7, 618–621. ISSN: 0966842X. DOI: [10.1016/j.tim.2022.03.011](https://doi.org/10.1016/j.tim.2022.03.011).
- MacArthur, R., 1955. Fluctuations of Animal Populations and a Measure of Community Stability. *Ecology* 36.3, 533–536. ISSN: 0012-9658, 1939-9170. DOI: [10.2307/1929601](https://doi.org/10.2307/1929601).
- 1970. Species Packing and Competitive Equilibrium for Many Species. *Theoretical Population Biology* 1.1, 1–11. ISSN: 00405809. DOI: [10.1016/0040-5809\(70\)90039-0](https://doi.org/10.1016/0040-5809(70)90039-0).
- Machado, D., Maistrenko, O.M., Andrejev, S., Kim, Y., Bork, P., Patil, K.R., and Patil, K.R., 2021. Polarization of Microbial Communities between Competitive and Cooperative Metabolism. *Nature Ecology & Evolution* 5.2, 195–203. ISSN: 2397-334X. DOI: [10.1038/s41559-020-01353-4](https://doi.org/10.1038/s41559-020-01353-4).
- Maki, K.A., Kazmi, N., Barb, J.J., and Ames, N., 2021. The Oral and Gut Bacterial Microbiomes: Similarities, Differences, and Connections. *Biological Research For Nursing* 23.1, 7–20. ISSN: 1099-8004, 1552-4175. DOI: [10.1177/1099800420941606](https://doi.org/10.1177/1099800420941606).
- Mangal, U., Noh, K., Lee, S., Cha, J.-K., Song, J.S., Cha, J.-Y., Lee, K.-J., Kim, K.-M., Kwon, J.-S., and Choi, S.-H., 2023. Multistability and Hysteresis in States of Oral Microbiota: Is It Impacting the Dental Clinical Cohort Studies? *Journal of Periodontal Research* 58.2, 381–391. ISSN: 0022-3484, 1600-0765. DOI: [10.1111/jre.13098](https://doi.org/10.1111/jre.13098).
- Manor, O., Dai, C.L., Kornilov, S.A., Smith, B., Price, N.D., Lovejoy, J.C., Gibbons, S.M., and Magis, A.T., 2020. Health and Disease Markers Correlate with Gut Microbiome Composition across Thousands of People. *Nature Communications* 11.1, 5206. ISSN: 2041-1723. DOI: [10.1038/s41467-020-18871-1](https://doi.org/10.1038/s41467-020-18871-1).
- Marchesi, J.R., Adams, D.H., Fava, F., Hermes, G.D.A., Hirschfield, G.M., Hold, G., Quraishi, M.N., Kinross, J., Smidt, H., Tuohy, K.M., Thomas, L.V., Zoetendal, E.G., and Hart, A., 2016. The Gut Microbiota and Host Health: A New Clinical Frontier. *Gut* 65.2, 330–339. ISSN: 0017-5749, 1468-3288. DOI: [10.1136/gutjnl-2015-309990](https://doi.org/10.1136/gutjnl-2015-309990).
- Mark Welch, J.L., Ramírez-Puebla, S.T., and Borisy, G.G., 2020. Oral Microbiome Geography: Micron-Scale Habitat and Niche. *Cell Host & Microbe* 28.2, 160–168. ISSN: 19313128. DOI: [10.1016/j.chom.2020.07.009](https://doi.org/10.1016/j.chom.2020.07.009).
- Markus Kalisch, Martin Mächler, Diego Colombo, Marloes H. Maathuis, and Peter Bühlmann, 2012. Causal Inference Using Graphical Models with the R Package pcalg. *Journal of Statistical Software* 47.11, 1–26. DOI: [10.18637/jss.v047.i11](https://doi.org/10.18637/jss.v047.i11).
- May, R.M., 1972. Will a Large Complex System Be Stable? *Nature* 238.5364, 413–414. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/238413a0](https://doi.org/10.1038/238413a0).

- 1976. Simple Mathematical Models with Very Complicated Dynamics. *Nature* 261.5560, 459–467. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/261459a0](https://doi.org/10.1038/261459a0).
- May, R.M. and Oster, G.F., 1976. Bifurcations and Dynamic Complexity in Simple Ecological Models. *The American Naturalist* 110.974, 573–599. ISSN: 0003-0147, 1537-5323. DOI: [10.1086/283092](https://doi.org/10.1086/283092).
- McBurney, M.I., Davis, C., Fraser, C.M., Schneeman, B.O., Huttenhower, C., Verbeke, K., Walter, J., and Latulippe, M.E., 2019. Establishing What Constitutes a Healthy Human Gut Microbiome: State of the Science, Regulatory Considerations, and Future Directions. *The Journal of Nutrition* 149.11, 1882–1895. ISSN: 00223166. DOI: [10.1093/jn/nxz154](https://doi.org/10.1093/jn/nxz154).
- McCann, K.S., 2000. The Diversity–Stability Debate. *Nature* 405.6783, 228–233. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/35012234](https://doi.org/10.1038/35012234).
- Medeiros, L.P., Allesina, S., Dakos, V., Sugihara, G., and Saavedra, S., 2023. Ranking Species Based on Sensitivity to Perturbations under Non-equilibrium Community Dynamics. *Ecology Letters* 26.1, 170–183. ISSN: 1461-023X, 1461-0248. DOI: [10.1111/ele.14131](https://doi.org/10.1111/ele.14131).
- Medeiros, L.P. and Saavedra, S., 2023. Understanding the State-dependent Impact of Species Correlated Responses on Community Sensitivity to Perturbations. *Ecology* 104.8, e4115. ISSN: 0012-9658, 1939-9170. DOI: [10.1002/ecy.4115](https://doi.org/10.1002/ecy.4115).
- Medeiros, L.P., Song, C., and Saavedra, S., 2021. Merging Dynamical and Structural Indicators to Measure Resilience in Multispecies Systems. *Journal of Animal Ecology* 90.9, 2027–2040. ISSN: 0021-8790, 1365-2656. DOI: [10.1111/1365-2656.13421](https://doi.org/10.1111/1365-2656.13421).
- Merz, E., Saberski, E., Gilarranz, L.J., Isles, P.D.F., Sugihara, G., Berger, C., and Pomati, F., 2023. Disruption of Ecological Networks in Lakes by Climate Change and Nutrient Fluctuations. *Nature Climate Change* 13.4, 389–396. ISSN: 1758-678X, 1758-6798. DOI: [10.1038/s41558-023-01615-6](https://doi.org/10.1038/s41558-023-01615-6).
- MetaHIT Consortium (additional members), Arumugam, M., Raes, J., Pelletier, E., Le Paslier, D., Yamada, T., Mende, D.R., Fernandes, G.R., Tap, J., Bruls, T., Batto, J.-M., Bertalan, M., Borruel, N., Casellas, F., Fernandez, L., Gautier, L., Hansen, T., Hattori, M., Hayashi, T., Kleerebezem, M., Kurokawa, K., Leclerc, M., Levenez, F., Manichanh, C., Nielsen, H.B., Nielsen, T., Pons, N., Poulain, J., Qin, J., Sicheritz-Ponten, T., Tims, S., Torrents, D., Ugarte, E., Zoetendal, E.G., Wang, J., Guarner, F., Pedersen, O., De Vos, W.M., Brunak, S., Doré, J., Weissenbach, J., Ehrlich, S.D., and Bork, P., 2011. Enterotypes of the Human Gut Microbiome. *Nature* 473.7346, 174–180. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature09944](https://doi.org/10.1038/nature09944).
- Momeni, B., Xie, L., and Shou, W., 2017. Lotka-Volterra Pairwise Modeling Fails to Capture Diverse Pairwise Microbial Interactions. *eLife* 6, e25051. ISSN: 2050-084X. DOI: [10.7554/eLife.25051](https://doi.org/10.7554/eLife.25051).
- Morin, M.A., Morrison, A.J., Harms, M.J., and Dutton, R.J., 2022. Higher-Order Interactions Shape Microbial Interactions as Microbial Community Complexity Increases. *Scientific Reports* 12.1, 22640. ISSN: 2045-2322. DOI: [10.1038/s41598-022-25303-1](https://doi.org/10.1038/s41598-022-25303-1).

- Munch, S.B., Brias, A., Sugihara, G., and Rogers, T.L., 2020. Frequently Asked Questions about Nonlinear Dynamics and Empirical Dynamic Modelling. *ICES Journal of Marine Science* 77.4. Gary Griffith, ed., 1463–1479. ISSN: 1095-9289. DOI: [10.1093/icesjms/fsz209](https://doi.org/10.1093/icesjms/fsz209).
- Munch, S.B., Poynor, V., and Arriaza, J.L., 2017. Circumventing Structural Uncertainty: A Bayesian Perspective on Nonlinear Forecasting for Ecology. *Ecological Complexity* 32, 134–143. ISSN: 1476945X. DOI: [10.1016/j.ecocom.2016.08.006](https://doi.org/10.1016/j.ecocom.2016.08.006).
- Munch, S.B., Rogers, T.L., Symons, C.C., Anderson, D., and Pennekamp, F., 2023. Constraining Nonlinear Time Series Modeling with the Metabolic Theory of Ecology. *Proceedings of the National Academy of Sciences* 120.12, e2211758120. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2211758120](https://doi.org/10.1073/pnas.2211758120).
- Murdoch, W.W., Avery, S., and Smyth, M.E.B., 1975. Switching in Predatory Fish. *Ecology* 56.5, 1094–1105. ISSN: 00129658. DOI: [10.2307/1936149](https://doi.org/10.2307/1936149).
- Nagpal, R., Mainali, R., Ahmadi, S., Wang, S., Singh, R., Kavanagh, K., Kitzman, D.W., Kushugulova, A., Marotta, F., and Yadav, H., 2018. Gut Microbiome and Aging: Physiological and Mechanistic Insights. *Nutrition and Healthy Aging* 4.4, 267–285. ISSN: 24519480, 24519502. DOI: [10.3233/NHA-170030](https://doi.org/10.3233/NHA-170030).
- Navas-Molina, J.A., Peralta-Sánchez, J.M., González, A., McMurdie, P.J., Vázquez-Baeza, Y., Xu, Z., Ursell, L.K., Lauber, C., Zhou, H., Song, S.J., Huntley, J., Ackermann, G.L., Berg-Lyons, D., Holmes, S., Caporaso, J.G., and Knight, R., 2013. “Advancing Our Understanding of the Human Microbiome Using QIIME.” *Methods in Enzymology*, vol. 531. (Elsevier), pp. 371–444. ISBN: 978-0-12-407863-5. DOI: [10.1016/B978-0-12-407863-5.00019-8](https://doi.org/10.1016/B978-0-12-407863-5.00019-8).
- Nguyen, J., Lara-Gutiérrez, J., and Stocker, R., 2021a. Environmental Fluctuations and Their Effects on Microbial Communities, Populations and Individuals. *FEMS Microbiology Reviews* 45.4, fuaa068. ISSN: 1574-6976. DOI: [10.1093/femsre/fuaa068](https://doi.org/10.1093/femsre/fuaa068).
- Nguyen, J., Pepin, D.M., and Tropini, C., 2021b. Cause or Effect? The Spatial Organization of Pathogens and the Gut Microbiota in Disease. *Microbes and Infection* 23.6-7, 104815. ISSN: 12864579. DOI: [10.1016/j.micinf.2021.104815](https://doi.org/10.1016/j.micinf.2021.104815).
- Nychka, D., Ellner, S., Gallant, A.R., and McCaffrey, D., 1992. Finding Chaos in Noisy Systems. *Journal of the Royal Statistical Society: Series B (Methodological)* 54.2, 399–426. ISSN: 0035-9246, 2517-6161. DOI: [10.1111/j.2517-6161.1992.tb01889.x](https://doi.org/10.1111/j.2517-6161.1992.tb01889.x).
- O’Brien, D.A., Deb, S., Gal, G., Thackeray, S.J., Dutta, P.S., Matsuzaki, S.-i.S., May, L., and Clements, C.F., 2023a. Early Warning Signals Have Limited Applicability to Empirical Lake Data. *Nature Communications* 14.1, 7942. ISSN: 2041-1723. DOI: [10.1038/s41467-023-43744-8](https://doi.org/10.1038/s41467-023-43744-8).
- O’Brien, D.A., Deb, S., Sidheekh, S., Krishnan, N.C., Sharathi Dutta, P., and Clements, C.F., 2023b. EWS-methods: An R Package to Forecast Tipping Points at the Community Level Using Early Warning Signals, Resilience Measures, and Machine Learning Models. *Ecography* 2023.10, e06674. ISSN: 0906-7590, 1600-0587. DOI: [10.1111/ecog.06674](https://doi.org/10.1111/ecog.06674).

- O’Callaghan, A. and Van Sinderen, D., 2016. Bifidobacteria and Their Role as Members of the Human Gut Microbiota. *Frontiers in Microbiology* 7. ISSN: 1664-302X. DOI: [10.3389/fmicb.2016.00925](https://doi.org/10.3389/fmicb.2016.00925).
- Odum, E.P., 1953. *Fundamentals of ecology (1953)* (Yale University Press), pp. 233–244.
- Packard, N.H., Crutchfield, J.P., Farmer, J.D., and Shaw, R.S., 1980. Geometry from a Time Series. *Physical Review Letters* 45.9, 712–716. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.45.712](https://doi.org/10.1103/PhysRevLett.45.712).
- Palmer, J.D. and Foster, K.R., 2022. Bacterial Species Rarely Work Together. *Science* 376.6593, 581–582. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.abn5093](https://doi.org/10.1126/science.abn5093).
- Panke-Buisse, K., Poole, A.C., Goodrich, J.K., Ley, R.E., and Kao-Kniffin, J., 2015. Selection on Soil Microbiomes Reveals Reproducible Impacts on Plant Function. *The ISME Journal* 9.4, 980–989. ISSN: 1751-7362, 1751-7370. DOI: [10.1038/ismej.2014.196](https://doi.org/10.1038/ismej.2014.196).
- Park, J., Smith, C., Sugihara, G., and Deyle, E., 2023. *rEDM: Empirical Dynamic Modeling ('EDM')*. ().
- Pearl, J., 2009. *Causality: Models, Reasoning, and Inference* (Cambridge, U.K. ; New York: Cambridge University Press). ISBN: 978-0-521-89560-6 978-0-521-77362-1.
- Peixoto, M., 1962. Structural Stability on Two-Dimensional Manifolds. *Topology* 1.2, 101–120. ISSN: 00409383. DOI: [10.1016/0040-9383\(65\)90018-2](https://doi.org/10.1016/0040-9383(65)90018-2).
- Pennekamp, F., Pontarp, M., Tabi, A., Altermatt, F., Alther, R., Choffat, Y., Fronhofer, E.A., Ganesanandamoorthy, P., Garnier, A., Griffiths, J.I., Greene, S., Horgan, K., Massie, T.M., Mächler, E., Palamara, G.M., Seymour, M., and Petchey, O.L., 2018. Biodiversity Increases and Decreases Ecosystem Stability. *Nature* 563.7729, 109–112. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/s41586-018-0627-8](https://doi.org/10.1038/s41586-018-0627-8).
- Pérez-Cobas, A.E., Artacho, A., Ott, S.J., Moya, A., Gosalbes, M.-a.J., and Latorre, A., 2014. Structural and Functional Changes in the Gut Microbiota Associated to Clostridium Difficile Infection. *Frontiers in Microbiology* 5. ISSN: 1664-302X. DOI: [10.3389/fmicb.2014.00335](https://doi.org/10.3389/fmicb.2014.00335).
- Perfect, E., Kay, B.D., Van Loon, W.K.P., Sheard, R.W., and Pojasok, T., 1990. Factors Influencing Soil Structural Stability within a Growing Season. *Soil Science Society of America Journal* 54.1, 173–179. ISSN: 0361-5995, 1435-0661. DOI: [10.2136/sssaj1990.03615995005400010027x](https://doi.org/10.2136/sssaj1990.03615995005400010027x).
- Perretti, C.T., Sugihara, G., and Munch, S.B., 2013. Nonparametric Forecasting Outperforms Parametric Methods for a Simulated Multispecies System. *Ecology* 94.4, 794–800. ISSN: 0012-9658, 1939-9170. DOI: [10.1890/12-0904.1](https://doi.org/10.1890/12-0904.1).
- Pimm, S.L., 1984. The Complexity and Stability of Ecosystems. *Nature* 307.5949, 321–326. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/307321a0](https://doi.org/10.1038/307321a0).
- Preston, F.W., 1948. The Commonness, And Rarity, of Species. *Ecology* 29.3, 254–283. ISSN: 0012-9658, 1939-9170. DOI: [10.2307/1930989](https://doi.org/10.2307/1930989).

- Qian, J.J. and Akçay, E., 2020. The Balance of Interaction Types Determines the Assembly and Stability of Ecological Communities. *Nature Ecology & Evolution* 4.3, 356–365. ISSN: 2397-334X. DOI: [10.1038/s41559-020-1121-x](https://doi.org/10.1038/s41559-020-1121-x).
- R Core Team, 2022. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. (Vienna, Austria).
- Rashidi, A., Ebadi, M., Rehman, T.U., Elhusseini, H., Nalluri, H., Kaiser, T., Holtan, S.G., Khoruts, A., Weisdorf, D.J., and Staley, C., 2021. Gut Microbiota Response to Antibiotics Is Personalized and Depends on Baseline Microbiota. *Microbiome* 9.1, 211. ISSN: 2049-2618. DOI: [10.1186/s40168-021-01170-2](https://doi.org/10.1186/s40168-021-01170-2).
- Ratzke, C., Barrere, J., and Gore, J., 2020. Strength of Species Interactions Determines Biodiversity and Stability in Microbial Communities. *Nature Ecology & Evolution* 4.3, 376–383. ISSN: 2397-334X. DOI: [10.1038/s41559-020-1099-4](https://doi.org/10.1038/s41559-020-1099-4).
- Raymond, F., Ouameur, A.A., Déraspe, M., Iqbal, N., Gingras, H., Dridi, B., Leprohon, P., Plante, P.-L., Giroux, R., Bérubé, È., Frenette, J., Boudreau, D.K., Simard, J.-L., Chabot, I., Domingo, M.-C., Trottier, S., Boissinot, M., Huletsky, A., Roy, P.H., Ouellette, M., Bergeron, M.G., and Corbeil, J., 2016. The Initial State of the Human Gut Microbiome Determines Its Reshaping by Antibiotics. *The ISME Journal* 10.3, 707–720. ISSN: 1751-7362, 1751-7370. DOI: [10.1038/ismej.2015.148](https://doi.org/10.1038/ismej.2015.148).
- Recknagel, F., 1985. Analysis of Structural Stability of Aquatic Ecosystems as an Aid for Ecosystem Control. *Ecological Modelling* 27.3-4, 221–234. ISSN: 03043800. DOI: [10.1016/0304-3800\(85\)90004-3](https://doi.org/10.1016/0304-3800(85)90004-3).
- Reichenbach, H., 1956. *The Direction of Time* Dover Books on Physics. (Mineola, N.Y: Dover). ISBN: 978-0-486-40926-9.
- Ricker, W.E., 1954. Stock and Recruitment. *Journal of the Fisheries Research Board of Canada* 11.5, 559–623. ISSN: 0015-296X. DOI: [10.1139/f54-039](https://doi.org/10.1139/f54-039).
- Rogers, T.L., Johnson, B.J., and Munch, S.B., 2022. Chaos Is Not Rare in Natural Ecosystems. *Nature Ecology & Evolution* 6.8, 1105–1111. ISSN: 2397-334X. DOI: [10.1038/s41559-022-01787-y](https://doi.org/10.1038/s41559-022-01787-y).
- Rohr, R.P., Saavedra, S., and Bascompte, J., 2014. On the Structural Stability of Mutualistic Systems. *Science* 345.6195, 1253497. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1253497](https://doi.org/10.1126/science.1253497).
- Ross, S.M., 2017. *Introductory Statistics* Fourth edition. (Amsterdam ; Boston: Elsevier/AP). ISBN: 978-0-12-804317-2.
- Saavedra, S., Bartomeus, I., Godoy, O., Rohr, R.P., and Zu, P., 2022. Towards a system-level causative knowledge of pollinator communities. en. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 377.1853, 20210159.
- Saavedra, S., Medeiros, L.P., and AlAdwani, M., 2020. Structural Forecasting of Species Persistence under Changing Environments. *Ecology Letters* 23.10. Carl Boettiger, ed., 1511–1521. ISSN: 1461-023X, 1461-0248. DOI: [10.1111/ele.13582](https://doi.org/10.1111/ele.13582).

- Saavedra, S., Rohr, R.P., Bascompte, J., Godoy, O., Kraft, N.J.B., and Levine, J.M., 2017. A Structural Approach for Understanding Multispecies Coexistence. *Ecological Monographs* 87.3, 470–486. ISSN: 0012-9615, 1557-7015. DOI: [10.1002/ecm.1263](https://doi.org/10.1002/ecm.1263).
- Saavedra, S., Rohr, R.P., Olesen, J.M., and Bascompte, J., 2016. Nested Species Interactions Promote Feasibility over Stability during the Assembly of a Pollinator Community. *Ecology and Evolution* 6.4, 997–1007. ISSN: 2045-7758, 2045-7758. DOI: [10.1002/ece3.1930](https://doi.org/10.1002/ece3.1930).
- Säterberg, T., Jonsson, T., Yearsley, J., Berg, S., and Ebenman, B., 2019. A Potential Role for Rare Species in Ecosystem Dynamics. *Scientific Reports* 9.1, 11107. ISSN: 2045-2322. DOI: [10.1038/s41598-019-47541-6](https://doi.org/10.1038/s41598-019-47541-6).
- Säterberg, T. and McCann, K.S., 2021. Detecting Alternative Attractors in Ecosystem Dynamics. *Communications Biology* 4.1, 975. ISSN: 2399-3642. DOI: [10.1038/s42003-021-02471-w](https://doi.org/10.1038/s42003-021-02471-w).
- Sauer, T., Yorke, J.A., and Casdagli, M., 1991. Embedology. *Journal of Statistical Physics* 65.3-4, 579–616. ISSN: 0022-4715, 1572-9613. DOI: [10.1007/BF01053745](https://doi.org/10.1007/BF01053745).
- Scheffer, M., Bascompte, J., Brock, W.A., Brovkin, V., Carpenter, S.R., Dakos, V., Held, H., Van Nes, E.H., Rietkerk, M., and Sugihara, G., 2009. Early-Warning Signals for Critical Transitions. *Nature* 461.7260, 53–59. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature08227](https://doi.org/10.1038/nature08227).
- Scheffer, M., Carpenter, S., Foley, J.A., Folke, C., and Walker, B., 2001. Catastrophic Shifts in Ecosystems. *Nature* 413.6856, 591–596. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/35098000](https://doi.org/10.1038/35098000).
- Schoener, T.W., 1993. “On the Relative Importance of Direct versus Indirect Effects in Ecological Communities.” *Mutualism and Community Organization*, Hiroya Kawanabe, Joel E Cohen, and Keiji Iwasaki, ed. (Oxford University Press/Oxford), pp. 365–411. ISBN: 978-0-19-854027-4 978-1-383-02689-4. DOI: [10.1093/oso/9780198540274.003.0023](https://doi.org/10.1093/oso/9780198540274.003.0023).
- Sharon, I., Quijada, N.M., Pasolli, E., Fabbri, M., Vitali, F., Agamennone, V., Dötsch, A., Selberherr, E., Grau, J.H., Meixner, M., Liere, K., Ercolini, D., De Filippo, C., Caderni, G., Brigidi, P., and Turrone, S., 2022. The Core Human Microbiome: Does It Exist and How Can We Find It? A Critical Review of the Concept. *Nutrients* 14.14, 2872. ISSN: 2072-6643. DOI: [10.3390/nu14142872](https://doi.org/10.3390/nu14142872).
- Shipley, B., 2016. *Cause and Correlation in Biology: A User's Guide to Path Analysis, Structural Equations and Causal Inference* 1. paperback ed., with corr., transferred to digital print. (Cambridge: Cambridge Univ. Press). ISBN: 978-0-521-52921-1 978-0-521-79153-3.
- Sireci, M., Muñoz, M.A., and Grilli, J., 2023. Environmental Fluctuations Explain the Universal Decay of Species-Abundance Correlations with Phylogenetic Distance. *Proceedings of the National Academy of Sciences* 120.37, e2217144120. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2217144120](https://doi.org/10.1073/pnas.2217144120).
- Smale, S., 1967. Differentiable Dynamical Systems. *Bulletin of the American Mathematical Society* 73.6, 747–817. ISSN: 0273-0979, 1088-9485. DOI: [10.1090/S0002-9904-1967-11798-1](https://doi.org/10.1090/S0002-9904-1967-11798-1).

- Smale, S., 1976. On the Differential Equations of Species in Competition. *Journal of Mathematical Biology* 3.1, 5–7. ISSN: 0303-6812, 1432-1416. DOI: [10.1007/BF00307854](https://doi.org/10.1007/BF00307854).
- Soetaert, K., Petzoldt, T., and Setzer, R.W., 2010. Solving Differential Equations in R: Package deSolve. *Journal of Statistical Software* 33.9, 1–25. DOI: [10.18637/Jss.V033.I09](https://doi.org/10.18637/Jss.V033.I09).
- Solé, R.V. and Valls, J., 1992. On Structural Stability and Chaos in Biological Systems. *Journal of Theoretical Biology* 155.1, 87–102. ISSN: 00225193. DOI: [10.1016/S0022-5193\(05\)80550-8](https://doi.org/10.1016/S0022-5193(05)80550-8).
- Sommer, F., Anderson, J.M., Bharti, R., Raes, J., and Rosenstiel, P., 2017. The Resilience of the Intestinal Microbiota Influences Health and Disease. *Nature Reviews Microbiology* 15.10, 630–638. ISSN: 1740-1526, 1740-1534. DOI: [10.1038/nrmicro.2017.58](https://doi.org/10.1038/nrmicro.2017.58).
- Song, C., Rohr, R.P., Vasseur, D., and Saavedra, S., 2020. Disentangling the Effects of External Perturbations on Coexistence and Priority Effects. *Journal of Ecology* 108.4. Akiko Satake, ed., 1677–1689. ISSN: 0022-0477, 1365-2745. DOI: [10.1111/1365-2745.13349](https://doi.org/10.1111/1365-2745.13349).
- Song, C. and Saavedra, S., 2018. Will a Small Randomly Assembled Community Be Feasible and Stable? *Ecology* 99.3, 743–751. ISSN: 0012-9658, 1939-9170. DOI: [10.1002/ecy.2125](https://doi.org/10.1002/ecy.2125).
- 2021. Bridging Parametric and Nonparametric Measures of Species Interactions Unveils New Insights of Non-equilibrium Dynamics. *Oikos* 130.7, 1027–1034. ISSN: 0030-1299, 1600-0706. DOI: [10.1111/oik.08060](https://doi.org/10.1111/oik.08060).
- Sprockett, D., Fukami, T., and Relman, D.A., 2018. Role of Priority Effects in the Early-Life Assembly of the Gut Microbiota. *Nature Reviews Gastroenterology & Hepatology* 15.4, 197–205. ISSN: 1759-5045, 1759-5053. DOI: [10.1038/nrgastro.2017.173](https://doi.org/10.1038/nrgastro.2017.173).
- Stark, J., 1999. Delay Embeddings for Forced Systems. I. Deterministic Forcing. *Journal of Nonlinear Science* 9.3, 255–332. ISSN: 0938-8974, 1432-1467. DOI: [10.1007/s003329900072](https://doi.org/10.1007/s003329900072).
- Stark, J., Broomhead, D., Davies, M., and Huke, J., 2003. Delay Embeddings for Forced Systems. II. Stochastic Forcing. *Journal of Nonlinear Science* 13.6, 519–577. ISSN: 0938-8974, 1432-1467. DOI: [10.1007/s00332-003-0534-4](https://doi.org/10.1007/s00332-003-0534-4).
- Stone, L., 2020. The Stability of Mutualism. *Nature Communications* 11.1, 2648. ISSN: 2041-1723. DOI: [10.1038/s41467-020-16474-4](https://doi.org/10.1038/s41467-020-16474-4).
- Strasser, B., Wolters, M., Weyh, C., Krüger, K., and Ticinesi, A., 2021. The Effects of Lifestyle and Diet on Gut Microbiota Composition, Inflammation and Muscle Performance in Our Aging Society. *Nutrients* 13.6, 2045. ISSN: 2072-6643. DOI: [10.3390/nu13062045](https://doi.org/10.3390/nu13062045).
- Sugihara, G., 1994. Nonlinear Forecasting for the Classification of Natural Time Series. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences* 348.1688, 477–495. ISSN: 0962-8428, 2054-0299. DOI: [10.1098/rsta.1994.0106](https://doi.org/10.1098/rsta.1994.0106).

- Sugihara, G., May, R., Ye, H., Hsieh, C.-h., Deyle, E., Fogarty, M., and Munch, S., 2012. Detecting Causality in Complex Ecosystems. *Science* 338.6106, 496–500. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1227079](https://doi.org/10.1126/science.1227079).
- Sugihara, G. and May, R.M., 1990. Nonlinear Forecasting as a Way of Distinguishing Chaos from Measurement Error in Time Series. *Nature* 344.6268, 734–741. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/344734a0](https://doi.org/10.1038/344734a0).
- Sundarraman, D., Hay, E.A., Martins, D.M., Shields, D.S., Pettinari, N.L., and Parthasarathy, R., 2020. Higher-Order Interactions Dampen Pairwise Competition in the Zebrafish Gut Microbiome. *mBio* 11.5. William Ludington and Edward G. Ruby, ed., e01667–20. ISSN: 2161-2129, 2150-7511. DOI: [10.1128/mBio.01667-20](https://doi.org/10.1128/mBio.01667-20).
- Suzuki, T.A. and Worobey, M., 2014. Geographical Variation of Human Gut Microbial Composition. *Biology Letters* 10.2, 20131037. ISSN: 1744-9561, 1744-957X. DOI: [10.1098/rsbl.2013.1037](https://doi.org/10.1098/rsbl.2013.1037).
- Swain, A., Fussell, L., and Fagan, W.F., 2022. Higher-Order Effects, Continuous Species Interactions, and Trait Evolution Shape Microbial Spatial Dynamics. *Proceedings of the National Academy of Sciences* 119.1, e2020956119. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2020956119](https://doi.org/10.1073/pnas.2020956119).
- Tabi, A., Gilarranz, L.J., Wood, S.A., Dunne, J.A., and Saavedra, S., 2023. Protection Promotes Energetically Efficient Structures in Marine Communities. *PLOS Computational Biology* 19.12. Rafael D’Andrea, ed., e1011742. ISSN: 1553-7358. DOI: [10.1371/journal.pcbi.1011742](https://doi.org/10.1371/journal.pcbi.1011742).
- Takens, F., 1981. “Detecting Strange Attractors in Turbulence.” *Dynamical Systems and Turbulence, Warwick 1980*, David Rand and Lai-Sang Young, ed. Vol. 898. (Berlin, Heidelberg: Springer Berlin Heidelberg), pp. 366–381. ISBN: 978-3-540-11171-9 978-3-540-38945-3. DOI: [10.1007/BFb0091924](https://doi.org/10.1007/BFb0091924).
- Tan, J., McKenzie, C., Potamitis, M., Thorburn, A.N., Mackay, C.R., and Macia, L., 2014. “The Role of Short-Chain Fatty Acids in Health and Disease.” *Advances in Immunology*, vol. 121. (Elsevier), pp. 91–119. ISBN: 978-0-12-800100-4. DOI: [10.1016/B978-0-12-800100-4.00003-9](https://doi.org/10.1016/B978-0-12-800100-4.00003-9).
- Tanaka, Y. and Benno, Y., 2015. Application of a Single-colony Coculture Technique to the Isolation of Hitherto Unculturable Gut Bacteria. *Microbiology and Immunology* 59.2, 63–70. ISSN: 0385-5600, 1348-0421. DOI: [10.1111/1348-0421.12220](https://doi.org/10.1111/1348-0421.12220).
- Taylor, L.R., 1961. Aggregation, Variance and the Mean. *Nature* 189.4766, 732–735. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/189732a0](https://doi.org/10.1038/189732a0).
- Thom, R., 1994. *Structural Stability And Morphogenesis* (CRC Press). ISBN: 978-0-429-49302-7. DOI: [10.1201/9780429493027](https://doi.org/10.1201/9780429493027).
- Thomas, W.R., Pomerantz, M.J., and Gilpin, M.E., 1980. Chaos, Asymmetric Growth and Group Selection for Dynamical Stability. *Ecology* 61.6, 1312–1320. ISSN: 0012-9658, 1939-9170. DOI: [10.2307/1939039](https://doi.org/10.2307/1939039).

- Tian, L., Wang, X.-W., Wu, A.-K., Fan, Y., Friedman, J., Dahlin, A., Waldor, M.K., Weinstock, G.M., Weiss, S.T., and Liu, Y.-Y., 2020. Deciphering Functional Redundancy in the Human Microbiome. *Nature Communications* 11.1, 6217. ISSN: 2041-1723. DOI: [10.1038/s41467-020-19940-1](https://doi.org/10.1038/s41467-020-19940-1).
- Tilman, D., 1996. Biodiversity: Population Versus Ecosystem Stability. *Ecology* 77.2, 350–363. ISSN: 0012-9658, 1939-9170. DOI: [10.2307/2265614](https://doi.org/10.2307/2265614).
- Turchin, P. and Taylor, A.D., 1992. Complex Dynamics in Ecological Time Series. *Ecology* 73.1, 289–305. ISSN: 0012-9658, 1939-9170. DOI: [10.2307/1938740](https://doi.org/10.2307/1938740).
- Turnbaugh, P.J., Ley, R.E., Hamady, M., Fraser-Liggett, C.M., Knight, R., and Gordon, J.I., 2007. The Human Microbiome Project. *Nature* 449.7164, 804–810. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature06244](https://doi.org/10.1038/nature06244).
- Turnbaugh, P.J., Ridaura, V.K., Faith, J.J., Rey, F.E., Knight, R., and Gordon, J.I., 2009. The Effect of Diet on the Human Gut Microbiome: A Metagenomic Analysis in Humanized Gnotobiotic Mice. *Science Translational Medicine* 1.6. ISSN: 1946-6234, 1946-6242. DOI: [10.1126/scitranslmed.3000322](https://doi.org/10.1126/scitranslmed.3000322).
- Ursell, L.K., Metcalf, J.L., Parfrey, L.W., and Knight, R., 2012. Defining the Human Microbiome. *Nutrition Reviews* 70, S38–S44. ISSN: 00296643. DOI: [10.1111/j.1753-4887.2012.00493.x](https://doi.org/10.1111/j.1753-4887.2012.00493.x).
- Ushio, M., Hsieh, C.-h., Masuda, R., Deyle, E.R., Ye, H., Chang, C.-W., Sugihara, G., and Kondoh, M., 2018. Fluctuating Interaction Network and Time-Varying Stability of a Natural Fish Community. *Nature* 554.7692, 360–363. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature25504](https://doi.org/10.1038/nature25504).
- Valdes, A.M., Walter, J., Segal, E., and Spector, T.D., 2018. Role of the Gut Microbiota in Nutrition and Health. *BMJ*, k2179. ISSN: 0959-8138, 1756-1833. DOI: [10.1136/bmj.k2179](https://doi.org/10.1136/bmj.k2179).
- Van De Guchte, M., Burz, S.D., Cadiou, J., Wu, J., Mondot, S., Blottière, H.M., and Doré, J., 2020. Alternative Stable States in the Intestinal Ecosystem: Proof of Concept in a Rat Model and a Perspective of Therapeutic Implications. *Microbiome* 8.1, 153. ISSN: 2049-2618. DOI: [10.1186/s40168-020-00933-7](https://doi.org/10.1186/s40168-020-00933-7).
- Van Meerbeek, K., Jucker, T., and Svenning, J.-C., 2021. Unifying the Concepts of Stability and Resilience in Ecology. *Journal of Ecology* 109.9, 3114–3132. ISSN: 0022-0477, 1365-2745. DOI: [10.1111/1365-2745.13651](https://doi.org/10.1111/1365-2745.13651).
- Vandeputte, D., Falony, G., Vieira-Silva, S., Tito, R.Y., Joossens, M., and Raes, J., 2016. Stool Consistency Is Strongly Associated with Gut Microbiota Richness and Composition, Enterotypes and Bacterial Growth Rates. *Gut* 65.1, 57–62. ISSN: 0017-5749, 1468-3288. DOI: [10.1136/gutjnl-2015-309618](https://doi.org/10.1136/gutjnl-2015-309618).
- Vano, J.A., Wildenberg, J.C., Anderson, M.B., Noel, J.K., and Sprott, J.C., 2006. Chaos in Low-Dimensional Lotka–Volterra Models of Competition. *Nonlinearity* 19.10, 2391–2404. ISSN: 0951-7715, 1361-6544. DOI: [10.1088/0951-7715/19/10/006](https://doi.org/10.1088/0951-7715/19/10/006).
- Verleysen, M. and François, D., 2005. “The Curse of Dimensionality in Data Mining and Time Series Prediction.” *Computational Intelligence and Bioinspired Systems*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu

- Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Dough Tygar, Moshe Y. Vardi, Gerhard Weikum, Joan Cabestany, Alberto Prieto, and Francisco Sandoval, ed. Vol. 3512. (Berlin, Heidelberg: Springer Berlin Heidelberg), pp. 758–770. ISBN: 978-3-540-26208-4 978-3-540-32106-4. DOI: [10.1007/11494669_93](https://doi.org/10.1007/11494669_93).
- Verma, T. and Pearl, J., 1990. “Equivalence and Synthesis of Causal Models.” *Probabilistic and Causal Inference*, Hector Geffner, Rina Dechter, and Joseph Y. Halpern, ed. 1st ed. (New York, NY, USA: ACM), pp. 221–236. ISBN: 978-1-4503-9586-1. DOI: [10.1145/3501714.3501732](https://doi.org/10.1145/3501714.3501732).
- Volterra, V., 1926. Fluctuations in the Abundance of a Species Considered Mathematically¹. *Nature* 118.2972, 558–560. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/118558a0](https://doi.org/10.1038/118558a0).
- Wade, W.G., 2013. The Oral Microbiome in Health and Disease. *Pharmacological Research* 69.1, 137–143. ISSN: 10436618. DOI: [10.1016/j.phrs.2012.11.006](https://doi.org/10.1016/j.phrs.2012.11.006).
- Wagg, C., Schlaeppli, K., Banerjee, S., Kuramae, E.E., and Van Der Heijden, M.G.A., 2019. Fungal-Bacterial Diversity and Microbiome Complexity Predict Ecosystem Functioning. *Nature Communications* 10.1, 4841. ISSN: 2041-1723. DOI: [10.1038/s41467-019-12798-y](https://doi.org/10.1038/s41467-019-12798-y).
- Wang, X.-W., Sun, Z., Jia, H., Michel-Mata, S., Angulo, M.T., Dai, L., He, X., Weiss, S.T., and Liu, Y.-Y., 2023. *Identifying Keystone Species in Microbial Communities Using Deep Learning*. Preprint. Microbiology. DOI: [10.1101/2023.03.15.532858](https://doi.org/10.1101/2023.03.15.532858).
- Weinans, E., Quax, R., Van Nes, E.H., and Leemput, I.A.V.D., 2021. Evaluating the Performance of Multivariate Indicators of Resilience Loss. *Scientific Reports* 11.1, 9148. ISSN: 2045-2322. DOI: [10.1038/s41598-021-87839-y](https://doi.org/10.1038/s41598-021-87839-y).
- Weiss, A.S., Niedermeier, L.S., Von Stempel, A., Burrichter, A.G., Ring, D., Meng, C., Kleigrewe, K., Lincetto, C., Hübner, J., and Stecher, B., 2023. Nutritional and Host Environments Determine Community Ecology and Keystone Species in a Synthetic Gut Bacterial Community. *Nature Communications* 14.1, 4780. ISSN: 2041-1723. DOI: [10.1038/s41467-023-40372-0](https://doi.org/10.1038/s41467-023-40372-0).
- Wickham, H., 2016. *ggplot2: Elegant Graphics for Data Analysis* (Springer-Verlag New York). ISBN: 978-3-319-24277-4.
- Wilbert, S.A., Mark Welch, J.L., and Borisy, G.G., 2020. Spatial Ecology of the Human Tongue Dorsum Microbiome. *Cell Reports* 30.12, 4003–4015.e3. ISSN: 22111247. DOI: [10.1016/j.celrep.2020.02.097](https://doi.org/10.1016/j.celrep.2020.02.097).
- Winter, S.E. and Bäumlér, A.J., 2023. Gut Dysbiosis: Ecological Causes and Causative Effects on Human Disease. *Proceedings of the National Academy of Sciences* 120.50, e2316579120. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2316579120](https://doi.org/10.1073/pnas.2316579120).
- Wolff, R., Shoemaker, W., and Garud, N., 2023. Ecological Stability Emerges at the Level of Strains in the Human Gut Microbiome. *mBio* 14.2. Rachel Whitaker and John W. Taylor, ed., e02502–22. ISSN: 2150-7511. DOI: [10.1128/mbio.02502-22](https://doi.org/10.1128/mbio.02502-22).

- Wood, S.N. and Thomas, M.B., 1999. Super-Sensitivity to Structure in Biological Models. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 266.1419, 565–570. ISSN: 0962-8452, 1471-2954. DOI: [10.1098/rspb.1999.0673](https://doi.org/10.1098/rspb.1999.0673).
- Worby, C.J., Sridhar, S., Turbett, S.E., Becker, M.V., Kogut, L., Sanchez, V., Bronson, R.A., Rao, S.R., Oliver, E., Walker, A.T., Walters, M.S., Kelly, P., Leung, D.T., Knouse, M.C., Hagmann, S.H.F., Harris, J.B., Ryan, E.T., Earl, A.M., and LaRocque, R.C., 2023. Gut Microbiome Perturbation, Antibiotic Resistance, and Escherichia Coli Strain Dynamics Associated with International Travel: A Metagenomic Analysis. *The Lancet Microbe* 4.10, e790–e799. ISSN: 26665247. DOI: [10.1016/S2666-5247\(23\)00147-7](https://doi.org/10.1016/S2666-5247(23)00147-7).
- Wu, G.D., Chen, J., Hoffmann, C., Bittinger, K., Chen, Y.-Y., Keilbaugh, S.A., Bewtra, M., Knights, D., Walters, W.A., Knight, R., Sinha, R., Gilroy, E., Gupta, K., Baldassano, R., Nessel, L., Li, H., Bushman, F.D., and Lewis, J.D., 2011. Linking Long-Term Dietary Patterns with Gut Microbial Enterotypes. *Science* 334.6052, 105–108. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1208344](https://doi.org/10.1126/science.1208344).
- Ye, H., Beamish, R.J., Glaser, S.M., Grant, S.C.H., Hsieh, C.-h., Richards, L.J., Schnute, J.T., and Sugihara, G., 2015. Equation-Free Mechanistic Ecosystem Forecasting Using Empirical Dynamic Modeling. *Proceedings of the National Academy of Sciences* 112.13. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1417063112](https://doi.org/10.1073/pnas.1417063112).
- Ye, H. and Sugihara, G., 2016. Information Leverage in Interconnected Ecosystems: Overcoming the Curse of Dimensionality. *Science* 353.6302, 922–925. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.aag0863](https://doi.org/10.1126/science.aag0863).
- Yodzis, P., 1988. The Indeterminacy of Ecological Interactions as Perceived Through Perturbation Experiments. *Ecology* 69.2, 508–515. ISSN: 0012-9658, 1939-9170. DOI: [10.2307/1940449](https://doi.org/10.2307/1940449).
- Zaura, E., Brandt, B.W., Teixeira De Mattos, M.J., Buijs, M.J., Caspers, M.P.M., Rashid, M.-U., Weintraub, A., Nord, C.E., Savell, A., Hu, Y., Coates, A.R., Hubank, M., Spratt, D.A., Wilson, M., Keijsers, B.J.F., and Crielaard, W., 2015. Same Exposure but Two Radically Different Responses to Antibiotics: Resilience of the Salivary Microbiome versus Long-Term Microbial Shifts in Feces. *mBio* 6.6. Alex Van Belkum, ed., e01693–15. ISSN: 2161-2129, 2150-7511. DOI: [10.1128/mBio.01693-15](https://doi.org/10.1128/mBio.01693-15).
- Zeeman, E.C., 1977. *Catastrophe Theory: Selected Papers, 1972 - 1977* 3. print. (Reading, Mass: Addison-Wesley). ISBN: 978-0-201-09015-4 978-0-201-09014-7.
- Zeileis, A. and Grothendieck, G., 2005. zoo: S3 Infrastructure for Regular and Irregular Time Series. *Journal of Statistical Software* 14.6, 1–27. DOI: [10.18637/Jss.V014.I06](https://doi.org/10.18637/Jss.V014.I06).
- Zhang, Z., Geng, J., Tang, X., Fan, H., Xu, J., Wen, X., Ma, Z., and Shi, P., 2014. Spatial Heterogeneity and Co-Occurrence Patterns of Human Mucosal-Associated Intestinal Microbiota. *The ISME Journal* 8.4, 881–893. ISSN: 1751-7362, 1751-7370. DOI: [10.1038/ismej.2013.185](https://doi.org/10.1038/ismej.2013.185).
- Zhao, Q., Van Den Brink, P.J., Xu, C., Wang, S., Clark, A.T., Karakoç, C., Sugihara, G., Widdicombe, C.E., Atkinson, A., Matsuzaki, S.-i.S., Shinohara, R., He, S., Wang, Y.X.G., and De Laender, F., 2023. Rela-

- tionships of Temperature and Biodiversity with Stability of Natural Aquatic Food Webs. *Nature Communications* 14.1, 3507. ISSN: 2041-1723. DOI: [10.1038/s41467-023-38977-6](https://doi.org/10.1038/s41467-023-38977-6).
- Zou, H. and Hastie, T., 2005. Regularization and variable selection via the elastic net. en. *J. R. Stat. Soc. Series B Stat. Methodol.* 67.2, 301–320.
- Zou, S., Fang, L., and Lee, M.-H., 2018. Dysbiosis of Gut Microbiota in Promoting the Development of Colorectal Cancer. *Gastroenterology Report* 6.1, 1–12. ISSN: 2052-0034. DOI: [10.1093/gastro/gox031](https://doi.org/10.1093/gastro/gox031).