

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Simplification and control of microbial ecosystems in theory and experiment

Permalink

<https://escholarship.org/uc/item/5kd0c75h>

Author

Jones, Eric

Publication Date

2020

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

**Simplification and control of microbial ecosystems
in theory and experiment**

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Eric William Jones

Committee in charge:

Professor Jean Carlson, Chair
Professor Roger Nisbet
Professor Zvonimir Dogic

June 2020

The Dissertation of Eric William Jones is approved.

Professor Roger Nisbet

Professor Zvonimir Dogic

Professor Jean Carlson, Committee Chair

June 2020

Simplification and control of microbial ecosystems
in theory and experiment

Copyright © 2020

by

Eric William Jones

For my granddad, Dr. John M. Richardson,
who inspired me to study physics

Acknowledgements

I would first like to thank my mom for her words of wisdom that she spoke to me as I was about to begin a grueling, wearisome, and at times overwhelming trek: “take things one day at a time.” She wasn’t talking about my PhD, but she might as well have been. Mom, thank you for being a trail angel in all things. Thanks to my dad for teaching me how to work hard at something until it is of high quality. Thanks to my brothers, with whom I have celebrated, toiled, and become closer with these last few years.

Thanks to Jean, who has been a wonderful PhD advisor. She taught me how to conduct myself as a professional, as a writer, as a teacher, and as a mentor.

Thanks to my PhD committee members. Roger, you are a role model to me both for your ecological mind and in how you interact with others. Zvonimir, I have thoroughly enjoyed meeting and talking with you.

Thanks to my collaborators, who joyfully shared their expertise with me. Will, thank you for always advocating for me, and for introducing me to the world of microbial ecology. I am thrilled that I get to continue working with you. Shenshen, thank you for your advice and guidance.

Thanks to the members of the Ecology, Evolution, and Marine Biology department that so readily accepted me into their theoretical ecology community. I have been privileged to learn from great ecologists like Steve, Cherie, Holly, and Ferdinand.

I acknowledge the National Science Foundation Graduate Research Fellowship Program (Grant No. 1650114), the David and Lucile Packard Foundation, and the Institute for Collaborative Biotechnologies (contract no. W911NF-09-D-0001 from the U.S. Army Research Office) for their role in funding this thesis.

To all the friends I met and travelled alongside with in graduate school, thanks for making it fun. Thanks for being a great roommate and friend, Dillon. Thanks Brianna

for caring about things. Thanks Josh for all of our walks. Thanks Seth for all of our talks. Thanks for the hikes, Alex. Thanks for skiing with me, Neelay. Thanks for surfing with me, Gabriel. Thanks Peter for the chess. And to everyone else with whom I have enjoyed my time in Santa Barbara, thank you.

Thanks to my friends who make it so that I always call Colorado “home.” Thanks John, for being up for anything. Thanks Alli, for all the adventures. And thanks to everyone else that I have ever been in the mountains with.

And lastly thanks to you, the reader, for engaging with this body of work. I hope you enjoy it.

Curriculum Vitæ

Eric William Jones

Education

2020	Ph.D. in Physics (Expected), University of California, Santa Barbara
2018	M.A. in Physics, University of California, Santa Barbara
2015	B.S. in Engineering Physics, Colorado School of Mines
2015	B.S. in Computational and Applied Mathematics, Colorado School of Mines

Publications

1. Z. Wang[†], **E. Jones**, J. Mueller, and J. Carlson (2020). “Control of ecological outcomes through deliberate parameter changes in a model of the gut microbiome.” *Physical Review E* 101(5):052402 [\[link\]](#)
2. **E. Jones** and J. Carlson (2019). “Steady-state reduction of generalized Lotka-Volterra systems in the microbiome.” *Physical Review E* 99(3):032403 [\[link\]](#)
3. A. Gould, V. Zhang, L. Lamberti, **E. Jones**, B. Obadia, N. Korasidis, A. Gavryushkin, J. Carlson, N. Beerenwinkel, and W. Ludington (2018). “Microbiome interactions shape host fitness.” *Proceedings of the National Academy of Sciences* 115(51):E11951 [\[link\]](#)
4. **E. Jones** and J. Carlson (2018). “In silico analysis of antibiotic-induced *Clostridium difficile* infection.” *PLoS Computational Biology* 14(2):e1006001 [\[link\]](#)
5. P. Diaz, P. Constantine, K. Kalmbach, **E. Jones**, and S. Pankavich (2018). “A modified SEIR model for the spread of Ebola in Western Africa and metrics for resource allocation.” *Applied Mathematics and Computation* 324:141 [\[link\]](#)
6. **E. Jones**, P. Roemer, M. Raghupathi, and S. Pankavich (2013). “Analysis and simulation of the three-component model of HIV dynamics,” *SIAM Undergraduate Research Online* 7:89 [\[link\]](#)

SUBMITTED:

7. P. Shankin-Clarke[†], **E. Jones**, and J. Carlson. “Efficient navigation and control of outcomes in generalized Lotka-Volterra systems in the microbiome.” [\[arXiv link\]](#)
8. **E. Jones**^{*}, J. Sheng^{*}, S. Wang, and J. Carlson. “Aging and fragility in a coupled innate-adaptive immune model.”

^{*}equal contribution; [†]undergraduate research advisee

Abstract

Simplification and control of microbial ecosystems
in theory and experiment

by

Eric William Jones

Over the last two decades an association between microbiome composition and some human diseases has been unambiguously established, and this discovery has provoked clinical interest into microbiome-based medical therapies. The advent of high-throughput sequencing has produced vast amounts of microbial abundance data, but mechanistic models that describe and predict this complex microbial ecosystem are not yet established.

In my doctoral research I employ novel theoretical approaches and tractable experimental systems to study simplified instances of the complex microbial dynamics of the microbiome. This research is motivated by a desire to inform the mechanism of action and development of microbiome-based bacteriotherapies. We use generalized Lotka-Volterra (gLV) models, a class of model that exhibits prototypical ecological behaviors, as a theoretical proxy for true microbial dynamics. We develop a numerical framework to predict how introduced foreign microbes (direct bacteriotherapy) and a modified microbiome environment (indirect bacteriotherapy) affect the composition of a microbiome. Additionally, we derive the dimensionality-reduction technique “steady-state reduction” (SSR), which compresses bistable dynamics in a high-dimensional gLV system into a reduced two-dimensional system. The insights gained from this reduced system inform the microbial dynamics of high-dimensional gLV systems, and therefore contribute to our knowledge of how bacteriotherapies function.

Additionally, in experimental work I construct simple models from first principles that describe how the measured physiological traits depend on microbiome composition in the fruit fly *Drosophila melanogaster*. The fruit fly microbiome naturally hosts only a few species of commensal bacteria. In our research we cultivate a core group of five commensal bacteria, and associate each possible combination of the five bacteria with a set of germ-free flies. Since the flies were identically reared, the flies associated with each of these 32 bacterial combinations carry a distinct microbiome-affiliated phenotype (e.g. lifespan or fecundity) that is a function of the complex interactions in the microbiome. We found that the complexity of these fly phenotypes was often reducible: we could approximate the phenotype of flies associated with more than one bacterial species by averaging the phenotypes of the flies with the corresponding single-species associations. Thus, we found that simple models can describe complex behaviors in the fly microbiome.

Contents

Curriculum Vitae	vii
Abstract	viii
1 Introduction	1
1.1 Ecology and engineering within the microbiome	1
1.2 Direct and indirect bacteriotherapies in generalized Lotka-Volterra models	3
1.3 Handling the complexity of the microbiome with steady-state reduction .	4
1.4 An experimentally-derived gLV model of <i>C. difficile</i> infection	6
1.5 Experimental findings in the fruit fly microbiome motivate bacteriotherapies	8
1.6 Thesis overview	9
1.7 Permissions and Attributions	10
2 Microbiome composition affects host physiology in the fruit fly	12
2.1 Abstract of “Microbiome interactions shape host fitness”	13
2.2 Significance	14
2.3 Introduction	15
2.4 Results	17
2.4.1 Microbiome diversity confers a life history tradeoff	17
2.4.2 Reproduction cannot be increased by mid-life microbiome addition	20
2.4.3 Microbiome interactions can change host physiology	21
2.4.4 Higher-order microbiome interactions change host physiology . . .	22
2.4.5 Bacterial presence-absence more than abundance impacts fly phys- iology	24
2.4.6 Live bacteria speed up fly development	28
2.4.7 Microbial abundance interactions correlate with host physiology interactions	29
2.5 Discussion	32
2.5.1 Bacterial abundance interactions may damage the host.	32
2.5.2 How much do higher-order interactions matter?	32

2.5.3	Microbiome interactions mediate a life history tradeoff between lifespan and fecundity	33
2.5.4	The <i>Drosophila</i> gut microbiome serves as an effective model of microbiome complexity	33
3	Implementing direct bacteriotherapies in an ecological model of the microbiome	35
3.1	Abstract of “ <i>In silico</i> analysis of antibiotic-induced <i>C. difficile</i> infection”	36
3.2	Significance	37
3.3	Introduction	38
3.4	Background	40
3.5	Models and Methods	41
3.5.1	The generalized Lotka-Volterra equations	41
3.5.2	Simulation of treatments for <i>C. difficile</i> infection	42
3.5.3	Implementation of simulated microbial transplants	45
3.5.4	Implementation of sporulation	45
3.5.5	Implementation of mutation	47
3.6	Results	49
3.6.1	Mapping system behaviors of the CDI model	49
3.6.2	Invadability of <i>C. difficile</i>	55
3.6.3	Simulated microbial transplants affect steady-state outcomes	57
3.6.4	Simulated antibiotic dosing alters gLV dynamics	61
3.6.5	Sporulation affects microbial dynamics	64
3.6.6	Mutation in a gLV model	65
3.7	Discussion	68
3.7.1	Applications towards experimental explanation, model validation, and model-motivated experiments	68
3.7.2	Unidentifiability of beneficial bacterial communities	69
3.7.3	Pharmacokinetic and pharmacodynamic approximations	70
3.7.4	Limitations of the gLV model	71
3.7.5	Analytic concerns of parameter fitting	72
3.7.6	Combining gLV and SIR techniques	72
3.8	Conclusion	73
4	Compression of high-dimensional bistable dynamics with steady-state reduction (SSR)	74
4.1	Abstract of “Steady-state reduction of generalized Lotka-Volterra systems in the microbiome”	76
4.2	Introduction	76
4.3	Compression of generalized Lotka-Volterra systems	78
4.4	Steady-state reduction applied to a microbiome model	81
4.5	Analysis of the 2D gLV equations	85

4.5.1	Explicit form of the separatrix	85
4.5.2	Dynamical landscape of the 2D gLV equations	87
4.6	SSR applied to fecal microbiota transplantation	93
4.7	Discussion	96
4.7.1	Compression of complex ecological systems	96
4.7.2	Simplification of gLV-based FMT frameworks	97
4.8	Conclusion	98
4.9	Derivation of steady-state reduction	99
5	Control of ecological outcomes with direct bacteriotherapies	104
5.1	Abstract for “Navigation and control of outcomes in a generalized Lotka-Volterra model of the microbiome”	105
5.2	Introduction	106
5.3	Background	108
5.3.1	Generalized Lotka-Volterra (gLV) models	108
5.3.2	Experimentally-derived gLV model of <i>C. difficile</i> infection	110
5.3.3	Approximation of bistable gLV dynamics	113
5.3.4	Transplant size and timing affect the efficacy of direct interventions	118
5.4	Results	120
5.4.1	Transplant compositions and their success rates	120
5.4.2	The dynamical landscape of a gLV system	123
5.4.3	Efficient navigation of an attractor network	125
5.4.4	Efficient construction of attractor networks with steady-state reduction	129
5.5	Discussion	131
5.5.1	Microbiome composition is associated with host health	131
5.5.2	Advances in modeling the microbiome	133
5.6	Conclusion	135
6	Control of ecological outcomes with indirect bacteriotherapies	136
6.1	Abstract of “Control of ecological outcomes through deliberate parameter changes in a model of the gut microbiome”	137
6.2	Introduction	138
6.3	Materials and methods	140
6.3.1	The generalized Lotka-Volterra equations	140
6.3.2	A gLV model fit to experimental data	141
6.3.3	SSR-guided parameter change (SPARC)	144
6.4	Results	150
6.4.1	Steady-state reduction (SSR) produces a 2D approximation to bistable dynamics in a high-dimensional gut microbiome model	151
6.4.2	Bifurcation analysis guides interaction parameter changes that modify steady-state outcomes in reduced 2D gLV systems	154

6.4.3	SSR maps low-dimensional bifurcation behavior to the high-dimensional system	156
6.4.4	SPARC generates a finite-time intervention that drives a disease-prone initial condition towards a healthy state in the CDI model	157
6.4.5	SPARC successfully changes steady state outcomes in synthetic gLV models	159
6.5	Discussion	164
6.5.1	SPARC is efficient and flexible	164
6.5.2	Perturbing ecological interactions indirectly controls steady-state outcomes	166
6.5.3	SPARC provides a lens for understanding the effect of the environment on microbial composition	167
6.6	Conclusion	168
7	Conclusion	170
7.1	Overview	170
7.2	Future work: A statistical mechanical treatment of the microbiome	171
7.3	Future work: Experimental validation in the fruit fly microbiome	174
7.4	Closing remarks	174
	Bibliography	176

Chapter 1

Introduction

1.1 Ecology and engineering within the microbiome

The combined might of ecology and engineering will be needed to combat the pressing environmental concerns that threaten our future. Traditionally, these fields have differed in their approaches to problem solving. Ecologists strive to model the behavior of living populations, and they collect data as inconspicuously as possible so that they do not disturb the natural order they seek to understand. Engineers design machines to perform a desired function, and actively intervene to control their system's behavior. Thus, developing theory that unifies these two fields requires a model system that reconciles their disparate approaches to problem solving: we need an ecological system that is easily replicated and measured so that accurate ecological models can be created, and which is also manipulable so that engineering control protocols may be tested and improved. We propose that the microbiome, which satisfies each of these prerequisites, is an ideal vehicle with which to unite ecology and engineering. In this thesis we address the engineering question of how to control the microbiome by analyzing the effects of bacteriotherapies in an ecological model of the microbiome.

Literally a microscopic biome, the microbiome is composed of trillions of microorganisms that compete to establish or colonize niches in precisely the same way that Darwin’s finches compete to thrive in a suitable habitat. Recent experimental breakthroughs have drastically improved our capacity for high-throughput measurement of the microbiome, revealing both which microbes inhabit the microbiome, and also (through metagenomics) which functions these microbes are carrying out. Especially with the rapid timescale of growth in microbial systems (the lifespan of a bacteria is roughly 12 hours), microbial ecosystems boast datasets whose breadth and precision would have been unimaginable to ecologists fifty years ago.

Beyond its ecological grandeur, the microbiome is linked to host health, and therefore studying the microbiome could feasibly lead to therapeutic outcomes. In particular, the development of bacteriotherapies— medicine for the microbiome— has rapidly emerged as a field of investigation over the past twenty years. Perhaps the most striking success of bacteriotherapies has been the treatment of *Clostridioides difficile* infection (CDI) with fecal microbiota transplantation (FMT), a type of bacteriotherapy. CDI commonly afflicts patients in health care facilities that recently received antibiotics, which deplete the microbiome and render it susceptible to infection by the pathogen *C. difficile*. The cure rate of CDI when FMT is used as a treatment is upwards of 90%, while the cure rate with antibiotics is only 35-70% [1]. FMT works by injecting a subset of a healthy person’s microbiome into a sick person’s microbiome, which replenishes the ailing microbiome with the healthy donor’s commensal bacteria. These newly colonized commensal bacteria then outcompete and exclude the pathogenic *C. difficile*. Thus, the mechanism of FMT is “subtraction by addition.”

Bacteriotherapies (including FMT) at their essence apply an engineering mindset to an ecological system. As more and more associations between the microbiome and host physiology are established— for example, the microbiome trains its host’s immune

system, and produces neurotransmitters as part of the “gut-brain axis”— the therapeutic potential of bacteriotherapies broadens. Many diseases are associated with imbalanced or “dysbiotic” microbiomes, and a key premise of bacteriotherapies is that improving a person’s microbiome composition will lessen the severity of the associated disease. Over the last five years, FMT has been proposed as a treatment for diseases like irritable bowel syndrome, inflammatory bowel disease, ulcerative colitis, and Crohn’s disease, with clinical trials returning promising but as yet inconclusive results [2, 3, 4, 5].

1.2 Direct and indirect bacteriotherapies in generalized Lotka-Volterra models

Motivated by the therapeutic potential of bacteriotherapies, in this thesis I present a numerical framework for the design of direct and indirect bacteriotherapies in a theoretical model of the gut microbiome. Specifically I consider the generalized Lotka-Volterra (gLV) equations, which model ecological dynamics by assuming they arise from pairwise interactions between idealized ecological species. These gLV systems are a classic model of theoretical ecology, and can exhibit prototypical ecological behaviors like competition, mutualism, and parasitism.

We implement direct bacteriotherapies by explicitly adding “healthy” microbes to a “diseased” microbiome state, and this intervention is considered successful when it drives the system towards a “healthy” state. Indirect bacteriotherapies are implemented by altering the environment of the microbiome, which changes the available microbial niches and thus can drive the microbiome’s composition towards a desired state. At their mathematical core, the direct and indirect bacteriotherapies examined in this thesis are control protocols that seek to drive the state of a system away from one equilibrium and

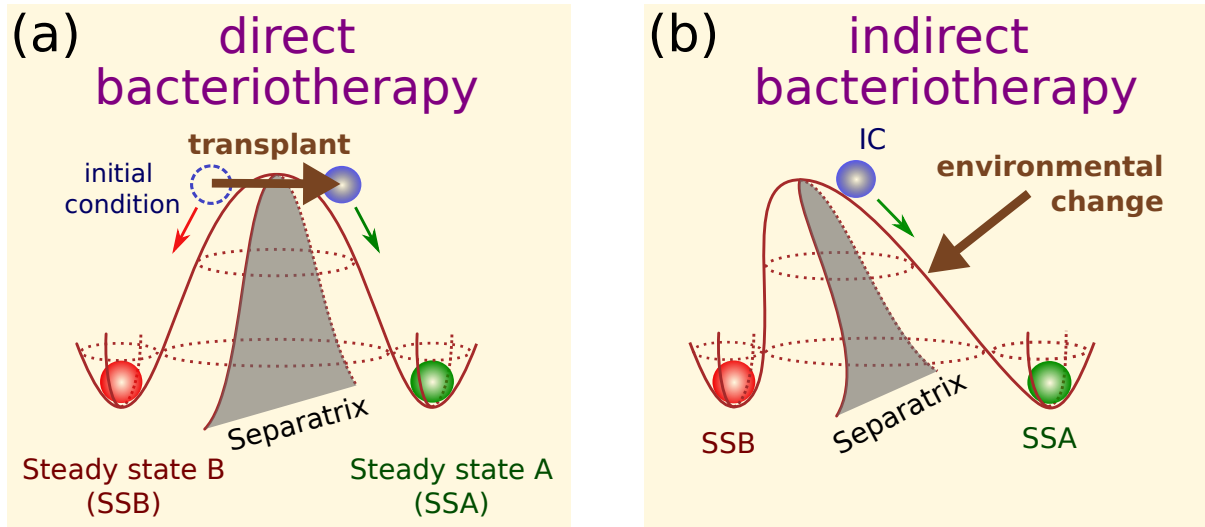


Figure 1.1: **Direct and indirect bacteriotherapies depicted schematically on a pseudoenergy landscape.** The separatrix delineates the basins of attraction of steady states A and B.

towards another target state. Direct control protocols alter the state of the system so that it switches its basin of attraction, while indirect control protocols alter the parameters of the system so that the basin of attraction of the target state is enlarged. In Fig. 1.1 these two control protocols are demonstrated schematically, with ecological dynamics construed as a ball rolling on a pseudoenergy landscape (formally, the landscape is the system's split Lyapunov function). From this perspective, direct bacteriotherapies function by translating the state of the system from one side of the hill to the other, while indirect bacteriotherapies alter the shape of the hill itself.

1.3 Handling the complexity of the microbiome with steady-state reduction

While the intuition behind direct and indirect bacteriotherapies displayed in Fig. 1.1 is relatively straightforward, the complexity inherent in such a multitudinous and inter-

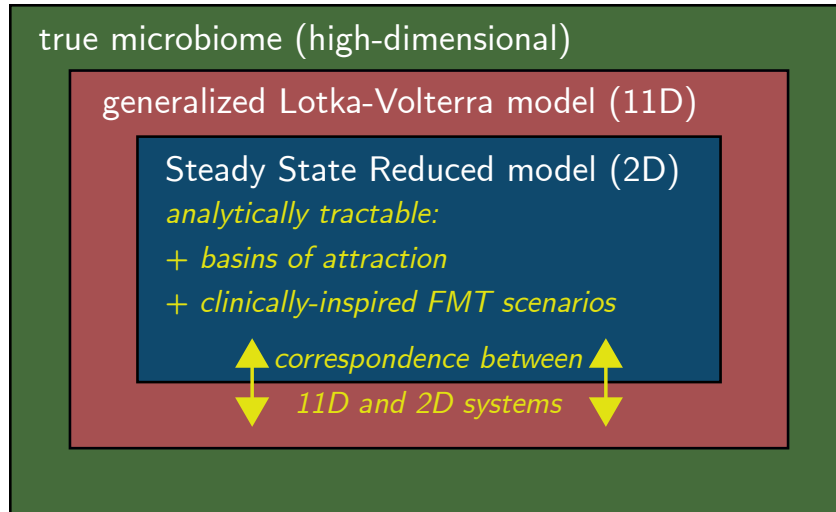


Figure 1.2: **Levels of complexity across ecological systems.** The true microbiome is massively complex, but ecological models (such as generalized Lotka-Volterra (gLV) models) attempt to describe these dynamics mathematically. The CDI model discussed in this thesis predicts the dynamics of 11 microbial populations, and is thus 11-dimensional (11D). Bistable regions of this 11D state space may be compressed with steady-state reduction into an analytically-tractable 2D system.

connected system complicates matters tremendously. Engineers typically study systems with a small number of control parameters whose effect is well-characterized. The microbiome contains so many degrees of freedom and is so interconnected that precisely understanding how a control protocol will affect microbial dynamics is very difficult. To grapple with this complexity, tractable approximations of the microbiome must be employed to guide the choice of intervention.

One such level of abstraction arises from the choice of ecological model— in our case the gLV equations— but even gLV models try to capture interactions between a large number of species, and are therefore high-dimensional and difficult to analyze analytically. In the past, high-dimensional gLV systems have been primarily explored with numerical simulations, but in my thesis I present a dimensionality-reduction technique that allows for high-dimensional behaviors to be compressed in a low-dimensional and analytically-tractable system. By compressing the ecological dynamics of the microbiome,

SSR identifies which “knobs” of the microbiome are the most sensitive. Remarkably, the pseudoenergy landscape of the SSR-reduced system is analytically tractable and qualitatively quite similar to the hill in Fig. 1.1. Thus, the insights gained by SSR directly inform the development of direct and indirect bacteriotherapies.

The levels of complexity of the actual microbiome, gLV models that seek to approximate the true microbial dynamics, and the SSR-reduced model are displayed in Fig. 1.2. Ecologists try to find models that match the empirical ecological dynamics as closely as possible, but engineers are more concerned with understanding which knobs of the system they can manipulate and how these interventions will affect the system. SSR bridges these two levels of abstraction by reducing the space of control parameters that need to be explored, without sacrificing the accuracy of the original high-dimensional model.

1.4 An experimentally-derived gLV model of *C. difficile* infection

In this thesis we primarily rely on a gLV model that was fit by Stein *et al.* to microbial abundance data from a mouse experiment performed by Buffie *et al.* studying antibiotic-induced CDI [7, 6]. The outcomes of this mouse experiment are showcased in Fig. 1.3: green signifies a healthy microbiome, yellow signifies an antibiotic-depleted microbiome, and red signifies a *C. difficile*-infected microbiome. In this experimental study, mice that were given antibiotics (which depleted their microbiomes) and subsequently exposed to *C. difficile* contracted CDI. However, mice that were exposed to *C. difficile* without having taken antibiotics did not contract CDI. Therefore this experimental model system reliably captured the clinical behavior of antibiotic-induced CDI, and demonstrated that the mouse’s gut microbiome state is responsible for whether it is CDI-susceptible or CDI-




Figure 1.3: A schematic description of the *C. difficile* infection (CDI) mouse experiment carried out by Buffie *et al.* [6]. Green mice have healthy microbiomes, yellow mice have antibiotic-depleted microbiomes, and red mice have CDI. CD represents inoculation with the pathogen *C. difficile*, RX represents treatment with the antibiotic clindamycin, and FMT represents administration of fecal microbiota transplantation. The gLV model fit by Stein *et al.* to this mouse experiment also exhibited the transitions between steady states indicated here.

resistant. Therefore, this experiment is a clear-cut demonstration of how microbiome composition can affect host health.

The fit gLV model, which I call the CDI model, also captures the experimentally-observed behavior of antibiotic-induced CDI: the CDI model possesses three steady states that are qualitatively similar to the steady states of the experimental system, and the transitions between the three states (as depicted in Fig. 1.3) are also consistent [7]. In this thesis we treat the CDI model as a mathematical proxy for the microbial dynamics responsible for the onset of CDI. This model system provides clinical motivation for the direct and indirect bacteriotherapies that we devise, and also serves as a testbed for their application.

1.5 Experimental findings in the fruit fly microbiome motivate bacteriotherapies

Finally, in work with experimental collaborator Will Ludington I study the microbiome of the fruit fly *Drosophila melanogaster*. The microbiome of the fruit fly naturally hosts only a few species of commensal bacteria (while the human microbiome hosts thousands). In our research we cultivate a core group of five commensal bacteria, and associate each possible combination of the five bacteria with a set of germ-free flies. Since the flies were identically reared, the flies associated with each of these 32 bacterial combinations carry a distinct microbiome-affiliated phenotype (e.g. lifespan or fecundity) that is a function of the complex interactions in the microbiome. Furthermore, we found that the complexity of these fly phenotypes was often reducible: we could approximate the phenotype of flies associated with more than one bacterial species by averaging the phenotypes of the flies with the corresponding single-species associations.

This research quantitatively established a causal link between microbiome composition and host fitness in the fruit fly. Since bacteriotherapies are predicated on this connection, it is motivating to see such unambiguous evidence in the fruit fly. Additionally, since we are interested in capturing the dynamics of more complex microbiomes with coarse-grained ecological models, it is encouraging that we were able to find simple models that could describe complex behaviors in the fly microbiome in this research. In the future, we hope that the experimental fruit fly system can be used to validate the theoretical direct and indirect control protocols that we discuss in this thesis.

1.6 Thesis overview

In chapter 2 we experimentally demonstrate how microbiome composition affects host health in the fruit fly *Drosophila melanogaster*. This work, carried out in collaboration with Will Ludington, quantitatively demonstrates that fly life span, fecundity, and development time significantly depend on their microbiome composition. Furthermore, we found that the physiological traits of flies with diverse microbiomes could be predicted by the measured traits of flies with simple microbiomes (e.g. containing only one species). This is an instance where complex microbiomes can be understood in terms of simpler microbiomes, which suggests that complexity in the microbiome can at least partially be captured by models. These two experimentally findings motivate the theoretical results discussed in the remainder of the thesis.

Chapter 3 characterizes the dynamical landscape of the CDI model through extensive numerical simulations. This chapter also formalizes our implementation of direct bacteriotherapies in gLV systems, and demonstrates that these simulated bacteriotherapies are capable of driving a disease-prone microbiome composition towards health. However, for direct bacteriotherapies, both the transplant timing and transplant composition are relevant to the efficacy of the treatment.

Chapter 4 introduces the steady-state reduction (SSR) method, and applies it to clinically-relevant scenarios in the CDI model. As a result, the complex ecological dynamics of bistable gLV systems (previously only accessible numerically) can be mapped to a low-dimensional approximate system that is analytically tractable. Specifically, the cause of the time- and composition-dependent efficacy of direct bacteriotherapies is revealed mathematically in terms of the reduced system. A derivation of SSR, proving that it is the best possible two-dimensional gLV approximation of high-dimensional bistable dynamics, is provided at the end of this chapter.

Chapters 5 and 6 employ the computational framework of chapter 3 and the SSR method of chapter 4 to inform the construction of direct and indirect bacteriotherapies that efficiently transition a system towards a target state. Chapter 5 introduces the “attractor network” as a means of compressing the dynamical landscape of a high-dimensional gLV system, by decomposing a multistable high-dimensional phase space into a web of bistable subsystems. This attractor network acts as a “map” of the dynamical landscape, and indicates how direct bacteriotherapies can be administered sequentially to minimize the total required intervention size.

Finally, chapter 6 employs SSR to design indirect bacteriotherapies that drive a bistable system towards a target state by modifying model parameters that alter the system’s dynamics. Since the parameters of gLV models represent idealized ecological interactions between microbial species, altering the parameters of these models corresponds to altering the environment of the microbiome. Due to the complexity of high-dimensional gLV systems, without a dimensionality-reduction tool like SSR, identifying which parameters should be changed to alter the system’s steady-state outcome is computationally costly. Fortunately, SSR provides a straightforward link between the reduced model parameters and the high-dimensional parameters, and can thus identify the most relevant and sensitive parameters of the original model that should be modified in an indirect bacteriotherapy. SSR is also capable of isolating the high-dimensional interactions that lead to bistability, which is an insight that would have been otherwise difficult to identify.

1.7 Permissions and Attributions

The content of chapter 2 is the result of a collaboration with Alison L. Gould, Vivian Zhang, Lisa Lamberti, Benjamin Obadia, Nikolaos Korasidis, Alex Gavryushkin, Niko Beerenwinkel, and William B. Ludington, and has previously appeared in the Proceedings

of the National Academy of Sciences under the title “Microbiome interactions shape host fitness” [8]. It is provided in an abbreviated form here with the permission of William Ludington, the corresponding author of the paper. In this article, I was responsible for the implementation of the averaging model and for analyzing how microbial interactions change as a function of microbiome diversity. For details about the materials and methods used in this experiment, we refer the reader to the original text.

The content of chapter 3 is adapted from the article “In silico analysis of antibiotic-induced *Clostridium difficile* infection,” which was previously published in PLOS Computational Biology [9]. The content of chapter 4 is adapted from the article “Steady-state reduction of generalized Lotka-Volterra systems in the microbiome,” which was previously published in Physical Review E [10]. The content of chapter 5 is adapted from the article “Efficient navigation and control of outcomes in generalized Lotka-Volterra systems in the microbiome,” which is under review for a Special Issue on Biological Systems Modeling as part of the AIMS Applied Mathematics Book Series [11]. Parker Shankin-Clarke, an undergraduate research mentee of mine, contributed to this publication and is second author. The content of chapter 6 is adapted from the article “Control of ecological outcomes through deliberate parameter changes in a model of the gut microbiome,” which was previously published in Physical Review E [12]. Zipeng Wang, an undergraduate research mentee of mine, was first author of this publication. If the reader is interested in the supplementary information for any article, we refer them to the original publication.

Chapter 2

Microbiome composition affects host physiology in the fruit fly

We begin with an explicit demonstration of how microbiome composition affects host physiology, carried out in the fruit fly with experimental collaborator Will Ludington. Bacteriotherapies, which seek to improve host health by modifying microbiome composition, rely on this fundamental connection. The fruit fly as an experimental model system is extremely manipulable and replicable: hundreds of flies can be precisely colonized with the same set of bacteria, and their behaviors then measured. The resulting statistics reveal quantitative and significant relationships between microbiome composition and host traits.

In addition to its experimental accessibility, the fruit fly is also a natural model system because its core microbiome naturally contains only a few microbial species (~ 10 compared to ~ 1000 in humans). Thus, the lab flies associated with between one and five bacterial species have microbiomes that are comparable with the microbiomes of wild fruit flies. In this chapter we demonstrate that the traits of flies with more complex microbiomes can be predicted by averaging together the traits of flies with simpler consti-

tutive microbiomes. This empirical evidence, showing that the fruit fly microbiome can be coarse-grained, is welcome news to those that wish to model the much more complex human microbiome.

The fruit fly is an excellent experimental vessel in which to study direct and indirect bacteriotherapies. Direct bacteriotherapies can be implemented by modifying the bacterial slurry that the flies eat, while indirect bacteriotherapies can be implemented by altering their microbiome environment (e.g. by changing the food's acidity or composition, or adding antibiotics). The success of these interventions can be determined by measuring the resulting microbiome composition of the fruit flies. By characterizing the microbial dynamics of the fruit fly microbiome, these experiments could lead to insights about how it can be reliably controlled.

2.1 Abstract of “Microbiome interactions shape host fitness”

Gut bacteria can affect key aspects of host fitness, such as development, fecundity, and lifespan, while the host in turn shapes the gut microbiome. However, it is unclear to what extent individual species versus community interactions within the microbiome are linked to host fitness. Here we combinatorially dissect the natural microbiome of *Drosophila melanogaster* and reveal that interactions between bacteria shape host fitness through life history tradeoffs. Empirically, we made germ-free flies colonized with each possible combination of the five core species of fly gut bacteria. We measured the resulting bacterial community abundances and fly fitness traits including development, reproduction, and lifespan. The fly gut promoted bacterial diversity, which in turn accelerated development, reproduction, and aging: flies that reproduced more died sooner.

From these measurements we calculated the impact of bacterial interactions on fly fitness by adapting the mathematics of genetic epistasis to the microbiome. Development and fecundity converged with higher diversity, suggesting minimal dependence on interactions. However, host lifespan and microbiome abundances were highly dependent on interactions between bacterial species. Higher-order interactions (involving 3, 4, and 5 species) occurred in 13% to 44% of possible cases depending on the trait, with the same interactions affecting multiple traits, a reflection of the life history tradeoff. Overall, we found these interactions were frequently context dependent and often had the same magnitude as individual species themselves, indicating that the interactions can be as important as the individual species in gut microbiomes.

2.2 Significance

All animals have associated microbial communities called microbiomes that can influence the physiology and fitness of their host. It is unclear to what extent individual microbial species versus interactions between them influence fitness of the host. Here we mapped all possible interactions between individual species of bacteria in *Drosophila melanogaster* fruit flies with host fitness traits. Our approach revealed that the same bacterial interactions that shape microbiome abundances also shape host fitness traits. The fitness traits of lifespan and fecundity showed a life history tradeoff, where equal total fitness can be gotten by either high fecundity over a short life or low fecundity over a long life. The microbiome interactions are as important as the individual species in shaping these fundamental aspects of fly physiology.

2.3 Introduction

In 1927, Steinfeld reported that germ-free flies live longer than their microbially colonized counterparts, suggesting that bacteria hinder host fitness. This observation — that the microbiome can impact aging — has been replicated in flies and vertebrates [13, 14]. However, a decrease in lifespan does not necessarily indicate a negative impact on the host. Organisms in their environment are selected for their fitness, which is a function of lifespan, fecundity, and development time [15]. Life history tradeoffs can, for instance, increase fecundity at the expense of lifespan [16, 17, 18] providing different strategies for equal fitness. These observations set up two major questions: what is the role of an individual bacterial species versus interactions between them in determining host lifespan, and how is the microbiome effect on lifespan related to overall host fitness?

Identifying the host effects of specific bacteria has been difficult, in part due to high gut diversity but also because interactions between bacteria can depend on context [19]. Non-additive effects of more than two variables are called higher-order interactions, and they indicate that interactions depend on context. For example, a bacterium may produce a specific B-vitamin in response to its neighbors [20, 21]. This response may impact the host, and host feedbacks can mitigate or exacerbate changes in the microbial community [22]. However, specific examples may be misleading, as the true complexity of a gut microbiome has never been exhaustively quantified. Thus, it remains an outstanding challenge to reverse engineer the interaction networks that characterize microbiome-host effects relative to host interactions with individual bacterial species. Doing so would allow us to address the role of microbial community complexity in shaping host fitness. However, quantifying the set of all possible interactions of n species is a combinatorial problem involving 2^n distinct bacterial communities. As n approaches the diversity of the mammalian gut with hundreds of species, this challenge becomes experimentally

unfeasible.

The gut microbiome of the fruit fly *Drosophila melanogaster* is an effective combinatorial model because as few as five species of bacteria consistently inhabit the gut of wild and laboratory flies [23, 24, 25], yielding 2^5 possible combinations of species. Because early work on the fruit fly microbiome suggested that it is a transient community consisting only of recently ingested bacteria [26], we set up our experiments to maintain bacterial colonization through frequent ingestion. However, newer studies demonstrate that a modified fly diet as well as specific bacterial strains make for a persistent gut microbiome [27, 28, 27], suggesting similarities with higher organisms. Here, we isolated the five core laboratory fly gut bacteria species in culture, *Lactobacillus plantarum* (*Lp*), *L. brevis* (*Lb*), *Acetobacter pasteurianus* (*Ap*), *A. tropicalis* (*At*), and *A. orientalis* (*Ao*). These fermentative lactic acid bacteria and acetic acid bacteria commonly occur in the wild fly gut [25, 29, 30], where they can maintain a stable association [28, 27]. We constructed germ-free flies by bleaching the embryos, and reinoculated the newly emerged adult flies via continuous association with defined flora using established protocols [27, 31]. We made the 32 possible combinations of the five bacterial species and then quantified the microbiome composition and resultant host phenotypes of (i) development time, (ii) reproduction, and (iii) lifespan to determine the relationship between gut microbe interactions and host fitness. We tested to what extent the presence and abundance of individual bacterial species account for the fly physiology phenotypes we measured.

2.4 Results

2.4.1 Microbiome diversity confers a life history tradeoff

We hypothesized that microbiome-induced lifespan changes might be due to changes in life history strategy, such as a tradeoff with fecundity. We therefore set up an experiment to measure how defined species compositions change each of the host fitness traits of lifespan, fecundity, and development time, which have been found to co-vary in life history tradeoffs [16, 17]. We measured these traits concomitantly in the same experiment so that we could sum them together to calculate overall fly fitness (Fig. 2.1A).

We first isolated each of the five species of bacteria found in our laboratory flies: *Lactobacillus plantarum* (*Lp*), *L. brevis* (*Lb*), *Acetobacter pasteurianus* (*Ap*), *A. tropicalis* (*At*), and *A. orientalis* (*Ao*). In order to test whether groups of bacteria have additive effects, we made each of the 32 possible combinations of the 5 species (including germ-free; Fig. 2.1A). We then made germ free flies and inoculated them with defined bacteria compositions at 5-7 days post-eclosure to reduce variation in development and gut maturation [32].

We performed five technical replicates of each experiment with 10 males and 10 females together in the same vial. The five replicates were performed over two separate biological replicates for a total of 100 adult flies per each of the 32 treatments. We transferred the flies every three days to fresh food that was inoculated with fresh bacteria in order to reduce the effects of bacterial growth on the food. To measure lifespan, we recorded the number of live flies daily. To measure fecundity, we kept the old vials that flies were transferred from and counted the number of emerged live adults. To measure development time in the population experiments where egg laying took place for three days, we counted the number of days for the first adult to emerge from a pupal case.

We first asked the role of individual bacterial species on fly lifespan. Consistent

with previous studies, our germ-free flies survived the longest (Fig. 2.1B; Table S1). However, only *Lp*, *At*, and *Ao* shortened lifespan, while flies aged with *Lb* and *Ap* had equivalent survival to germ free flies. We next asked the effect of microbial diversity on fly lifespan. Germ-free flies survived ~20% longer than flies colonized by all five bacteria (mean lifespan \pm standard error of the mean, 53.5 ± 1.5 germ-free vs. 43.5 ± 1.1 for 5-species gnotobiotics). Overall, we found a decrease in survival over many bacterial associations as we increased gut diversity (Fig. 2.1B inset, S1, S2; $r=-0.54$, $p=0.002$, $n=32$, Spearman correlation), consistent with the gut microbiome having a pathogenic effect on the host.

We next asked whether the reduction in lifespan was offset by a life history tradeoff in fecundity. Decreased lifespan corresponded to an increase in fecundity for female flies (total daily fecundity vs. lifespan: $r=-0.50$, $p=0.003$, $n=32$, Spearman correlation; Fig. 2.1C; Table S1) and is not explained by differences in fly activity (Fig. S3). Such life history tradeoffs are well-documented in the literature and are believed to constitute a differential allocation of resources between long-term body maintenance and reproduction [15, 33].

A true tradeoff is one that allows an individual organism to adapt its life history strategy with equal overall fitness. Fitness is a function of fecundity, development, and lifespan, which gives an estimate of the maximum rate of population growth. We wondered whether the observed differences in lifespan were balanced by differential rates of fecundity and development or whether these differences in fly physiology actually made flies with distinct microbiome compositions more and less fit. To address this question, we combined our data for development, fecundity, and lifespan (Table S1) in a Leslie matrix [34], a classical model of discrete population growth, to calculate organismal fitness under each bacterial association. Overall, fitness was constant across many distinct bacterial associations (Fig. 2.1D). Thus, the changes in lifespan we observed are consis-

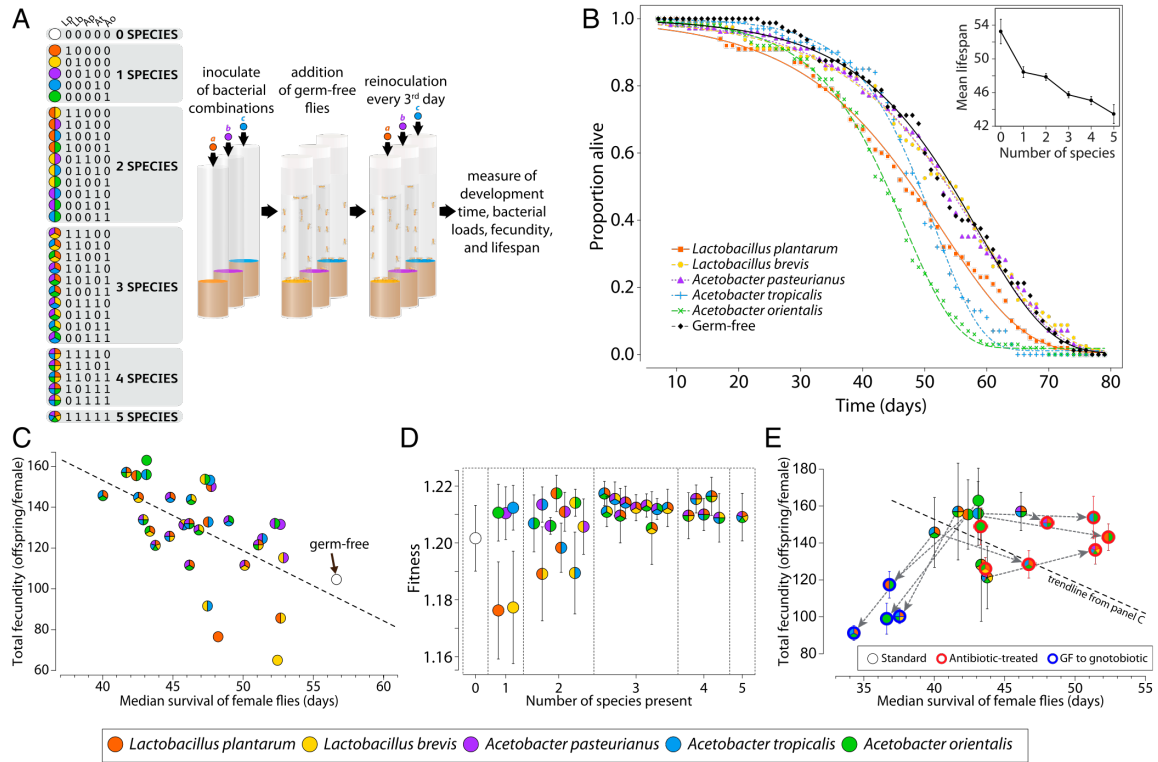


Figure 2.1: The microbiome induces a life history tradeoff between lifespan and reproduction. (A) Experimental design. The multi-color pies indicate which species are present in a given combination along with the corresponding binary code. Each species abbreviation (Lp, Lb, Ap, At, Ao) is indicated above its corresponding locus in the binary string. Both notations, colored pies and binary are used consistently throughout the chapter. The color code is included redundantly in the figures to aid the reader. (B) Single bacterial associations decrease fly lifespan. (B inset) Microbiome diversity decreases fly lifespan. Error bars S.E.M. (C) In agreement with prior reports, higher total fecundity is associated with shorter lifespan. This tradeoff is apparent for average daily fecundity as well as total fecundity per female. See Table S1 for S.E.M. (D) Fitness calculations using a Leslie matrix reveals roughly constant fitness across different microbiomes. Error bars are standard error of the estimate. (E) The lifespan-fecundity tradeoff can be broken by putting flies on antibiotics after their peak reproduction (red circles = gnotobiotic flies treated with antibiotics; see Methods) after 21 days, which encompasses the natural peak fecundity (Fig. S4). Note the shifts in lifespan between the regular treatment, the antibiotic treatment, and the late-life bacterial inoculation treatment. Lifespan was significantly extended, whereas total fecundity stayed high. Shifting germ-free flies to gnotobiotic treatment after 21 days post-eclosion decreased lifespan without increasing reproduction (blue circles = germ-free flies made gnotobiotic 21 days post-eclosion). $n=100$ flies per treatment for the standard and antibiotic-treated experiments $n=60$ flies per treatment for the germ-free switched to gnotobiotic experiment. Error bars show S.E.M.

tent with a differential allocation of resources to reproduction, a true life history tradeoff, meaning that microbiome composition sets a fly's life history strategy to maximize either reproduction or longevity. In a co-submitted manuscript, Walters *et al.* (2018) show evidence that such microbiome-based fitness tradeoffs also occur in wild flies.

2.4.2 Reproduction cannot be increased by mid-life microbiome addition

The life history tradeoff suggests that a fly born into stark conditions in the wild could maximize its fitness by first acquiring a longevity-promoting microbiome and then converting to a fecundity-promoting one when environmental conditions improve. Female flies are primarily reproductive in the first part of their life, with a gradual decay in fecundity approaching middle age (Fig. S4). To test whether individual flies can switch life history strategy to match their microbiome, we aged germ-free flies for 21 days (roughly middle age) and then associated these flies with fecundity-promoting bacteria. There was no significant increase in total fecundity for these flies and a significant decrease in lifespan compared to germ-free flies (Fig. 2.1E; $p=0.054$ for fecundity, $n=275$ flies pooled across four bacterial combinations, two sample one-sided t-test, $p>0.05$ for all pairwise combinations after Tukey's multiple comparison correction of two-sample one-sided t-tests; $p=2\times 10^{-7}$ for lifespan, $n=400$ flies pooled across four bacterial combinations, two-sample one-sided t-test; $p<0.001$ for 4 of 4 combinations after Tukey's correction for multiple pairwise comparisons, $n=100$ flies per combination, two-sample one-sided t-tests; see Fig. 2.1E for specific bacterial combinations). These results are consistent with the simple hypothesis that a fly's reproductive window cannot be extended by late-life improvement in nutrition.

2.4.3 Microbiome interactions can change host physiology

We hypothesized that the microbiome may shorten lifespan through a process independent of reproduction. To examine this hypothesis, we used antibiotics to remove the microbiome of high-fecundity female flies and measured the resulting change in lifespan. We first allowed female flies with high-fecundity microbiomes to reproduce for 21 days (to a level greater than the total lifetime fecundity of germ-free flies; Fig. S4), and we subsequently eliminated the microbiome using an antibiotic cocktail (ampicillin, tetracycline, rifamycin and streptomycin). In general, the midlife elimination of gut flora lengthened the female fly lifespan by roughly 15% compared to flies continuously fed live bacteria (Fig. 2.1E; $p=9 \times 10^{-7}$, $n=560$ flies pooled across bacterial combinations; $p < 0.05$ for 4 of 7 combinations after Tukey's correction for multiple pairwise comparisons, $n=80$ flies per combination, two-sample one-sided t-test). Total fecundity decreased slightly (Fig. 2.1E; $p=0.01$, $n=560$ flies pooled across bacterial combinations; $p > 0.05$ for all 7 combinations after Tukey's correction for multiple pairwise comparisons, $n=80$ flies per combination, two-sample one-sided t-test). This result demonstrates that the life history tradeoff is not necessarily fixed and suggests that fly lifespan is shortened by some aspect of the bacteria rather than by reproduction. However, two specific bacterial combinations yielded no increase in lifespan when removed from their host by antibiotics: *Ao* and *Lp+Lb+Ao*, suggesting a memory in host physiology induced by these two combinations. Interestingly, the intermediate microbiome composition, *Lp+Ao*, did not show this memory, nor did the similar composition *Lp+At+Ao* (Fig. 2.1E) (with antibiotic elimination of the microbiota extending lifespan) suggesting specificity of the microbiome composition in this physiological memory. These experiments demonstrate that interactions between bacteria can significantly impact the host's ability to adjust its physiology.

2.4.4 Higher-order microbiome interactions change host physiology

We next calculated to what extent microbiome interactions change fly physiology. We applied a multivariate linear regression model, a common statistical test for interactions between experimental variables [35]. Here our variables are the 5 bacterial species. We detected evidence of widespread pairwise interactions in the data (Tables S2-5). Higher-order interactions are non-additive effects of more than two variables, which would complicate efforts to predict physiology of hosts with high diversity microbiomes. We checked for 3-way, 4-way, and 5-way interactions using the same statistical approach. Higher-order interactions for lifespan, fecundity, development, and bacterial composition were evident at each level of diversity (Tables S2-5), indicating that species interactions rather than just their direct effects change host phenotypes. Many of these interactions have equivalent magnitude to the impacts of individual species. For instance, the average lifespan of germ-free flies is 53 days (Fig. 2.1B). Individually, *Ao* can shorten lifespan by 10 days. Pairwise interactions can change mean lifespan by 8 days (Table S2). Likewise, flies colonized by all five species of bacteria survive an average of 43 days. Microbiome interactions account for a 13 day (28%) increase in lifespan over the additive prediction (Table S2). Overall, these findings demonstrate that microbiome interactions can have major impacts on host physiology.

To confirm these interactions, we next asked whether the physiology of flies with more than one bacterial species could be predicted by simply averaging the phenotypes of flies with the corresponding single species associations (Fig. 2.2; e.g. for combination *Lp-Lb-Ap*, phenotypes of flies living with either *Lp*, *Lb*, or *Ap* were averaged; Math Supplement Section 10). This model showed minor predictive power for development time (27% total: 3/10 2-way; 3/10 3-way; 0/5 4-way; 1/1 5-way interactions predicted), better prediction

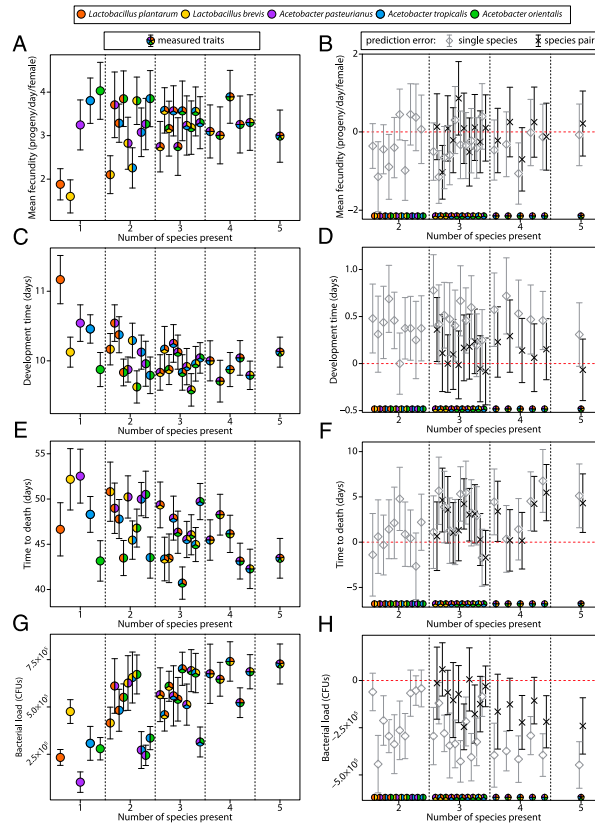


Figure 2.2: **Microbiome interactions impact host lifespan and bacterial load.**

Mean fecundity per female per day was measured concomitantly with development time and adult survival over the flies' lifespans. (A) Variation in fecundity decreases as gut diversity increases. Median $n=65$ vials measured per bacterial treatment. (B) As described in Math Supplement Section 10, daily fecundity in multispecies bacterial combinations can be predicted by averaging either the corresponding phenotypes of the single species associations or the corresponding phenotypes of the pairwise species associations. Error in the predictions (averaging prediction minus measured trait value) is displayed. Single species averaging predictions in gray. Species pair averaging predictions in black. Error bars are 95% confidence intervals. See Math Supplement Section 10. (C) The number of days to adulthood was measured as the first pupa to emerge from an individual fly vial during the lifespan experiment. Median $n=24$ per bacterial treatment (Fig. S2). (D) Averaging models as in B applied to development data. (E) Lifespan decreases as gut diversity increases. Median $n=100$ flies per bacterial treatment. (F) Averaging models as in B applied to lifespan data. (G) Mean bacterial load averaged over 48 replicates per combination. (H) Averaging models as in B applied to bacterial load. Error bars for all plots are 95% confidence intervals. Colored pies on the x-axis of B, D, F, H indicated bacteria combinations and are consistently ordered with panels A, C, E, G.

of lifespan (65% total: 9/10 2-way; 6/10 3-way; 2/5 4-way; 0/1 5-way phenotypes predicted), and reasonably accurate prediction of average daily fecundity (81% total: 7/10 2-way; 9/10 3-way; 4/5 4-way; 1/1 5-way phenotypes predicted). We also measured total bacterial abundances in the flies (Table S1), which had little predictive power in the simple averaging model (20% total: 4/10 2-way; 1/10 3-way; 0/4 4-way; 0/1 5-way phenotypes predicted). We next asked whether averaging data from the corresponding species pairs could predict phenotypes of 3-, 4- and 5-way combinations (Fig. 2.2; e.g. for combination *Lp-Lb-Ap*, phenotypes of flies living with either *Lp-Lb*, *Lp-Ap*, or *Lb-Ap* were averaged; Math Supplement Section 10). This model overall correctly predicted 78% of the fly traits (Fig. 2.2; 15/16 development; 9/16 lifespan; 15/16 fecundity; 11/16 bacterial load), indicating that pairwise interactions account for a majority of host variation. However, in the remaining cases (up to 44% unpredicted for lifespan), failure of this simple model indicates higher-order interactions within the gut microbiome. Taking these analyses together with the life history tradeoff, microbiome-host interactions (including those of higher order) can significantly impact fly fitness traits. In a later section, we analyze these interactions comprehensively to show how the context of bystander species influences these interactions.

2.4.5 Bacterial presence-absence more than abundance impacts fly physiology

The differences in host physiology we observed resulting from different microbiome compositions could be due not only to which species are present but also to their abundances. We reasoned that if a particular bacterial species drives a host physiological trait, then its abundance should be correlated with that trait. We therefore measured the abundances of individual bacterial species in the flies in order to determine the re-

relationship to different fly physiologies. We first prepared gnotobiotic flies as before by inoculating 5-7 day old mated germ free flies with defined bacterial compositions. Flies were transferred to fresh food inoculated with fresh bacteria every third day for a total duration of 10 days before they were washed in 70% ethanol, crushed, plated, and colony-forming units (CFUs) enumerated (Fig. 2.3A). The experiments were performed in two biological replicates for a total of 12 female and 12 male flies that were analyzed for each of the 32 bacterial combinations (Fig. 2.3B). The total bacterial load was higher when more species were present ($r=0.63$, $p=0.0001$, $n=31$ bacterial combinations, Pearson correlation). However, on a species-by-species basis, abundance stayed constant or decreased as species diversity increased (Fig. 2.3C; *Lp*: $r=-0.07$, $p=0.8$; *Lb*: $r=-0.37$, $p=0.2$; *Ap*: $r=-0.50$, $p=0.06$; *At*: $r=-0.59$, $p=0.02$; *Ao*: $r=-0.55$, $p=0.03$; Spearman correlations), suggesting competition plays a role in the interactions. To quantify the robustness of bacterial association in our experiments, we prepared a parallel experiment with the only difference being that after the initial 10 days of inoculation, flies were transferred daily to fresh, germ free food for five subsequent days before CFUs enumeration as before. Only very minor differences occurred between the two experiments (Fig. S5), with the flies transferred daily to germ free food for 5 days surprisingly having slightly higher CFU counts than flies plated directly after day 10 of inoculation (Wilcoxon rank sum test, median CFUs for flies transferred to germ free food: $10^{5.65}$ CFUs vs flies directly plated after day 10 of inoculation: $10^{5.59}$ CFUs, $p=0.01$, $N=1536$ individual flies). Because only minor differences were observed, we merged the two experiments to increase statistical power. Median total bacterial load ranged from 49,000 CFUs per fly for *Ap* alone to 737,000 CFUs per fly for *Lb+At+Ao*, with an overall median of 425,000 CFUs per fly (Fig. 2.3B). The robust colonization observed despite daily transfer to germ free food indicates the gut microbiome is persistent under these conditions, which is in contrast to some previous reports [26, 36]. Two variables that could account for this difference

are different bacterial strains [28, 27] and our use of a fly food with minimal microbial growth inhibitors [37].

To test whether bacterial abundances drive fly physiology, we next compared the individual species abundances and total bacterial abundances in adult flies with the fly physiology phenotypes (Fig. 2.2A,C,E,G, S2). We first calculated the correlation between individual species abundances and each host physiology trait (Fig. S6). Out of 20 possible correlations, two significant correlations were found, (i) between *Lp* abundance and total female fecundity (Fig. 2.3D; $r=0.52$, $p=0.04$, $n=16$) and (ii) between *Ao* abundance and decreased lifespan (Fig. 2.3E; $r=-0.53$, $p=0.03$, $n=16$), indicating that these two individual species can explain 27% and 28% of the variation in fecundity and lifespan respectively. We did not detect other significant relationships between bacterial load and host physiology, leaving the remaining variation (73% of fecundity and 72% of lifespan) unexplained by individual species abundances. However, as we showed in Fig. 2.1E, the interaction between *Ao* and *Lp* can dramatically alter the fly's ability to adjust its physiology when treated with antibiotics, with a 21% change in lifespan (Fig. 2.1E). Thus, individual bacterial species loads are not necessarily expected to determine impacts on the host.

As a secondary test that bacterial abundance drives fly phenotypes, we examined both load and phenotype variation. If the load of individual bacterial species drives host physiology traits, we would expect that higher variation in bacterial load would correspond to higher variation in host traits, yielding a positive correlation. When we calculated the relationship between bacterial load variation and host trait variation, we found no statistical evidence for an association between host bacterial load and host physiology traits (Fig. S7). Taken together these results suggest that the long term presence of bacterial species is more indicative of their effect on host physiology than their abundances.

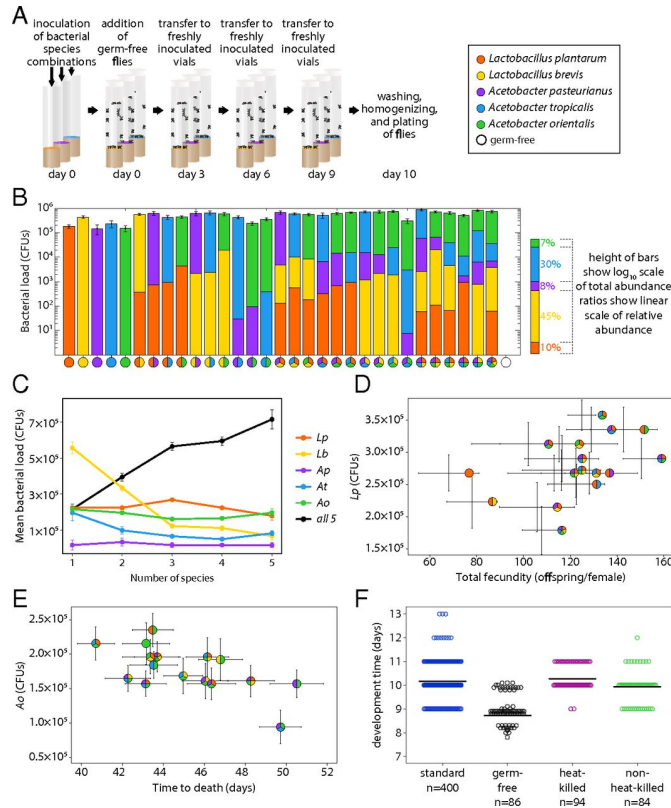


Figure 2.3: Microbiome abundances correlate with some host physiology traits. (A) Gnotobiotic flies were associated with defined bacterial flora for 10 days before washing, crushing, and CFU enumeration. (B) Mean microbiome load (\log_{10} scale) and relative abundances of the different species (linear scale) for all 32 possible combinations of the five species. $N=24$ replicate flies from 2 independent biological replicates were measured per combination. (C) Total bacterial load increases as the number of species increases but *Lb* abundance drops. Mean abundances were calculated from B as a function of the number of species present (see Fig S5 for complete data). Black line indicates mean total bacterial load per fly. (D) *Lp* abundance (from B) correlates with increased female fly fecundity (from Fig. 2.1C). (E) *Ao* abundance (from B) correlates with decreased fly lifespan (from Fig. 2.1C). (F) Development time from embryo to adult is accelerated by live bacteria. Development assay from Fig. 2.2B was repeated with variation in food preparation and source of embryos. ‘Standard’: data from Fig. 2.2B fitness experiment, ‘germ-free’: embryos from germ-free females placed directly on fresh food inoculated with defined bacteria; ‘heat-killed’ and ‘non-heat-killed’: vials from fitness experiment cleared of flies and either seeded directly with germ-free embryos (non-heat-killed) or placed at 60°C for 1 hour and checked for sterility (heat-killed) before being seeded with germ-free embryos. The number below the x axis indicates number of replicate vials assessed. See Fig. S9 for complete bacterial combinations and individual replicates of F. All error bars S.E.M.

2.4.6 Live bacteria speed up fly development

We did find one notable exception to the presence-absence rule: *Ao* abundance in the food sped up larval development time significantly ($r=-0.95$, $p=0.003$, $n=7$ bacterial combinations with *Ao*, Spearman correlation; Fig. 2.2C, S8). However, consistent with the adult results, there was no correlation with fly physiology for the four other species. We next tested whether there was a maternal effect on development time by removing the maternal bacterial association. We prepared the vials for this development experiment by first setting up a replicate fitness experiment (as in Fig. 2.1A). After the first transfer to fresh vials, we took the used vials, allowed all larvae to form pupae, and then removed the pupae. Flies developing in these vials had an equal rate of development to the fitness experiment (Fig. 2.3F; paired sample t-test, $p>0.18$, $n=500$), indicating no maternal effect. We then tested whether live bacteria aid the flies under these conditions. We performed a duplicate experiment but heat-killed the vials in a humidified (to prevent drying) 60° C chamber for 1 hour (and tested for sterility). All the sterile vials were inoculated with ~30 germ-free embryos each. Flies in heat-killed vials developed ~8 hours more slowly (Fig. 2.3F; $p<0.005$, $n=16$, paired sample t-test), suggesting that active bacterial metabolism [38] speeds up fly development (Fig. 2.3F, S9). Finally, we asked whether bacteria degrade the fly food. We harvested eggs from germ-free flies, associated them with all 32 bacterial combinations on fresh food, and measured development times. This experiment uncovered a significant acceleration in development time compared with the fitness experiment (Fig. 2.1A). These experiments demonstrate context-dependence in terms of timing in microbiome experiments.

2.4.7 Microbial abundance interactions correlate with host physiology interactions

Are interactions between bacteria linked with host physiology? We first tested whether microbial interactions detected through CFU counts were correlated with interactions detected through development, fecundity and lifespan (Fig. S13). Focusing on the statistically significant interactions, there is a strong correlation between the interaction strengths across these distinct phenotypes (Fig. S13D), indicating that the same microbiome abundance interactions also are associated with fly physiology interactions and the life history tradeoff. Thus, interactions calculated from bacterial species abundances may be predictive of fly traits. This relationship is notably in contrast to the relationship between the individual bacterial species abundances and fly physiology phenotypes, where only two weak correlations were established (Figs. 2.3D,E, and Fig. S6-7).

We next asked how the pairwise interactions between individual bacteria species change under increasing numbers of species present. We used the abundance data for individual species (Fig. 2.3) to calculate the pairwise interaction strengths between the 5 species.

We first calculated the pairwise correlations in species abundances as a function of the total number of species present in the gut (Fig. 2.4A), correlations became more negative for individual species pairs as diversity increased ($p=0.03$, $n=10$ species pairs, Kendall's Tau and Wilcoxon signed rank; see Math Supplement, Sections 11.3-11.4), consistent with stronger competition at higher diversity.

We then calculated the directional interactions (i.e. $A \rightarrow B$ vs. $B \rightarrow A$) using Paine's classic approach [39] where interaction strength is based on the change in abundance of one species when a second species is removed (Fig. 2.4B,C, S14A,B; see Math Supplement, Section 10.1). Comparing the pairwise interaction maps at high and low

diversity, we found that interactions are generally positive when only two species are present, consistent with interactions between two species *in vitro* (Fig. S14C). However, interactions become more negative at higher diversity, consistent with increasing competition. An alternate approach to calculate the interactions by fitting the classic generalized Lotka-Volterra model (see Math Supplement Section 10.2) gave qualitatively similar results (Fig. S14D,E). However, parametrizing the model on low diversity data did a poor job of estimating the bacterial abundances at higher diversity with $n \geq 3$ species ($p=0.8$, binomial test, $N=16$; see Math Supplement Section 11.3), in agreement with the changing interaction landscape at higher diversity (Fig. 2.4, S14).

Lastly, we asked if the interaction networks we calculated are consistent with the maintenance of diversity we observe. We calculated the asymmetry in the interaction network using the approach of Bascompte *et al.* [40], where asymmetry of interactions is indexed from 0 (perfectly symmetric) to 2 (exactly opposite). For the low diversity case the mean asymmetry is 1.04 (SD = 0.13), and for the high diversity case the mean asymmetry is 0.77 (SD = 0.08) (see Math Supplement Section 11.1), indicating significant asymmetry. Furthermore, analysis of the variation in total bacterial load between individual flies showed a decreased coefficient of variation for high diversity ($p=0.02$, Wald test; Fig. S15, S16; see Math Supplement Section 10.5). Together with the strength of interactions (Fig. 2.4), these calculations are consistent with community stability at higher diversity [41].

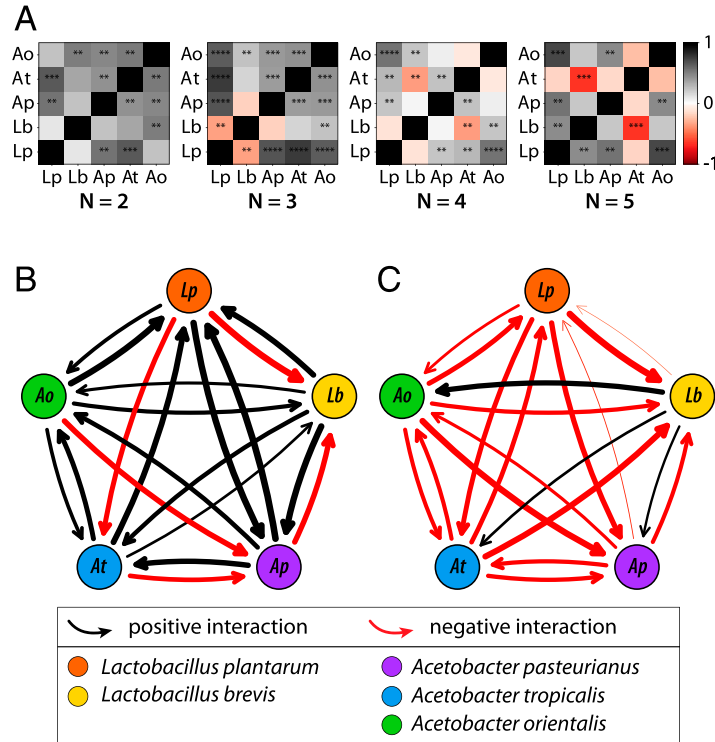


Figure 2.4: **Microbiome interactions stabilize diversity in the fly gut.** (A) Pairwise correlations in abundance for the five species of bacteria in fly guts with totals of two, three, four, and five species present. More positive correlations are apparent at low diversity, whereas more negative correlations occur as diversity increases ($p=0.03$; see Math Supplement Section 10.4). Direct calculation of interaction strength [39] at low (B, one to two species) and high (C, four to five species) diversity based on CFU abundance data (see Figs. 2.3B, S5) revealed asymmetric interactions that decrease in strength at higher diversity (see Math Supplement Section 11.1; Fig. S14). Consistent with the correlations in (A), more negative interactions occur in more diverse guts.

2.5 Discussion

2.5.1 Bacterial abundance interactions may damage the host.

Taking the epistatic interactions and the directional interactions together, we found that the biological interactions determining the bacterial community in the fly gut involve more than just pairs of species [42]. Considering just the bacterial abundance interactions, these generally become weaker and more negative as diversity increases, which is consistent with community stability through competition. Time series experiments could be a better way to evaluate community stability [43]. We also acknowledge that this 5 species gut community may have been selected for its stability. However, host immunity is another potential factor shaping the detected interactions [22]. The negative interactions we detect in the microbiome are associated with shorter lifespans in the host flies, suggesting that negative bacterial interactions may damage the host. Consistent with this finding, microbiome removal by antibiotic treatment (Fig. 2.1E) extended lifespan. We speculate that molecular mechanisms for microbial damage to the host could include nutrient depletion (Fig. S9), toxic secondary metabolite production, and physical injury through e.g. bacterial secretion systems, which have been shown to trigger fly mortality during to bacterial interactions [44].

2.5.2 How much do higher-order interactions matter?

While we found that higher-order interactions occur and are responsible for significant changes in fly physiology, we also found that lower order interactions between bacterial combinations can predict many phenotypes in 3-, 4-, and 5-way bacterial combinations. And simply averaging lower order traits can predict >50% of higher diversity traits, a substantial decrease in microbiome uncertainty. The convergence of traits such as fe-

cundity under increasing diversity indicates we may harness good predictive power from low diversity microbiomes. However, the lack of convergence in traits such as lifespan and bacterial abundances suggests different rules may apply to different phenotypes, and it is unclear what rules will apply to more diverse host-microbiome systems. Different conditions, such as diet composition could may drastically change the microbiome interactions. Decomposing interactions in increasingly diverse systems remains an important goal for future studies.

2.5.3 Microbiome interactions mediate a life history tradeoff between lifespan and fecundity

Overall, we found that interactions in the fruit fly gut microbiome structure both the fitness of the fly and the composition of the microbiome. The magnitudes of these interactions are equivalent to the effects of individual species. Thus, microbiome interactions (and not just individual species) can be a major driver of host physiology. Many studies have documented changes in fly lifespan as a function of various factors including diet, host genetics, and microbiome composition [13, 45, 46, 47]. Our study suggests that microbiome composition and the timing of association with the microbiome can have major impacts on lifespan as well as life history tradeoffs [33]. In a paper submitted concurrently with this one, Walters *et al.* show the consequences of this tradeoff for ecology and evolution of wild flies [48].

2.5.4 The *Drosophila* gut microbiome serves as an effective model of microbiome complexity

A pervasive challenge in host-microbiome science is the complexity of most host-associated microbiomes. *Drosophila melanogaster* has a naturally low diversity micro-

biome, which facilitates the study of complexity. Regarding suitability of this model, a major question is whether such a simple system with just five species can recapitulate the complex phenotypes associated with higher diversity microbiomes such as humans and plants. The fact that we observe emergent properties in this simple and tractable five species community makes it an attractive model for studies of microbiome complexity. Based on our empirical results, we argue that interacting groups rather than just individual species may be fundamental building blocks of microbiome-host relationships.

Chapter 3

Implementing direct bacteriotherapies in an ecological model of the microbiome

In this chapter we lay out the mathematical foundation, based on a generalized Lotka-Volterra (gLV) model, with which we will study bacteriotherapies. This chapter, developed before the creation of the steady-state reduction technique, primarily uses numerical approaches to investigate a gLV model fit by Stein *et al.* to microbial abundance time-series data collected by Buffie *et al.* in a mouse experiment of *C. difficile* infection (CDI). This experimentally-fit CDI model serves as a clinically-relevant point of departure for our analyses: in this chapter we map the steady states of this model that are reachable from experimentally-measured initial conditions, we characterize the resilience of these steady states to *C. difficile* (CD), and we show how potentially relevant behaviors of CD—such as sporulation or the development of antibiotic-resistance— may be incorporated in a gLV model.

Most importantly, this chapter formalizes the implementation of direct bacteriother-

apies in gLV models. We deal specifically with fecal microbiota transplantation (FMT) in which the transplants are composed of experimentally-measured mouse microbiome compositions. The major result is that FMT is capable of altering the steady state outcome of a CDI-prone initial condition, so that it attains a health steady state instead of a CD-infected steady state. However, the efficacy of FMT is highly variable and dependent on the transplant composition (some transplants drive the system *towards* the CDI state) and timing (transplants are more effective when administered sooner than later).

At a fundamental level, these questions about FMT efficacy correspond to questions about the dynamical landscape of high-dimensional gLV systems. The successful administration of FMT indicates that the intervention switched which basin of attraction the system was in. The state of the system changes over time according to the system's ecological dynamics, and these dynamics give rise to the time-dependent efficacy of direct bacteriotherapies. Direct bacteriotherapies endow the rich mathematical behaviors of gLV systems with clinically-relevant motivations, and the initial analyses performed in this chapter motivate our subsequent analytical work in chapters 4, 5, and 6.

3.1 Abstract of “*In silico* analysis of antibiotic-induced *C. difficile* infection”

In this chapter we study antibiotic-induced *C. difficile* infection (CDI), caused by the toxin-producing *C. difficile* (CD), and implement clinically-inspired simulated treatments in a computational framework that synthesizes a generalized Lotka-Volterra (gLV) model with SIR modeling techniques. The gLV model uses parameters derived from an experimental mouse model, in which the mice are administered antibiotics and subsequently dosed with CD. We numerically identify which of the experimentally measured initial

conditions are vulnerable to CD colonization, then formalize the notion of *CD susceptibility* analytically. We simulate fecal transplantation, a clinically successful treatment for CDI, and discover that both the transplant timing and transplant donor are relevant to the efficacy of the treatment, a result which has clinical implications. We incorporate two nongeneric yet dangerous attributes of CD into the gLV model, sporulation and antibiotic-resistant mutation, and for each identify relevant SIR techniques that describe the desired attribute. Finally, we rely on the results of our framework to analyze an experimental study of fecal transplants in mice, and are able to explain observed experimental results, validate our simulated results, and suggest model-motivated experiments.

3.2 Significance

The burgeoning integration of big data and medicine is a portent of personalized healthcare. There is a need for accurate, predictive, and mechanistic models that can be relied upon to forecast the course of a disease, test treatments in-silico, and ultimately inform the doctor's prescription. These models, still nascent, are buoyed by rich datasets available due to recent advances in experimental methods (e.g. 16S rRNA high-throughput sequencing); one such model, which we build upon in this chapter, was developed by Stein *et al.* to predict the growth of the infectious *C. difficile* (CD) and 10 other microbial genera. In this chapter we extend the existing model to capture clinical treatments and biologically relevant phenomena. First, we incorporate fecal transplants and identify the mechanism by which they treat *C. difficile* infection (CDI). Then, we develop a methodology that endows a microbe with nongeneric attributes within the existing framework; specifically, we add CD sporulation and the development of antibiotic-resistant strains of CD. By better reflecting the clinically relevant properties of CDI we

can “personalize” a mathematical model to a given disease; this construction of generic yet customizable models will be relevant for personalized healthcare models in years to come.

3.3 Introduction

Microbiota are covertly instrumental in bodily functions including immune response [49] and colonization resistance [50, 51]. Some diseases are associated with an imbalanced microbiome, due to disproportionate regulatory action of the host in response to the microbiome composition [52]. Ironically, another pathway to disease is through antibiotic administration, which can dramatically alter microbial composition and diversity, hinder colonization resistance, and subsequently allow for pathogen infection. Specifically in this chapter, we focus on antibiotic-induced *C. difficile* infection (CDI), a prevalent nosocomial disease [53, 54].

The advent of high-throughput sequencing provides cheap and accurate time-series abundance data of interacting microbial populations, which can then inform dynamic models that extrapolate system behavior [55, 56]. One idealization of interacting species is the generalized Lotka-Volterra (gLV) model, which assumes that the competitive dynamics of a system are entirely captured through pairwise (inter-species) and self (intra-species) interactions [57]. The gLV model ignores explicit external factors like availability of organic compounds, temperature, or location, but it is the most general possible second order differential equation that describes interacting populations, with some reasonable biological constraints.

Approximating microbiome dynamics as a gLV system is a first step towards quantifying the complex interactions between competing microbes. Inarguably this model misses many subtle, non-competition based, interactions: for example, a non-abundant type of

bacteria (e.g. *Escherichia*) may produce proteins vital to general bacterial function (e.g. pili production) [58], but this contribution would not explicitly appear in the model.

In this chapter we simulate the prevalence of *C. difficile* (CD) in the microbiome with a generalized Lotka-Volterra model. The work by Stein *et al.* [7] and Buffie *et al.* [6] serves as a point of departure, from which we develop a framework for evaluating the efficacy of different treatment protocols for CDI. This framework develops causal relationships between simulated therapies and microbiome compositions and also explores how bacterial adaptations such as sporulation and antibiotic-resistant mutation may be added to the gLV model. These clinically motivated approaches explain distinct qualitative aspects of CDI that are otherwise unexplored or inconsistent with previous models.

We begin by discussing the clinical background and existing models of CD infection, including the mathematical model we use in this chapter, and by describing our in-silico implementations of CD treatments. Then we numerically construct phase diagrams that depict the available behaviors of the simulated system, implement in-silico clinical therapies for CDI, and quantitatively track the efficacies of these therapies. Lastly we describe how to include mechanisms for sporulation and mutation in our model, and evaluate their impacts on the efficacy of antibiotic treatment. Through these techniques, we reveal the importance of timing on the efficacy of fecal microbiota transplantation (FMT) and additionally recover the clinical recommendation for pulsed antibiotic administration when treating CD. Finally, we wield this framework to explain experimental FMT outcomes [59], validate simulated results, and propose future experiments.

The era of personalized medicine and prevalence of high-throughput sequencing will demand accurate microbiome models that can predict, diagnose, and recommend treatment for microbiome disease, and the framework developed in this chapter builds upon existing models [60] to progress towards this goal.

3.4 Background

CD is a spore-forming bacterium that can produce toxins which cause CD associated diarrhea, afflicting three million people each year [61]. CDI is especially common in the elderly and in patients who are prescribed antibiotics, since antibiotics deplete the microbiome so that ingested spores of CD— often acquired in healthcare facilities or nursing homes— may invade the vulnerable microbiome [62].

The link between antibiotic treatment, CDI, and microbiome composition was investigated by Buffie *et al.* [6] in a study that gathered mouse time-series phylogenetic data via high-throughput 16S rRNA sequencing. In the study three scenarios were considered, in which the mice were either left alone as a control, exposed to CD, or dosed with the antibiotic clindamycin and subsequently exposed to CD. Each scenario was performed in triplicate and consisted of around 10 time points spanning four weeks, and each time point consisted of thousands of phylogenetic 16S rRNA gene sequences which were mapped to taxonomic species and tallied. The study found that after antibiotic administration of clindamycin the mouse microbiome was less diverse (in terms of the Shannon diversity index) and vulnerable to CDI, which is consistent with clinical observations of humans who develop CDI [61, 62]. Because the anatomies of mice and humans are similar [63] and the microbiomes of both species react to changes in diet in a similar manner [64], it is common to treat the mouse model as a proxy for human CDI.

In a first attempt to model the relationship between CDI and antibiotic treatment, Stein *et al.* [7] proposed a generalized Lotka-Volterra (gLV) model to explain the interactions between different microbes. The parameters for this model were fit with the previously mentioned data from Buffie *et al.* [6]. To reduce dimensionality, Stein *et al.* assumed that bacteria within a given genus behave similarly, and consolidated the species-level data into genus-level data. The parameter fitting procedure was tested on

in-silico data, and the fitted parameters satisfied biologically reasonable restrictions. This model—described in more detail in the text surrounding Eq (3.1)—produces microbiome composition trajectories which allow for simulated antibiotic treatment or exposure to CD. The Spearman rank correlation, a measure comparing the predicted microbe abundances with the experimentally measured abundances, was 0.62 (the largest achievable value is 1), and simulated trajectories for each microbe typically matched experimental trajectories within an order of magnitude. Especially, the model preserved the clinical and experimental conclusion that microbiomes treated with the antibiotic clindamycin were vulnerable to CDI.

In this chapter, we start from a gLV model with previously fitted parameters [7], analyze the steady states, and then build upon this model to explore clinically motivated adaptations. In particular, we focus on simulated remedial treatments that can avoid or reverse *C. difficile* infected steady states, which we interpret as microbiomes suffering CDI.

3.5 Models and Methods

3.5.1 The generalized Lotka-Volterra equations

The generalized Lotka-Volterra equations track the abundance of N populations x_i through time; in our case, the populations are $N - 1$ genera plus the bacterial species CD. They read, for $i \in 1, \dots, N$,

$$\frac{d}{dt}x_i(t) = x_i(t) \left(\mu_i + \sum_{j=1}^N M_{ij}x_j(t) + \varepsilon_i u(t) \right). \quad (3.1)$$

The dynamics of each population are of the same form, so the distinct individual trajectories are entirely determined by the choices of parameters and initial conditions. The

parameters and initial conditions that are used to generate each figure are given in Table A of the Supplementary Information. For a population x_i , μ_i describes that population's self-growth while M_{ij} describes the pairwise effect of population j on population i , an interaction that can be interpreted as mutualistic, commensalistic, or parasitic. Lastly, $\varepsilon_i u(t)$ is an external forcing term, which in our model represents the effect of an administered antibiotic $u(t)$ operating with efficacy ε_i . In all, Eq (3.1) accounts for zeroth, first, and second-order terms, and approximates the competitive dynamics as a power series of the individual populations.

The procedure for parameter fitting is explained in detail and performed by Stein *et al.* [7]. Briefly, the fitted parameter values satisfy $\mu_i > 0$ and $M_{ii} < 0$ for each i , so that in isolation each population will grow and eventually self-limit. Most but not all microbial groups are inhibited by the antibiotic clindamycin. Since the interactions between populations have no clear hierarchy, we interpret the gLV model as microbes on the same trophic level competing for a shared resource—the pairwise interactions, then, effectively describe a food web which we visualize in Fig 5.1. While dynamical systems such as this one may in principle display an array of behaviors, with these fitted parameters we have only observed trajectories that approach biologically reasonable steady states (e.g. no periodic orbits have been observed); if we interpret the negative values in M_{ij} as negative covariances between populations, then this stability is consistent with the covariance effect [65].

3.5.2 Simulation of treatments for *C. difficile* infection

In clinical practice, CDI is defined by the presence of toxigenic CD or of CD toxins in a patient experiencing diarrhea—since there are asymptomatic carriers of CD the mere presence of CD is not sufficient for diagnosis [62]. However, since the model Eq (3.1) does




Figure 3.1: **Pairwise interactions between bacterial populations may be interpreted as a microbial food web.** An arrow from population j to population i represents the effect of j on the growth of i , which we equate to the interaction term M_{ij} in a generalized Lotka-Volterra model, Eq (3.1). The width and opacity of an arrow are proportional to $|M_{ij}|$, and positive interactions ($M_{ij} > 0$) are green while inhibitory interactions ($M_{ij} < 0$) are red. M_{ij} was fit in [7] using experimental mouse data from [6]. To reduce dimensionality, bacterial species of the same genus are consolidated into one population; the exception is *C. difficile* (CD), which is a single bacterial species. CD, the culprit behind *C. difficile* infection (CDI), is colored red and located in the center of the food web.

not predict toxigenicity or toxin production, for the purposes of this chapter we equate CDI to the prolonged presence of CD in a simulated microbiome.

Stein *et al.* [7] investigated the existence and stability of steady states for the system Eq (3.1). Additionally, they found that for some initial compositions, antibiotic administration can alter a microbial composition to the degree that the composition becomes susceptible to CD colonization. Building upon their work, we propose the following three clinically relevant interventions and their corresponding in-silico implementations:

1. **inoculation** with CD at time t_I , corresponding to $\mathbf{x}(t_I) \mapsto \mathbf{x}(t_I) + \mathbf{x}_c$ where \mathbf{x}_c is purely composed of CD,
2. **antibiotic administration**, corresponding to a $u(t)$ that is (unless otherwise specified) a unit pulse of concentration c at $t = 0$, and
3. **transplantation** of CD-resilient microbiota into a CD-susceptible microbiome at time t_T , corresponding to $\mathbf{x}(t_T) \mapsto \mathbf{x}(t_T) + \mathbf{x}_{IC}$, where \mathbf{x}_{IC} is a transplant composed of a CD-resilient initial condition.

We refer to a simulation which implements any combination of these external interventions as a *treatment scenario*.

Simulations are run in Python with the `scipy` package and the `scipy.integrate.odeint` function, which uses ordinary differential equation solver `lsoda` from `odepack`, written in FORTRAN. This solver adaptively switches between stiff and non-stiff solvers, and simulations are run with an absolute tolerance of 10^{-12} . The code used to generate the figures in this chapter is freely available at https://github.com/erijones/simulated_CDI_with_gLV.

3.5.3 Implementation of simulated microbial transplants

Clinically, an external microbial transplant seeks to rejuvenate an unhealthy microbiome by infusing “healthy” microbes into the unhealthy patient. The infused samples typically consist of probiotics or a microbiome (often fecal) sample from a healthy subject [66]. Microbial transplants can confer attributes (e.g. obesity) from the donor to the donee [67], so in some sense a microbiome transplant is seeking to confer *CD-colonization resistance* from a CD-resilient donor to a CD-susceptible donee. Since antibiotics tend to be ineffective in treating CDI and additionally can facilitate the growth of drug-resistant mutant strains of CD by providing them with a selective advantage, fecal transplants are becoming an increasingly popular CDI treatment [61].

In our implementation we simulate transplants made of CD-resilient initial conditions, and demonstrate how these treatments can guide the system into a desired (i.e. noninfective) steady state. We model the administration of a transplant of some external microbial source \mathbf{v} at time t^* as

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{f}(\mathbf{x}) + \mathbf{v} \delta(t - t^*), \quad (3.2)$$

where $\mathbf{f}(\mathbf{x})$ entirely encapsulates the right-hand sides of the gLV equations of Eq (3.1) in vector form and $\delta(t)$ is the Dirac delta function, which will serve to instantaneously add the transplant \mathbf{v} to the microbial community \mathbf{x} at time t^* .

3.5.4 Implementation of sporulation

Under environmental pressures CD can sporulate, entering a defensive state of dormant spores that maintain the genetic information of CD while functioning at a fraction of the vegetative cell’s metabolism. These spores are resilient to antibiotics, and CD

sporulation may be induced by environmental stressors such as heat [68] and alcohol [62]. While the entire gamut of environmental conditions that induce sporulation is not yet known [69], there is some evidence that in murine models antibiotics may induce sporulation [61]. The toxin-producing types of CD prevalent in nosocomial infections are notoriously difficult to kill, and their resilience has in part been attributed to sporulation [61].

Mathematically, sporulation can be modeled by creating a population of spores that, through conversion of active CD, grows when environmental conditions are harsh and declines when conditions are mild. This implementation is inspired by the treatment of latently infected T-cells in SIR models of HIV, in which the latently infected T-cells effectively hide from the immune response in the same way that the inert spore cells are unaffected by the presence of antibiotics and other microbes [70]. To capture sporulation, we augment the basic model Eq (3.1) by introducing a spore compartment $s(t)$ so that the populations of the original gLV model become

$$\begin{aligned} \frac{d}{dt}x_i(t) &= x_i(t) \left(\mu_i + \sum_j M_{ij}x_j(t) + \varepsilon_i u(t) \right), \\ \frac{d}{dt}x_c(t) &= x_c(t) \left(\mu_c + \sum_j M_{cj}x_j(t) + \varepsilon_c u(t) \right) + \beta s(t)[u(t) < u_{spor}], \quad \text{and} \\ \frac{d}{dt}s(t) &= \alpha x_c(t)[u(t) \geq u_{spor}] - \beta s(t)[u(t) < u_{spor}], \end{aligned} \quad (3.3)$$

where the terms in square brackets should be interpreted as conditional statements that return 1 if true and 0 if false.

In Eq (3.3), we assume that the background microbes (which we define as the bacteria that are not CD) are unaffected by the presence of the inert spores. In the presence of antibiotics bacterial growth often acts as a step function, growing or not growing if

the antibiotic concentration is lower or higher than the bacteria's minimum inhibitory concentration (MIC)[71]. We similarly model the inflow and outflow of spores as a step function, where sporulation or germination occurs if the antibiotic concentration is larger or smaller than some threshold u_{spor} . Since the spores are robust, we assume they have no death rate. We assume that some proportion α of the CD normally killed by antibiotics are converted to spores, so there is no explicit α term in the CD growth term, and as a consequence of this we require $\alpha < \varepsilon_c u(t)$. The experimental methods used to measure CD sporulation are not yet standardized, so there is no clear consensus on the rate of CD sporulation [68]; therefore, the sporulation parameters α , β , and u_{spor} must be considered in a qualitative fashion.

3.5.5 Implementation of mutation

The final augmentation we add to the gLV model is antibiotic-resistant mutation, which is culpable for many of the difficulties in treating CDI [72]. Existing antibiotic resistance models for both within-host [73] and between-host [74] versions of antibiotic-resistance typically only consider isolated bacterial systems which include only the native and mutant strains of a single bacterial species. Since we consider mutation in the gLV framework, in this chapter we are able to probe the more realistic scenario of mutation occurring within a complex microbial community.

We modify the standard gLV model in Eq (3.1) to include terms that allow for mutation of CD into an antibiotic-resistant mutant strain of CD, denoted $x_m(t)$, so that the

microbial dynamics are described by

$$\begin{aligned}\frac{d}{dt}x_i(t) &= x_i(t) \left(\mu_i + \sum_j M_{ij}x_j(t) + \varepsilon_i u(t) \right), \\ \frac{d}{dt}x_c(t) &= x_c(t) \left(\mu_c + \sum_j M_{cj}x_j(t) + \varepsilon_c u(t) \right) - k x_c(t), \quad \text{and} \\ \frac{d}{dt}x_m(t) &= x_m(t) \left(\mu_m + \sum_j M_{mj}x_j(t) \right) + k x_c(t).\end{aligned}\tag{3.4}$$

In addition to the standard gLV pairwise interactions, the background microbes x_i of Eq (3.4) now interact with the CD mutant x_m via the M_{im} term. Following existing mutation models [74], we (1) group all potential antibiotic-resistant mutations into the one mutant population x_m and (2) neglect the possibility of mutation from a mutant strain x_m back to the native strain x_c . Furthermore, we assume that the mutations are fully resistant to antibiotics and so we omit the ε_m term in Eq (3.4). While other candidate models for antibiotic-resistant mutation exist and have been examined [75], here we focus on embedding this particular implementation of single-strain mutation into the gLV framework; other types of mutation models may be implemented in a similar way.

Since we are extrapolating beyond the mouse data collected in [6], it is not surprising that the mouse microbiome data does not distinguish between native and mutant strains of CD. Antibiotic resistant strains of CD are already rampant: one survey found that close to half of tested CD strains were resistant to at least one antibiotic, and about one quarter of tested strains were resistant to multiple antibiotics [76]. However, since the antibiotic susceptibility of CD ε_c is non-zero, we assume that the administered CD used to inoculate the mice is antibiotic-sensitive.

3.6 Results

3.6.1 Mapping system behaviors of the CDI model

We first demonstrate the available behaviors of the system described by Eq (3.1). In Fig 3.2 we evolve our system from the nine distinct initial conditions experimentally measured by Stein *et al.* [7] for one particular treatment scenario, in which all initial conditions are initially treated with antibiotics and later inoculated with CD. All but one of these initial conditions are free of CD, and the remaining initial condition (IC 8) has a trace amount of CD. Despite the diverse composition of the initial conditions, under this treatment scenario the simulated trajectories evolve into only two steady states.

Then, in Fig 5.2 we apply four different treatment scenarios to one initial condition and identify three different reachable steady states, indicating that the initial conditions can be sensitive to which treatment scenario is applied. In this chapter, within a single simulation microbe counts can vary by more than two orders of magnitude. For clarity, in our figures we plot the total microbe count on a log scale (where the total microbe count is the sum of all of the microbes in each microbial population), and then at each time we linearly color each microbial population according to its proportion at that time, so that at a given time regions of equal size correspond to equal microbe counts. The treatment scenarios that result in Fig 5.2 mirror the experimental mouse treatments [6] and include a control, high dosing with antibiotic (the inset of Fig 5.2b depicts the initial microbial response to antibiotics), low dosing with antibiotic followed by inoculation with CD, and high dosing with antibiotic followed by inoculation with CD. While the log scaling disguises changes in total microbe count between the different steady states, the steady state of Fig 5.2a contains seven times as many microbes as the depleted (in microbe count) steady state of Fig 5.2b, and contains more than twice as many microbes as the infected steady state of Fig 5.2d (for details on steady state compositions refer to Table

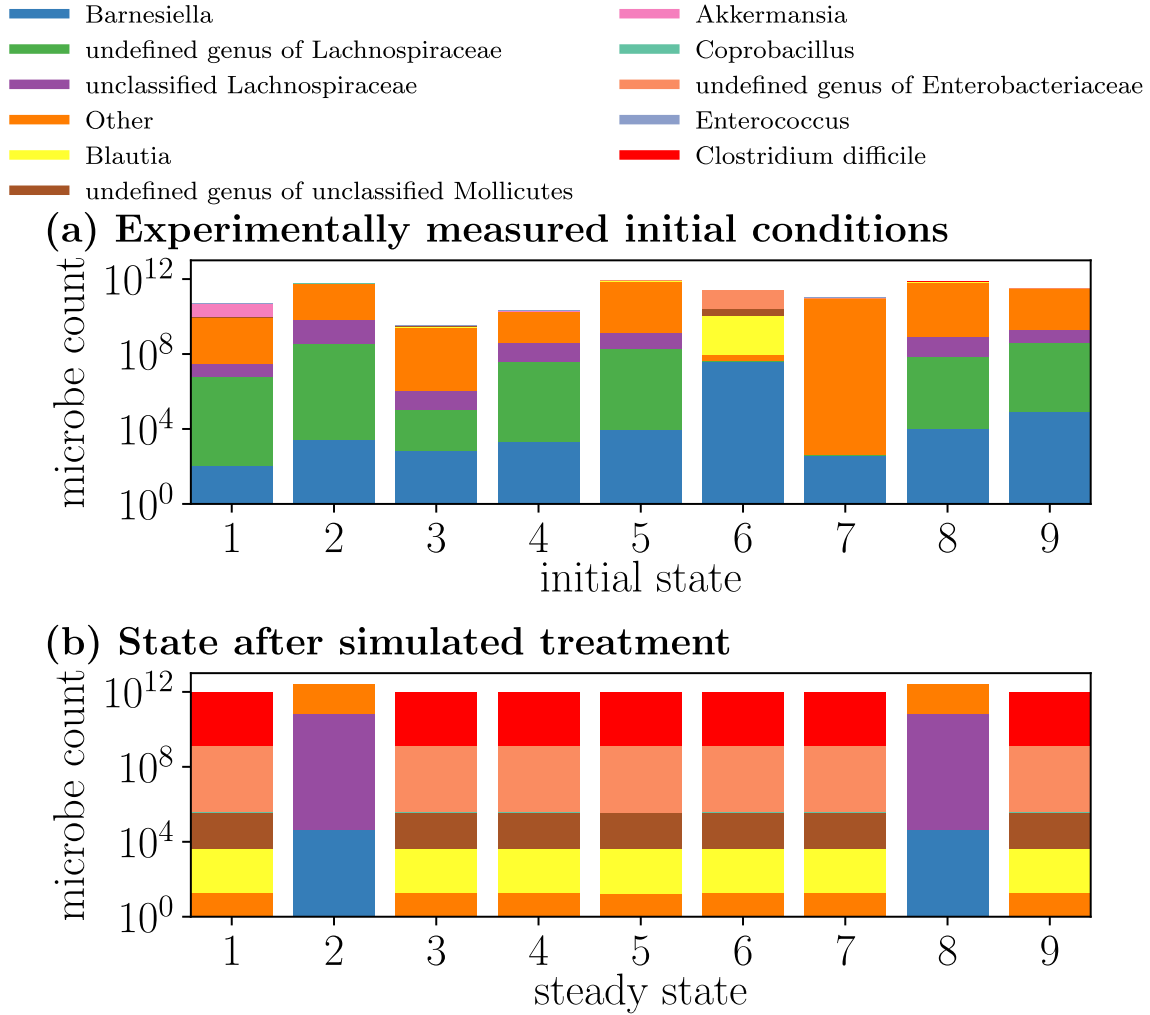


Figure 3.2: **Diverse initial conditions respond similarly for a simulated treatment.** Microbiome compositions are simulated forward in time from experimentally determined and diverse initial conditions (a), but all initial conditions eventually equilibrate to only two steady states (b) for this particular treatment scenario. Initial conditions are experimentally measured microbiome compositions from mice [6] and are time evolved according to the generalized Lotka-Volterra model, Eq (3.1), with parameters fit in [7]. In this simulated scenario, the system is administered 1 dose of antibiotic on day 0 and inoculated with the infectious CD (colored red) on day 10.

B of the Supplementary Information). This figure elucidates the mechanism for CDI: antibiotic-induced microbiome depletion followed by opportunistic CD colonization.

Taken together the complementary results of Figs 3.2 and 5.2 indicate that (1) for a given treatment scenario there are a limited number of achievable steady states across all initial conditions, and (2) for a given initial condition there are a variety of steady states that may be achieved across different treatment scenarios. Since the model was fit with data collected over a 30 day period but the obtained steady states are often slow to equilibrate (e.g. around 100 days in Fig 5.2), we should proceed with caution when extrapolating the model [77]. However, since the collected experimental data [6] roughly equilibrates by day 30, and because experimental validation on longer time scales is difficult to obtain, we follow convention [78] and study long-term system behavior through steady state analyses.

In the four weeks before the mouse experiment the mice were identically housed and fed, and during the experiment the microbial compositions of mice in the control group were approximately constant over time [6]. Hence, what we consider “initial conditions” may also be interpreted as steady states compositions of the mice before any external intervention. However, the gLV model Eq (1) does not capture these initial conditions as steady states. Over the course of the 13-day control group experiment the measured bacterial abundances maintained a relatively stable composition, with the 7 or 8 colonized bacteria varying by less than an order of magnitude over the course of the experiment. However, the model Eq (1) predicts that the same control group initial conditions (ICs 2, 5, and 8) will tend towards a simpler steady state that consists of only 3 bacteria.

This inconsistency demonstrates two limitations of the gLV model: the paucity of steady states, and the likelihood of their stability. For a generalized Lotka-Volterra system of N species there are 2^N steady states, each corresponding to a different subset of bacteria— hence, there is just one steady state that consists exclusively of the 7

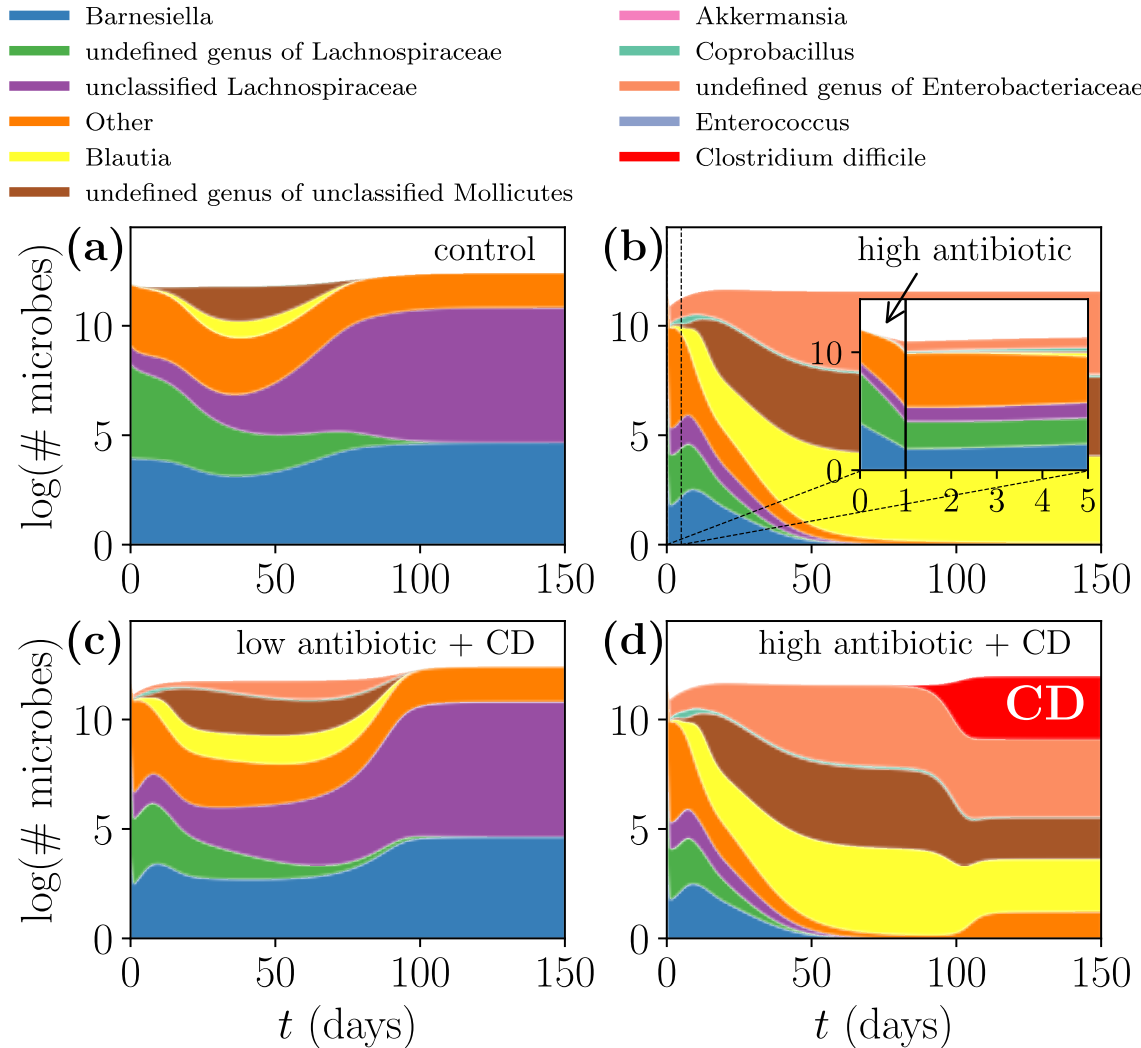


Figure 3.3: **External intervention can alter the steady state a given initial condition achieves.** All panels originate from the same initial condition, but different panels correspond to different interventions: (a) no interventions occur; (b) one dose of antibiotic is administered at day 0 (inset: the microbial dynamics during the first 5 days, in response to the one-day administration of antibiotics); (c) half of a dose of antibiotic is administered at day 0, and then at day 10 the system is inoculated with CD; (d) one dose of antibiotic is administered at day 0, and then at day 10 the system is inoculated with CD. The growth of *C. difficile* (colored red) is encouraged by antibiotic treatment, since the antibiotics deplete the other microbes to a level at which *C. difficile* gains a foothold.

overlapping bacteria of the control group. Since there is variation between the control experiments, there can be no steady state that would simultaneously and precisely fit all three control trials. Furthermore, even if this steady state were relatively accurate for each trial it is unlikely that it would be stable: Stein *et al.* [7] found that 98% of the steady states of this system were unstable. Despite the fact that unperturbed initial conditions are not stable steady states, other qualitative features of the model (including antibiotic-induced depletion of the microbiome and opportunistic CDI) indicate the model's utility in modeling CDI.

To summarize the available system dynamics, we construct the phase diagrams in Fig 3.4 by systematically sweeping through treatment scenarios for each initial condition; specifically, we vary the concentration of antibiotic treatment and whether the system is exposed to a small amount of CD. Though we simulate nine initial conditions (ICs), the phase diagrams for some initial conditions are redundant. We classify the phase diagrams of Fig 3.4 as (a) *CD-susceptible*, ICs which become infected upon exposure by CD regardless of antibiotic usage; (b) *CD-resilient*, ICs which are not infected by CD regardless of antibiotic usage; and (c) *CD-fragile*, ICs which switch from CD-resilient to CD-susceptible upon sufficient administration of antibiotic (an antibiotic concentration of approximately 0.71). We label the five reachable steady states A through E, categorize them as CD-infected or CD-uninfected, and plot their compositions in Fig. S1. Each phase diagram is composed of a number of treatment scenarios; for each treatment scenario, a 1-day pulse of antibiotic with varying antibiotic concentration is administered on day 0, and then a small amount of CD may be administered on day 10 depending on whether the scenario is with or without CD. For reference, the experimental antibiotic dose was normalized in [7] to a 1-day pulse of antibiotic concentration 1.

With the phase diagrams of Fig 3.4, we may now identify the initial condition plotted in Fig 5.2 as CD-fragile. Furthermore, the steady states of Fig 5.2a-d correspond, respec-

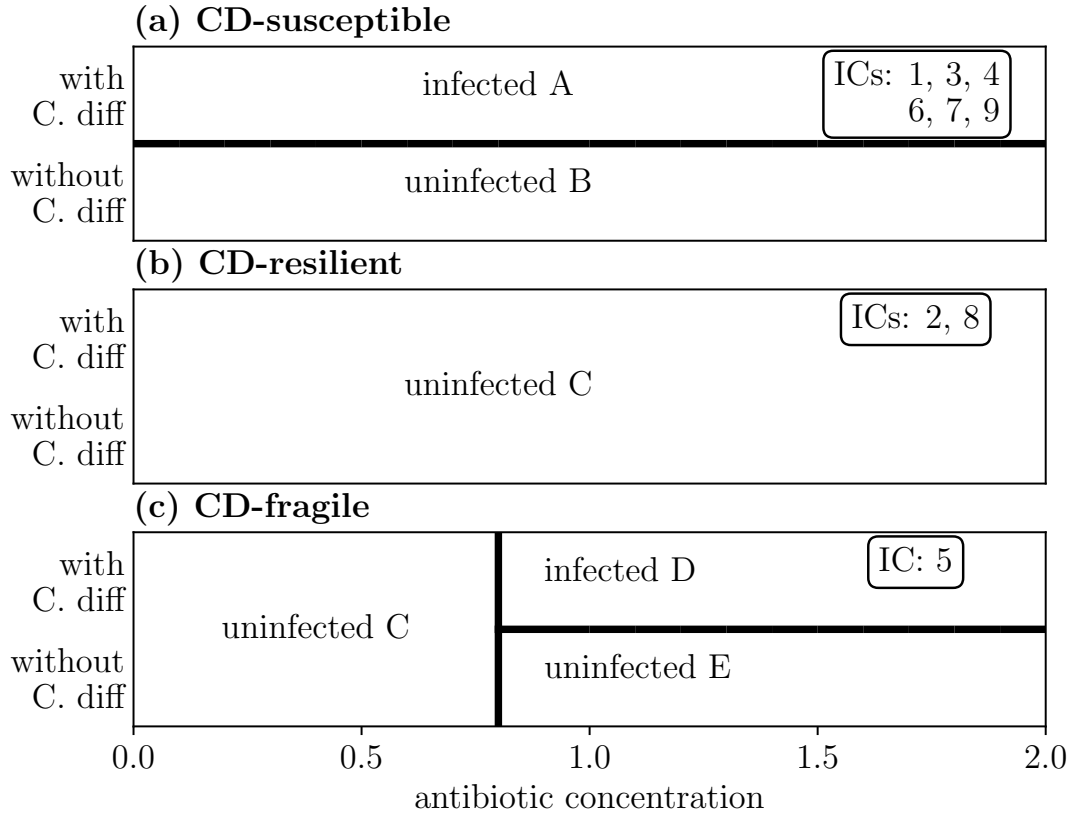


Figure 3.4: **Phase diagram of reachable steady states from initial conditions.** (a) Six initial conditions (ICs) are susceptible to *C. difficile* (*CD-susceptible*), resulting in infected steady state A if inoculated with any amount of CD. (b) Two initial conditions are *CD-resilient* and always remain in uninfected steady state C regardless of CD exposure. (c) One initial condition displays more complex behavior, becoming susceptible to CD only after being treated with a sufficient dose of antibiotics (*CD-fragile*). The steady states of the four external interventions of Fig 5.2 correspond to different regions of the CD-fragile phase diagram (c). The phase boundaries of Fig 3.4 are robust to the amount of CD inoculum (ranging from 10^{-10} to 1 in nondimensionalized units) as well as to the timing of CD inoculation (ranging from on day 1 to on day 100). For details, refer to Table A of the Supplementary Information.

tively, to steady states C, E, C, and D of Fig 3.4. Notably, IC 8 is CD-resilient despite the fact that the initial condition contains a small amount of CD; in fact, according to the fitted interactions the presence of CD promotes the growth of microbes that inhabit the uninfected steady state. Therefore, the isolated presence of CD inhibits colonization of an infected steady state.

One key takeaway from this survey of model behaviors is that there is no *a priori* obvious predictor for whether an initial condition will be CD-susceptible, CD-resilient, or CD-fragile, even with knowledge about the microbial food web. Often, the complex interplay of microbial interactions can lead to unexpected and even counterintuitive results.

3.6.2 Invadability of *C. difficile*

In the numerical phase diagrams of Fig 3.4 we observe different regimes for different initial conditions, but we can substantiate this phenomenon analytically as well. We label steady state A in Fig 3.4 by \mathbf{x}_A^* , and similarly label all other steady states. After all antibiotic has been administered, we perform a perturbative analysis of the uninfected steady states by introducing a small amount of CD (notated $x_c(t)$) to the uninfected steady state \mathbf{x}^* . This CD will invade the steady state only if $\left[\frac{d}{dt}x_c(t)\right]_{\mathbf{x}^*} > 0$. Since the introduced $x_c(t)$ is positive, we may discern the invadability of an uninfected steady state \mathbf{x}^* by the sign of $I(\mathbf{x}^*)$, defined to be

$$I(\mathbf{x}^*) \equiv \frac{1}{x_c(t)} \left[\frac{d}{dt}x_c(t) \right]_{\mathbf{x}^*} = (\boldsymbol{\mu} + M \mathbf{x}^*(\mathbf{t}))_c . \quad (3.5)$$

Here, we have rearranged Eq (3.1), removed the antibiotic dependence $u(t)$, consolidated all the μ_i and M_{ij} into their respective vector and matrix forms $\boldsymbol{\mu}$ and M , and consolidated the individual populations $x_i(t)$ into their vector form $\mathbf{x}(\mathbf{t})$. Notationally, the

Table 3.1: **Analytic justification for CD-susceptibility.** The ability of CD to invade the three uninfected steady states \mathbf{x}_B^* , \mathbf{x}_C^* , and \mathbf{x}_E^* depends upon the sign of $I(\mathbf{x}^*) \equiv \frac{1}{x_c(t)} \left[\frac{d}{dt} x_c(t) \right]_{\mathbf{x}^*}$: a positive value indicates a CD-susceptibility, while a negative value indicates CD-resilience. This result follows from analysis of Eq (3.1).

steady state	interpretation ($I(\mathbf{x}^*)$)	$\ \mathbf{x}^*(t)\ _1$
\mathbf{x}_B^*	CD-susceptible (0.24)	3.238
\mathbf{x}_C^*	CD-resilient (-0.86)	24.770
\mathbf{x}_E^*	CD-susceptible (0.28)	3.546

subscript c denotes the value of a vector corresponding to the index of CD. While magnitude of the invadability $|I(\mathbf{x}^*)|$ will correspond to the initial rate at which CD will grow or decay, only the sign of $I(\mathbf{x}^*)$ is relevant in determining long-term susceptibility or resilience to CD.

In Table 3.1 we compute and compile this invadability for each of the three uninfected steady states \mathbf{x}_B^* , \mathbf{x}_C^* , and \mathbf{x}_E^* . This table also provides the size of each steady state, where size is interpreted as the sum of all the bacterial populations (here written as the 1-norm). These conclusions provide analytic justification for why some initial conditions are susceptible to CD while others are not, and complement the phase diagrams in Fig 3.4.

CD is predominantly inhibited by the existence of other microbes (mostly, $M_{cj} < 0$) and so a larger $|\mathbf{x}^*(t)|$ will tend to inhibit the growth of CD. Additionally, microbes tend to be inhibited by antibiotics (mostly, $\varepsilon_i < 0$). Together, these tendencies allude to a mechanism of CDI whereby antibiotic administration depletes the microbiome and induces CD susceptibility.

While Table 3.1 indicates that the reachable CD-susceptible steady states are smaller than CD-resilient steady states, the size of the initial condition had little effect on the overall steady state profile: growing or shrinking the initial condition sizes only marginally modified the resulting phase diagrams. Hence, the different steady states are robust to variations in initial condition size.

Having exhaustively explored the basic behavioral regimes of Eq (3.1), we now implement in-silico two commonly administered real-world medical interventions: *fecal microbiome transplantation* and *antibiotic administration*.

3.6.3 Simulated microbial transplants affect steady-state outcomes

Following Eq (5.2), we choose a microbial transplant \mathbf{v} that is derived from a CD-resilient donor so that \mathbf{v} is proportional to the composition of a CD-resilient initial condition, and we choose the donee microbiome to be the CD-fragile initial condition so that the effects of the transplant are more apparent. In the simulation we choose the timing of the treatment scenario to match the clinical counterpart of CDI, in which CD attempts to colonize a microbiome that has been recently depleted by antibiotics: we administer antibiotics on day 0, inoculate with CD on day 1, and insert a transplant on day d . By categorizing the resultant steady state as CD-infected or CD-uninfected and sweeping over antibiotic concentrations, relative transplant sizes, and transplant times, we realize the phase diagram in Fig 3.5.

Fig 3.5 demonstrates how a transplant can alter the steady state behavior of a system exposed to CD. We can bias the initial condition towards a CD-uninfected steady state with a proper fecal transplant via the mechanism of *steady state conversion*, wherein a transplant can convert a state from CD-susceptible to CD-resilient. This result, consistent with clinical practice, supplies a numerical framing for microbial transplants, narrowing the gap between real-world practice and simulation.

For transplants that are applied after antibiotic administration, this figure indicates that shorter transplant delays lead to more effective transplants. However, a transplant applied concurrently with antibiotic administration on day 0 (labeled $d=0$ in Fig 3.5)

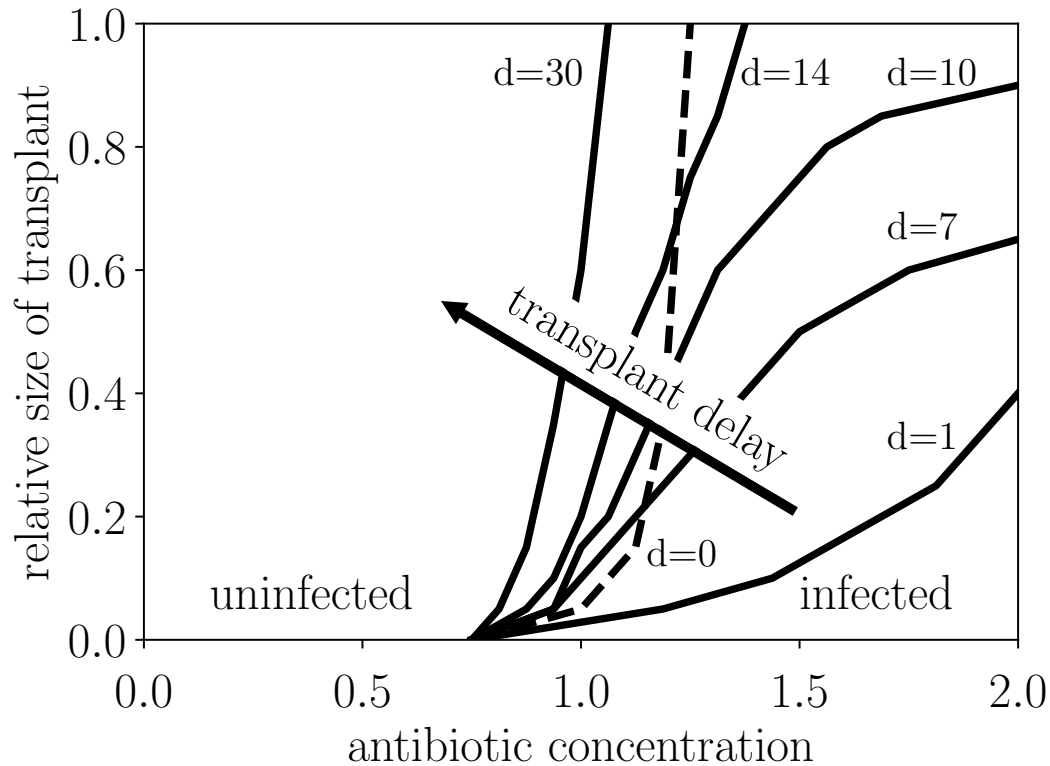


Figure 3.5: **Administration of microbial transplants can ward off infected steady states.** Starting from the CD-fragile initial condition, antibiotics of varying antibiotic concentration are administered on day 0, and the system is exposed to CD on day 1. Then, a “healthy” transplant made up of the CD-resilient initial condition 8 is infused on day d . The *infected* region corresponds to infected steady state D, and the *uninfected* region corresponds to uninfected steady state E. Note that a transplant on day 0 (dashed line), concurrent with the start of antibiotic administration, is less effective than a transplant on day 1. A relative transplant size of 1 corresponds to a transplant that has the same size as the initial condition that the transplant was derived from. The phase boundaries of Fig 5 are robust for small amounts of CD inoculum (ranging from 10^{-10} to 10^{-8} in nondimensionalized units), but for larger amounts of CD inoculum (ranging from 10^{-5} to 10^{-2}) transplants become more effective at all timings, requiring a smaller transplant to overcome a larger antibiotic dose. For details, refer to Table A of the Supplementary Information.

is less effective than a transplant applied after antibiotics on day 1. This reflects that antibiotic administration depletes *all* microbes, so a transplant on day 1 will be unsullied by antibiotics whereas applying a transplant on day 0 will lead to the depletion of the aggregate composition.

In Fig 3.6 we examine the effect of transplant timing for a fixed antibiotic concentration and transplant size. Steady state conversion is most effective immediately after antibiotic administration, when the depleted microbiome has room to grow. During this time the malleable microbiome is especially responsive to transplants, and introduction of the right collection of microbes can direct the microbiome towards an infection-free steady state. However, as indicated in Fig 5.2, without any transplant the CD-fragile IC will naturally evolve towards a CD-susceptible steady state: hence, the timing of the transplant is critical, with more immediate transplants being more effective.

We found that out of the measured ICs, the collection of microbes that best deter CDI are derived from IC 8. This transplant replenishes the unclassified Lachnospiraceae (colored purple), which promote constituents of the uninfected steady state while inhibiting *Blautia* (colored yellow), a key member of the infected steady state. More surprisingly, the existence of CD in IC 8 amplifies the effect of the transplant—the same transplant but without CD was a functional but substantially less effective treatment, and similarly mediocre results were obtained with a transplant derived from the other CD-resilient initial condition (IC 2) as displayed in Fig. S2. This result is due to the deleterious and contradictory effect of CD on the CD-infected steady state. As an aside, note that since IC 8 contains CD, the transplant on day 0 effectively inoculates the system with CD on day 0 rather than on day 1.

While appropriately derived and implemented transplants are effective at reversing CDI, if we had mistakenly used a CD-susceptible donor instead, simulation confirms the intuitive expectation that these results would be flipped. Since these initial conditions

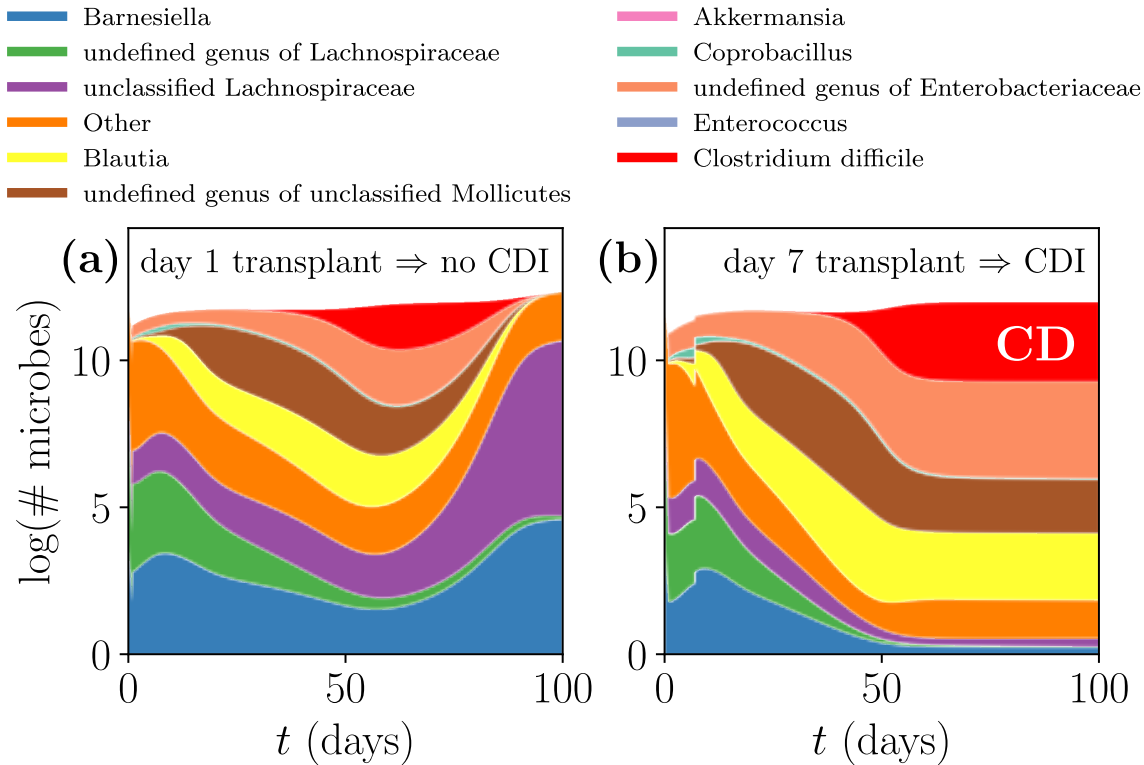


Figure 3.6: **Mechanism of *steady state conversion***. Microbiome compositions are identically dosed with antibiotic on day 0 and inoculated with CD (colored red) on day 1. Transplants of the same size are administered on (a) day 1, leading to a CD-uninfected steady state, and (b) day 7, leading to a CD-infected steady state. Following the antibiotic-induced microbiome depletion, the transplant serves to replenish the microbes responsible for the CD-uninfected steady state (e.g. unclassified Lachnospiraceae, colored purple) while curbing the growth of those responsible for the CD-infected steady state (e.g. Blautia, colored yellow).

are *a priori* unidentifiable as CD-resilient or CD-susceptible, this prompts a clinically relevant caution of whether some donor’s microbiome will be beneficial or detrimental to another’s microbiome.

In a recent experimental study by Buffie *et al.* [59] CD-vulnerable mice exposed to CD were given transplants consisting of a known microbial composition, and the transplant efficacy for each composition was measured. Our work on simulated transplants, which resembles the experimental study, provides context and explanation for the mechanism of the experimental transplants. In conjunction, simulated and experimental transplants could direct the development of model-guided and experimentally-validated “designer” transplants.

3.6.4 Simulated antibiotic dosing alters gLV dynamics

Antibiotic administration has traditionally been the standard approach to fight infection, but antibiotics have struggled to control CD infection: CDI has a recurrence rate of 30-65% following antibiotic treatment, while fecal transplantation has cure rates upwards of 90% [1]. Nonetheless, the Society for Healthcare Epidemiology of America (SHEA) and the Infectious Diseases Society of America (IDSA) jointly recommend treating CDI with antibiotics— often vancomycin— administered in one of three dosing regimens: a constant dosing regimen, a pulsed dosing regimen, or a tapered dosing regimen [62]. Other studies have found that vancomycin administered in tapered or pulsed doses reduced the likelihood of recurrent infections of CD, compared with treatment at a constant dosage [79]. Our model, which allows arbitrary control over the dosing schedule and concentration $u(t)$, provides a computational framework on which we can compare the efficacy of different dosing schedules: our implementations of the three dosing regimens are plotted in Fig. S3.

Over short time scales of 1-2 days we found that given the same total amount of antibiotic, the rate at which antibiotics were administered (e.g. .5 doses for 2 days vs. 2 doses for .5 days) did not affect the eventual steady state. Over longer time scales of around 2 weeks, we observed similar behavior—the model does not capture long-term differences between different dosing regimens as long as the total amount of administered antibiotic is the same, reflecting that the time-scale for microbial growth is longer than the period over which antibiotics are typically administered.

In modeling the different dosing regimens, we are faced with one main complication: only one antibiotic, clindamycin, was fit in [7], and furthermore clindamycin was acting to *induce* CDI rather than halt it. The antibiotic efficacy parameter ϵ therefore does not serve as a realistic proxy for vancomycin or metronidazole, antibiotics which are used to eliminate CD [79]. To simulate the effect of an antibiotic which eliminates CD, we introduce an artificial “targeted antibiotic” $\tilde{\epsilon}$, which by construction only inhibits CD; specifically, $\tilde{\epsilon}_c = -1$ and $\tilde{\epsilon}_i = 0$ for $i \neq c$.

Even with this targeted antibiotic our model does not capture significant differences between the treatment regimens, which is contrary to the clinical recommendation that pulsed or tapered dosing be preferred over constant dosing [62]. In Figs 3.7a and 3.7b we administer the same amount of targeted antibiotic via constant and pulsed dosing to the CD-infected steady state and find that the two dosing regimens produce near-identical microbe trajectories (a similar result, shown in Fig. S4, was found with tapered dosing). We propose sporulation (which acts on a much shorter time-scale) as a biologically relevant mechanism that could explain this inconsistency.

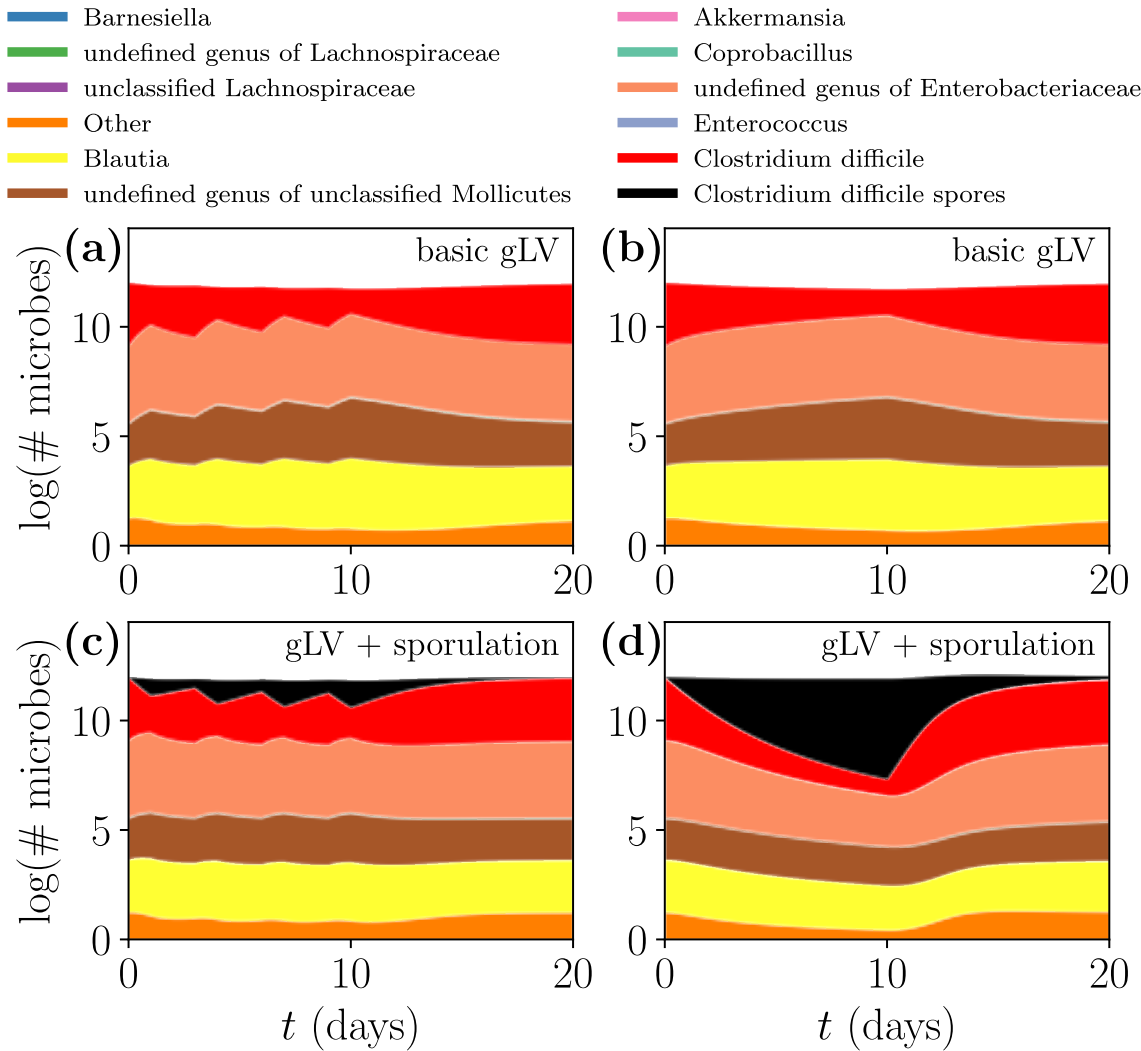


Figure 3.7: **Different antibiotic treatment regimens influence transient CD dynamics when considering sporulation.** All scenarios start from infected steady state D. Over 10 days, the same volume of “targeted” antibiotic is administered via a pulsed (a,c) or constant (b,d) dosing regimen. The top panels use the original gLV model Eq (3.1) while the bottom panels use the sporulation gLV model Eq (3.3). For details about the parameters used in this figure, refer to Table A of the Supplementary Information.

3.6.5 Sporulation affects microbial dynamics

In considering the model for sporulation detailed in Eq (3.3), the steady state analysis we previously performed is still relevant since all steady states will eventually be spore-free—we assume that the antibiotics will eventually cease, so the sporulation term of Eq (3.3) will eventually decay exponentially. However, the naïve expectation that including sporulation would make CD-infected steady states more common is incorrect; once again, due to the interactions between CD and other background microbes (mediated by the interaction matrix M), the presence of CD encourages growth of the microbes that populate the infection-free steady state, and so increasing the prevalence of CD through sporulation only entrenches the non-infective steady state. Since the steady states and phase diagrams are mostly unchanged by the inclusion of sporulation, we concentrate on the dynamics of CD and CD spores on shorter time scales.

In Fig 3.7 we compare the effects of the standard gLV model Eq (3.1) (top panels) and the sporulation model Eq (3.3) (bottom panels) under constant and pulsed antibiotic dosing regimens. Here, we use the targeted antibiotic previously described and apply all treatments to the CD-infected steady state D. Sporulation causes spores to form as antibiotics are administered, and germinate once the antibiotics cease, which is on display in the pulsed dosing regimen scenario of Fig 3.7c. After targeted antibiotic administration CD recovers slightly more quickly with sporulation than without, and we interpret this expedited resurgence as a more robust CDI. For details about the parameters used in Fig 3.7, refer to Table A of the Supplementary Information.

While many of the in-host dynamics and the biological mechanisms that underlie CD sporulation and germination remain under active investigation, studies have identified that both spores and vegetative CD colonize and persist in the gut [80], and other studies have discerned the role of bile acids in promoting spore germination [81]. Our model does

not allow for the long-term establishment of spores because we assume that germination always occurs in the absence of antibiotics, and we include no mechanism for germination induced by bile acids. However, more detailed sporulation models (e.g. models that include bile acid-induced germination) may extend our basic model to build upon the qualitative features of CDI it possesses.

We emphasize that sporulation is simply a proposed biological mechanism that would modify the model’s predictions to better match clinical observations, and so these results should be interpreted in a qualitative manner. However, by including sporulation we regain (at least for short time scales) the clinically expected result [62] that pulsed dosing is more effective than constant dosing at eliminating CD— comparing the top panels with the bottom panels of Fig 3.7 indicates that a pulsed dosing regimen dramatically reduces the buildup of CD spores compared to constant dosing.

3.6.6 Mutation in a gLV model

The mechanism of mutation, introduced in Eq (3.4), introduces new unconstrained parameters for the mutation rate k as well as for pairwise interactions M_{im} and M_{mi} . Here we identify an intuitive parameter choice that reflects the underlying biology, discuss the resultant steady states, and then demonstrate the effects of mutation on transient microbe dynamics.

Antibiotic-resistant mutations typically incur a fitness cost in the absence of antibiotics since resources are being allocated for defense against antibiotics rather than growth [82, 83, 84], so we choose $\mu_m = .9\mu_c < \mu_c$. Our choice of M assumes that the background microbes interact with mutant and native types identically (i.e. $M_{mi} = M_{ci}$ and $M_{im} = M_{ic}$ for $i \neq c, m$). In real systems, the mutation rate k is variable and depends on factors including the concentration of antibiotic, the type of antibiotic, the

native strain type, and other environmental pressures. In our model we approximate the mutation rate as a constant $k = 2 * 10^{-6}$ (in units of 1/day), a choice which is in the range of measured mutation rates of some bacteria [85], but for our purposes mostly serves to highlight the effects of mutation. For details about the parameters used in our simulations of the mutation model, refer to Table A of the Supplementary Information.

Due to our parameter choices the steady states of the background microbes are largely unchanged between the mutation model and the basic model (the CD-infected steady states of the standard and mutation models are explicitly compared in Table B of the Supplementary Information), but the transient dynamics shown in Fig 3.8 differ. In these plots the same amount of targeted antibiotic is applied to the same initial state, but Fig 3.8a uses the standard gLV model Eq (3.1) while Fig 3.8b uses the modified mutant gLV model Eq (3.4). The targeted antibiotic severely inhibits CD in the standard gLV model, but in the mutation model the antibiotic-resistant mutant compensates for the antibiotics and reinforces the colonization of CD despite the antibiotic pressures.

At the scale of a single bacterium experiments now track the growth and decline of individual lineages of bacteria when confronted with antibiotics [86], and at larger scales experiments track the spread and fixation of mutations across an entire bacterial community [87]. Since the gLV model considers populations of bacteria rather than individual cells, the individual lineages cannot be resolved. However, our model does capture the tendency of microbes with a selective advantage to outcompete microbes with lower fitness (in our case, CD mutants outcompete native CD in the presence of antibiotics), and these simulations resemble the selective sweeps found in experimental data [87].

Existing mutation models have studied native and mutant strains of bacteria in isolation, but by embedding mutation within a gLV framework we can probe the complex behaviors of mutant strains within a microbial consortia. Accordingly, the wealth of

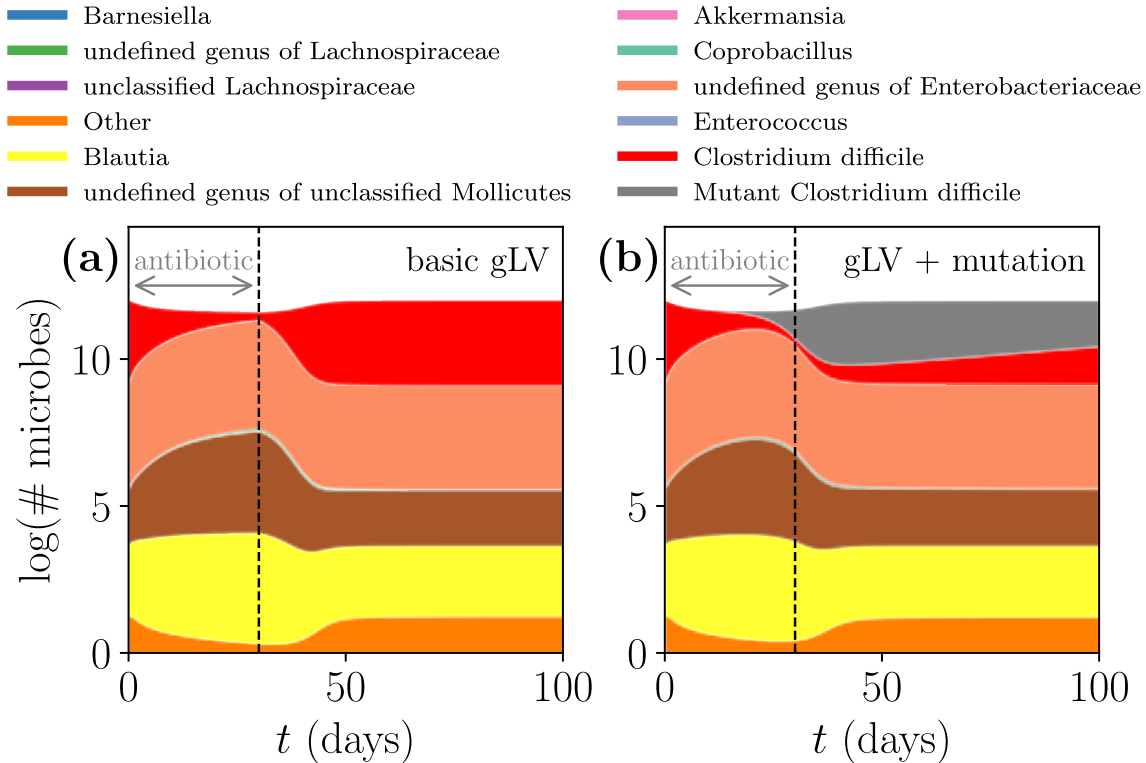


Figure 3.8: Antibiotic-resistant mutation improves CD resilience to antibiotics in transient dynamics. These simulations are identical except that (a) uses the original gLV model Eq (3.1) while (b) uses the mutation gLV model Eq (3.4). This scenario starts from the CD-infected steady state D and administers an idealized “targeted” antibiotic (that only inhibits CD) for 30 days. For details about the parameters and initial conditions used in this figure, refer to Table A of the Supplementary Information.

behaviors present in the simpler mutation models [74] may be observed within the gLV model with mutation Eq (3.4). This comprehensive and community-level view is essential in identifying, understanding, and resolving the role of antibiotic-resistant mutants in disease.

3.7 Discussion

3.7.1 Applications towards experimental explanation, model validation, and model-motivated experiments

A study by Buffie *et al.* [59] follows the modeling method of Stein *et al.* [7] and fits a gLV model to both mouse and human experimental time-series data in order to predict the growth of CD following antibiotic administration. In this study they identify the microbes anticorrelated with CDI in experimental data as well as the microbes that most inhibit the growth of CD according to the interaction matrix M of the gLV model. They create and administer transplants made up of a subset of these identified microbes: four transplants consist of individual microbes in isolation while another consists of a combination of all four microbes.

Buffie *et al.* [59] find that of the four transplants made up of isolated microbes only one microbe is effective at curing CDI, despite the fact that the other three microbes were *a priori* supposed to inhibit CD. We provide two explanations for their findings, motivated by the results of our chapter: (1) the ability of CD-resilient transplants to confer CD-colonization resistance is largely variable and depends on the transplant composition (e.g. the variation in transplant efficacies between Fig 3.5 and Fig. S2), and (2) inhibiting the growth of CD does not necessarily inhibit CD-infected steady states, since the presence of CD inhibits some of the microbes that populate the CD-infected steady state. By

applying the results of our simulations to microbiome transplant experiments, we can offer a computational context for experimental results.

Many of the microbes identified by Buffie *et al.* [59] as potential transplants were of the genus *Clostridium*; in fact, the only isolated-microbe transplant that was effective in curing CDI was *Clostridium scindens*. If we resolve only to the genus level (as assumed when constructing the gLV model in Stein *et al.* [7]), this experimental result is consistent with our own transplant simulations in which a transplant made up of the CD-resilient IC 8 was significantly more effective with CD than without (Fig 3.5). Hence, the seemingly contradictory computational result—that the presence of CD inhibits CD-infected steady states—is validated by experiment.

Finally, we can formulate experimental questions that are couched in our computational framework. Our results point to the importance of timing when administering microbial transplants, an area that is mostly unexplored both experimentally and therapeutically, and experiments could elucidate how the the timing of transplants effects their efficacy. While Buffie *et al.* [59] inferred microbial interactions from a gLV model, when predicting the CD-inhibiting microbes their analysis did not include dynamic simulations; applying our method of simulated transplants to such experiments could inform the selection of “personalized” transplants, and the corresponding experiments could then be used to inform the model, the model’s limitations, and additional experiments.

3.7.2 Unidentifiability of beneficial bacterial communities

In this chapter, the principle driver of CDI was whether a given microbial composition was CD-resilient or CD-susceptible: for example, when administering a fecal transplant, the effectiveness of the treatment depended on the properties of the donor’s microbiome. In general these properties are unknown *a priori*, so picking the right donor is a gamble.

In clinical practice, the screening process for potential fecal donors consists primarily of avoiding those with impaired microbiomes (e.g. due to recent antibiotic therapy) or poor health, with only about 10% of prospective donors being accepted [88, 89]. While fecal transplantation has been more successful at curing CDI than traditional antibiotic treatments [90], predictive models are not currently being implemented to quantitatively select an optimal donor. Eventually, predictive models could allow for “designer” fecal transplants that are engineered to optimally confer colonization resistance. Until donor selection methods consist of searching for optimal donors rather than excluding diseased donors, our model warns that donor selection— even of seemingly healthy donors— can have unexpected consequences.

3.7.3 Pharmacokinetic and pharmacodynamic approximations

In this chapter, we follow Stein *et al.* [7] and model the pharmacokinetics (the in-host concentration of the antibiotic $u(t)$) as a pulse. In reality, clindamycin pharmacokinetics are characterized by an initial spike in the in-host antibiotic concentration, after which antibiotics are cleared from the system (driven by uptake and deterioration of the antibiotic) with a half-life of approximately 4 hours [91]. However, in our simulations we found that over short durations (1-14 days) it is the total amount of administered antibiotic that determines the long-term dynamics of Eq (1) rather than the shape of the dosing regimen $u(t)$ (meaning that administering .5 doses for 2 days leads to the same outcome as administering 2 doses for .5 days). This insensitivity to the form of $u(t)$ justifies our simplified pharmacokinetic form.

Additionally, we model the pharmacodynamics (the microbial response or killing rate due to antibiotics) as a linear response $-\varepsilon_i u(t)$, while more realistic models use a saturating Hill function [92, 93]. However, we only use one antibiotic concentration for each

simulation, corresponding to one killing rate for each simulation. For any killing rate in the range of the saturating Hill function, one may find an effective antibiotic concentration that achieves this killing rate via either the linear response or by the saturating Hill function. Since both the linear response and the saturating Hill function are monotonic, there is a nonlinear scaling for $u(t)$ between the two response functions, meaning that our results—acquired with the linear response function—may be extended to a model that uses a saturating Hill function as long as the antibiotic scaling is observed (e.g. for the phase diagram of Fig 3.5, stretch the antibiotic axis). Since a linear antibiotic response qualitatively captures the same long-term dynamics that a saturating Hill function would, we are justified in using a simplified pharmacodynamic model.

3.7.4 Limitations of the gLV model

The gLV model idealizes interspecies interaction, and this simplification imposes limitations on our framework. The gLV model does not explicitly model *why* populations grow or decay (due to the underlying resource excesses or limitations) [94], and populations are assumed to respond instantly to changes in other populations, failing to account for the time required to respond to change [95]. The number of parameters required for a gLV model scales as N^2 for N species, and even with high-throughput sequencing, the number of data points per parameter is still low (e.g. roughly 5 data points per parameter in [7]). Despite these drawbacks, gLV models are commonly implemented to describe microbial growth [60] since they are predictive, manipulable, and often capture the qualitative characteristics of microbial consortia. Our framework attempts to resolve some of these limitations by treating the gLV model as a base model, then offering extensions to the model that incorporate nongeneric and mechanistic features in order to more accurately portray microbial growth.

3.7.5 Analytic concerns of parameter fitting

There are many techniques that fit parameters to data [96, 97], but it is difficult to know that these fitted parameters are indeed the true parameters. Stein *et al.* [7] fit the parameters used in this chapter with regularized linear regression with a Tikhonov regularization, but other fitting methods exist, such as LIMITs [98], a software specifically designed for fitting microbial time-series data to a gLV model. Analytically, there are sufficient conditions on the model parameters that ensure the Lyapunov stability of fixed points of generalized Lotka-Volterra systems [57], but the fitted parameter values in this chapter do not satisfy all of these conditions. Leveraging fitting methods to simultaneously fit parameters to data while maintaining the analytic properties that ensure stability would alleviate the potential for non-biological divergences in microbe count (divergences which are not impossible in the given system since M has a single positive eigenvalue). Regardless of this possibility, no unstable behavior was observed in any of the simulations run for this chapter, perhaps due to the relatively few symbiotic relationships.

3.7.6 Combining gLV and SIR techniques

In this chapter we fuse standard SIR techniques with the gLV model, thereby introducing specific mechanisms for sporulation and mutation. In this way, our framework allows for non-generic attributes of populations to be captured and simulated, and the resulting analyses provide qualitative insights into different mechanisms. Effectively, this allows for the entire family of SIR methods to be used in conjunction with the gLV model.

3.8 Conclusion

As the era of personalized medicine approaches, there is a growing need for accurate computational models that reflect human biology and can predict the progress of disease. This pursuit will be aided by the availability of “big data” in medicine, but this data needs to be harnessed in a useful way. This chapter addresses initial steps in developing these computation models by constructing a framework at the interface between computational models and clinical therapies. This modular framework allows for “plug-and-play” implementations of clinical techniques and observed phenomena: in this chapter, we implement fecal transplant therapy, antibiotic treatment regimens, sporulation, and mutation.

Our in-silico implementations of clinical treatments were mostly congruent with the actual clinical realizations— there exist initial conditions that become susceptible to CD after exposure to antibiotics; administration of a fecal transplant can halt CDI; and (once sporulation is included) pulsed dosing is more effective at eliminating CD than constant dosing, though fecal transplants are more effective than antibiotic administration in the long run. Introducing mechanisms for antibiotic-resistant mutations and sporulation strengthens the resilience of CD to remedial treatments. In all, this framework captures the intention and qualitatively the results of real-world clinical techniques.

There are many avenues stemming from this framework that may be explored in the future, including research into “designer transplants” or of bile acid-mediated germination of CD spores. Eventually this framework could be used to suggest clinical practices, but first more experiments, better data, and novel modeling are needed. As we recognize the advancement of gene sequencing in the past few years, it is not inconceivable that user-specific personalized medicine programs, built upon mathematical models of human health, will be accessible in the future.

Chapter 4

Compression of high-dimensional bistable dynamics with steady-state reduction (SSR)

Chapter 3 revealed that in the CDI model, the efficacy of FMT is dependent on transplant composition and the timing of administration. These results were gathered through extensive numerical simulations and, while intriguing, stop short of explicitly identifying the mechanisms within the gLV model that give rise to those behaviors. In this chapter we introduce the dimensionality-reduction technique steady-state reduction (SSR), which maps the bistable ecological dynamics of a high-dimensional gLV system to a compressed two-dimensional (2D) gLV system that approximates the high-dimensional dynamics. Critically, this reduced system retains the bistability of the high-dimensional system: SSR maps a pair of stable steady states \vec{y}_a and \vec{y}_b in the high-dimensional system (each of which are N -dimensional vectors) to a pair of stable steady states $(\|\vec{y}_a\|, 0)$ and $(0, \|\vec{y}_b\|)$ in the two-dimensional model, where $\|\cdot\|$ is the 2-norm. The development of SSR is the key analytical result of this thesis.

With SSR, the ecological dynamics that occur between a pair of stable steady states—say, the dynamics between a diseased and healthy steady state— may be analyzed in terms of the reduced 2D model. The high-dimensional dynamics mapped onto the plane spanned by \vec{y}_a and \vec{y}_b are similar to the 2D dynamics of the SSR-reduced system, even though the high-dimensional system has $\sim N^2$ parameters and the SSR-reduced system only uses 6 parameters. Indeed, we prove that the choice of SSR parameters (each a linear combination of the high-dimensional parameters and the steady states \vec{y}_a and \vec{y}_b) minimizes the error between the projected high-dimensional dynamics and the 2D gLV dynamics. Therefore, there is an explicit correspondence between high-dimensional bistable dynamics and the dynamics of their SSR-reduced counterpart.

With this correspondence, the high-dimensional dynamics that give rise to the previously observed composition- and timing-dependent efficacy of direct bacteriotherapies may be mapped to behaviors in the SSR-reduced model. After applying SSR to the transition between the healthy and CDI-susceptible steady states of the CDI model, we find that the separatrix of the SSR-reduced system closely aligns with the location of the separatrix in the high-dimensional system. Notably, the separatrix of the 2D system may be derived analytically (and therefore efficiently) as a power series expansion about a semistable fixed point. Determining the location of the separatrix in the high-dimensional system is computationally expensive. The ability to instantly approximate the location of the high-dimensional separatrix via SSR is handy, and will be valuable in chapter 5 when it is used to efficiently map the dynamical landscape of high-dimensional gLV systems. In chapter 6 this two-dimensional/high-dimensional correspondence will be used once again to guide the creation of indirect bacteriotherapies.

Finally, as a forward-looking comment, note that the 2D gLV system can be interpreted as a normal form of bistability. While SSR employs this normal form to analyze bistability in high-dimensional gLV systems, other mathematical models that feature

bistability (e.g. resource-consumer models or gene regulatory networks) could feasibly be represented by this 2D normal form as well. Thus, the direct and indirect control protocols that we apply to gLV systems in chapters 5 and 6 might one day be relevant for other types of bistable systems.

4.1 Abstract of “Steady-state reduction of generalized Lotka-Volterra systems in the microbiome”

The generalized Lotka-Volterra (gLV) equations, a classic model from theoretical ecology, describe the population dynamics of a set of interacting species. As the number of species in these systems grow in number, their dynamics become increasingly complex and intractable. We introduce steady-state reduction (SSR), a method that reduces a gLV system of many ecological species into two-dimensional (2D) subsystems that each obey gLV dynamics and whose basis vectors are steady states of the high-dimensional model. We apply this method to an experimentally-derived model of the gut microbiome in order to observe the transition between “healthy” and “diseased” microbial states. Specifically, we use SSR to investigate how fecal microbiota transplantation, a promising clinical treatment for dysbiosis, can revert a diseased microbial state to health.

4.2 Introduction

The long-term behaviors of ecological models are proxies for the observable outcomes of real-world systems. Such models might try to predict whether a pathogenic fungus will be driven to extinction [99], or whether a microbiome will transition to a diseased state [60]. In this chapter we explicitly account for this outcome-oriented perspective with *steady-state reduction* (SSR). This method compresses a generalized Lotka-Volterra

(gLV) model of many interacting species into a reduced two-state gLV model whose two unit species represent a pair of steady states of the original model.

This reduced gLV model is defined on the two-dimensional (2D) subspace spanned by a pair of steady states of the original model, and the subspace itself is embedded within the high-dimensional ecological phase space of the original gLV model. We prove that the SSR-generated model is the best possible gLV approximation of the original model on this 2D subspace. The parameters of the reduced model are weighted combinations of the parameters of the original model, with weights that are related to the composition of the two high-dimensional steady states. We note that SSR could be extended to encompass three or more steady states, but the resulting reduced systems would quickly become analytically opaque. In Section 4.3 we describe SSR and its implementation in detail.

We apply this method to the microbiome, which consists of thousands of microbial species in mammals [49], and which exhibits distinct “dysbiotic” microbial states that are associated with diseases ranging from inflammatory bowel disease to cancer [100]. Microbial dynamics are mediated by a complex network of biochemical interactions (e.g. cellular metabolism or cell signaling) performed by microbial and host cells [101, 102]. Ecological models, including the gLV equations, seek to consolidate these myriad biochemical mechanisms into nonspecific coefficients that characterize the interactions between microbial populations. We consider one particular genus-level gLV model of antibiotic-induced *C. difficile* infection (CDI), which was fit with microbial abundance data from a mouse experiment [7, 6].

This CDI model exhibits steady states that correspond to experimentally-observed outcomes of health (i.e. resistance to CDI) or dysbiosis (i.e. susceptibility to CDI). The transition between these healthy and diseased states is difficult to effectively probe due to the high dimensionality of the system, so previous analyses have been largely limited to numerical methods [9]. By reducing the dimensionality of the original gLV model,

SSR enables this transition to be investigated with analytic dynamical systems tools. We demonstrate the fidelity of SSR as applied to this CDI model in Section 4.4, and describe the analytic tools accessible to reduced gLV systems in Section 4.5.

Finally, we use SSR to analyze the clinically-inspired scenario of antibiotic-induced CDI. Specifically, we examine the bacteriotherapy *fecal microbiota transplantation* (FMT), in which gut microbes from a healthy donor are engrafted into an infected patient, and which has shown remarkable success in treating recurrent CDI [103]. In Section 4.5 we implement FMT in the reduced model and successfully revert a disease-prone state to health, and also find that the efficacy of FMT depends on the timing of its administration. In Section 4.6 we show that this dependence on FMT timing, also present in the experimentally-derived CDI model [9], is preserved under SSR.

4.3 Compression of generalized Lotka-Volterra systems

We begin by demonstrating how to compress the high-dimensional ecological dynamics of the generalized Lotka-Volterra (gLV) equations, given in Eq. (4.1), into an approximate two-dimensional (2D) subspace. This process, called steady-state reduction (SSR), is depicted schematically in Fig. 5.4. The idea behind SSR is to recast a pair of fixed points of a high-dimensional gLV model as idealized ecological species in a 2D gLV model, and to characterize the interactions between these two composite states by taking a weighted average over the species interactions of the high-dimensional system. Within this subspace, these reduced dynamics constitute the best possible 2D gLV approximation of the high-dimensional gLV dynamics.




Figure 4.1: Schematic of steady State Reduction (SSR). A gLV system of N species (Eq. (4.1)) exhibits two steady states \vec{y}_a and \vec{y}_b , characterized as diseased (red) and healthy (green). SSR identifies the two-dimensional (2D) gLV system defined on the 2D subspace spanned by the two high-dimensional steady states (Eq. (5.3)) that best approximates the high-dimensional system. Specifically, SSR prescribes 2D parameters (Eq. (4.3)) that minimize the deviation between the N -dimensional gLV dynamics $d\vec{y}/dt$ and the embedded 2D SSR-reduced dynamics $d\vec{x}/dt$.

The gLV equations model the populations of N interacting ecological species y_i as

$$\frac{d}{dt}y_i(t) = y_i(t) \left(\rho_i + \sum_{j=1}^N K_{ij}y_j(t) \right), \quad (4.1)$$

for $i \in 1, \dots, N$. In vector form, these microbial dynamics are written $\frac{d\vec{y}}{dt} = \sum_{i=1}^N \frac{dy_i}{dt} \hat{y}_i$. Here, the growth rate of species i is ρ_i , and the effect of species j on species i is given by the interaction term K_{ij} . In the following derivation, we assume this model observes distinct stable fixed points \vec{y}_a and \vec{y}_b .

Define variables x_a and x_b in the direction of unit vectors \hat{x}_a and \hat{x}_b that parallel the two steady states according to $\hat{x}_a \equiv \vec{y}_a / \|\vec{y}_a\|_2$, and $\hat{x}_b \equiv \vec{y}_b / \|\vec{y}_b\|_2$, where $\|\cdot\|_k$ is the k -norm. The 2D gLV dynamics on the subspace spanned by \hat{x}_a and \hat{x}_b are given by

$$\begin{aligned} \frac{dx_a}{dt} &= x_a(\mu_a + M_{aa}x_a + M_{ab}x_b), \quad \text{and} \\ \frac{dx_b}{dt} &= x_b(\mu_b + M_{ba}x_a + M_{bb}x_b). \end{aligned} \quad (4.2)$$

The *in-plane dynamics* on this subspace in vector form are defined to be $\frac{d\vec{x}}{dt} = \frac{dx_a}{dt} \hat{x}_a + \frac{dx_b}{dt} \hat{x}_b$.

SSR links the parameters of the in-plane dynamics to the high-dimensional gLV dynamics by setting

$$\begin{aligned} \mu_\gamma &= \frac{\vec{\rho} \cdot (\vec{y}_\gamma^{\circ 2})}{\|\vec{y}_\gamma\|_2^2}, \quad \text{for } \gamma \in a, b, \quad \text{and} \\ M_{\gamma\delta} &= \frac{(\vec{y}_\gamma^{\circ 2})^T K \vec{y}_\delta}{\|\vec{y}_\gamma\|_2^2 \|\vec{y}_\delta\|_2}, \quad \text{for } \gamma, \delta \in a, b. \end{aligned} \quad (4.3)$$

Here, the Hadamard square represents the element-wise square of a vector, defined as $\vec{v}^{\circ 2} = [v_1^2, v_2^2, \dots, v_N^2]^T$. The parameter definitions in Eq. (4.3) are valid when \vec{y}_a and \vec{y}_b are orthogonal; when they are not, the cross-interaction terms M_{ab} and M_{ba} become more complicated, and are given in Eqs. (4.28) and (4.29) of the Appendix.

This choice of parameters minimizes the deviation between the in-plane and high-

dimensional gLV dynamics $\epsilon = \|\frac{d\vec{y}}{dt} - \frac{d\vec{x}}{dt}\|_2$ for any point on the subspace spanned by \hat{x}_a and \hat{x}_b . This is proved in the Appendix by showing that, when evaluated with the SSR-prescribed parameter values of Eq. (4.3), $\frac{\partial \epsilon}{\partial c_i} = 0$ for every coefficient $c_i \in \{\mu_a, \mu_b, M_{aa}, M_{ab}, M_{ba}, M_{bb}\}$, and that $\frac{\partial^2 \epsilon}{\partial c_i \partial c_j} > 0$ for every pair of coefficients c_i and c_j .

Under this construction, the high-dimensional steady states \vec{y}_a and \vec{y}_b have in-plane steady state counterparts at $(\|\vec{y}_a\|_2, 0)$ and $(0, \|\vec{y}_b\|_2)$, respectively. It is for this reason we call this method *steady-state reduction*. Further, if \vec{y}_a and \vec{y}_b are stable and orthogonal, then the corresponding 2D steady states are stable as well, which guarantees the existence of a separatrix in the reduced 2D system. These properties are shown in the Supplementary Information, which includes many other calculations that accompany the results of this chapter. We provide a Python module that implements SSR on arbitrary high-dimensional gLV systems in the Supplementary Code¹.

If the ecological dynamics of the system lie entirely on the plane spanned by \hat{x}_a and \hat{x}_b , the SSR approximation is exact. In this case, the plane contains a slow manifold on which the ecological dynamics evolve. Therefore, the dynamics generated by SSR result from a linear approximation of the slow manifold.

4.4 Steady-state reduction applied to a microbiome model

Thousands of microbial species populate the gut microbiome [49], but for modeling purposes it is common to coarse-grain at the genus or phylum level. Recently, many experimentally derived gLV microbiome models have been constructed with tools such as

¹Supplementary Code used to implement SSR and generate Fig. 2 available at github.com/erijones/ssr_module.

MDSINE, a computational pipeline that infers gLV parameters from time-series microbial abundance data [104]. SSR is applicable to any of these gLV systems, so long as it exhibits two or more stable steady states.

We consider one such experimentally derived gLV model, constructed by Stein et al. [7], that studies CDI in the mouse gut microbiome. This model takes the same form as Eq. (4.1) and tracks the abundances of 10 different microbial genera and the pathogen *C. difficile* (CD), all of which can inhabit the mouse gut. The 11-dimensional (11D) parameters of this model were fit with data from an experimental mouse study [6]. The parameters of this model, along with a sample microbial trajectory, are provided in the Supplementary Information.

Despite the fact that this model did not resolve individual bacterial species, it still captured the clinically- and experimentally-observed phenomenon of antibiotic-induced CDI, suggesting that the true microbiome’s dimensionality could be approximated by an 11-dimensional model. SSR further simplifies the dimensionality of the microbiome: instead of thousands of microbial species or even eleven dominant genera, with SSR steady states of the microbiome (each of which are multi-species equilibrium populations) are idealized as individual ecological populations.

This CDI model exhibits five steady states that are reachable from experimentally measured initial conditions [7]. In previous work, we identified which of these steady states were susceptible or resilient to invasion by *C. difficile* (CD) [9]. Based on this classification, we interpret a CD-susceptible steady state \vec{y}_a of the 11D model as “diseased,” and interpret a CD-resilient steady state \vec{y}_b as “healthy.” Explicit details about each of these states are provided in the Supplementary Information. These two states are used to demonstrate SSR.

The reduced 2D parameters are generated according to Eq. (4.3). We introduce new scaled variables, $z_a = x_a / \|\vec{y}_a\|_2$, and $z_b = x_b / \|\vec{y}_b\|_2$, so that the 2D system exhibits steady

states at $(1, 0)$ and $(0, 1)$. In Fig. 4.2, trajectories of the reduced system (solid lines) that originate from four initial conditions and tend toward either the healthy (green) or diseased (red) steady states are plotted. The 2D separatrix is also plotted (light grey), which divides the basins of attraction of the two steady states, and which is derived in Eq. (4.6) of the Section 4.5.

To compare the original and reduced models, we consider 11D trajectories that originate from the 11D embedding of the four 2D initial conditions. The projections of these 11D trajectories onto the 2D subspace spanned by \vec{y}_a and \vec{y}_b (dashed lines) are shown alongside the corresponding 2D trajectories in Fig. 4.2. The in-plane 11D separatrix is also shown (dark grey), which is numerically constructed by tracking the steady state outcomes of a grid of initial conditions on the plane.

We note that \vec{y}_a and \vec{y}_b are nearly orthogonal. However, in the Supplementary Information we demonstrate that the high-dimensional and SSR-reduced trajectories and basins of attraction agree for four different implementations of SSR; in two of these implementations the pairs of steady states were orthogonal, and in the other two they were not. It is important to understand when SSR is a good approximation, and under what conditions it may be successfully applied— this issue will be addressed in a future publication (in progress).

In the five realizations of SSR explored in this chapter and in the supplement, the basins of attraction and microbial trajectories are largely preserved through SSR. Since the 11D system has been compressed (from 132 parameters to 6), it is not surprising that the low- and projected high-dimensional trajectories do not exactly match. Even so, the basins of attraction agree almost entirely, and the dynamical trajectories share clear similarities. The deviation between the original and reduced systems is examined in more detail in the Supplementary Information. The close agreement between the original and reduced systems intimates the reductive potential of SSR.

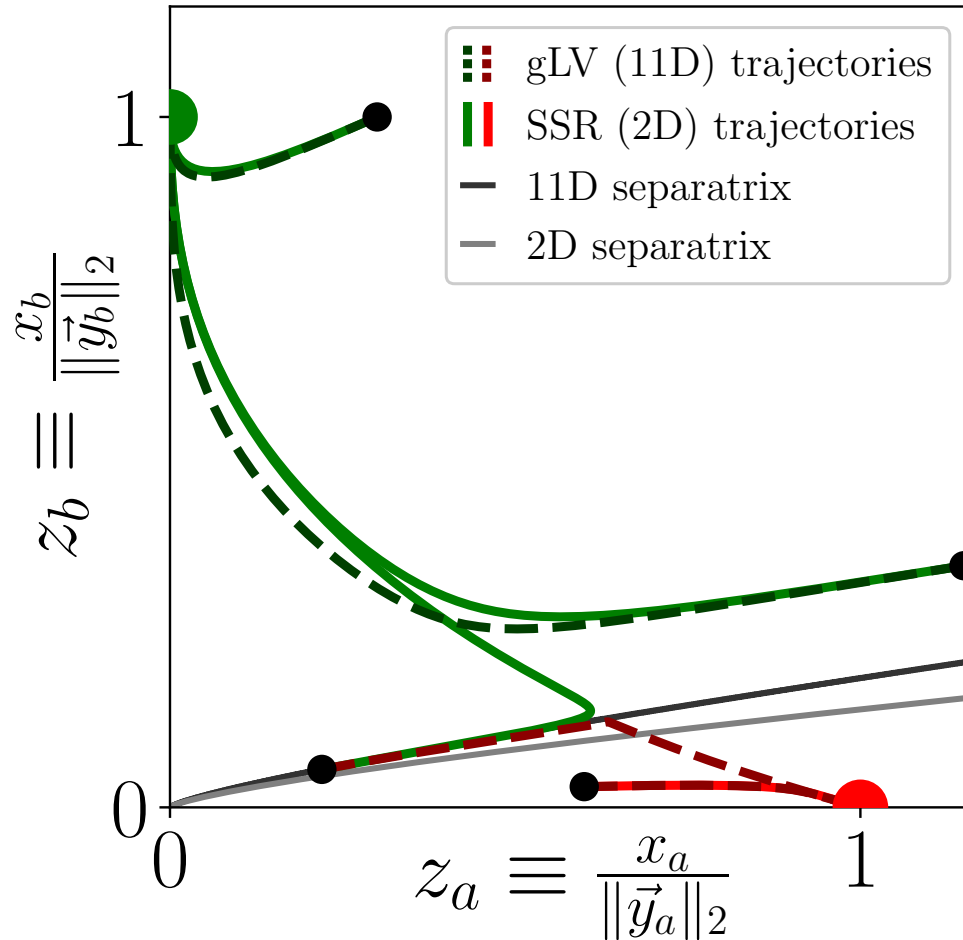


Figure 4.2: Fidelity of steady-state reduction (SSR). SSR is applied to an experimentally-derived 11-dimensional (11D) gLV model of *C. difficile* infection (CDI) [7]. This model exhibits steady states \vec{y}_a and \vec{y}_b that are vulnerable (diseased, red) and resilient (healthy, green) to invasion by the pathogen *C. difficile*. We consider 11D microbial trajectories whose initial conditions lie the plane spanned by these two steady states, and plot the in-plane projection of these trajectories (dashed lines). The 2D SSR-generated dynamics (solid lines) are plotted alongside these high-dimensional trajectories. The separatrix of each system is also plotted: as a proxy for the 11D separatrix (actually a 10-dimensional surface), the in-plane separatrix (dark grey) is numerically generated and plotted; the 2D separatrix is exact and given in Eq. (4.6) (light grey). The code used to generate this figure is available in the Supplemental Code.

4.5 Analysis of the 2D gLV equations

Having demonstrated a method of linking a high-dimensional gLV system to a 2D gLV system via SSR, we now take advantage of the analytic accessibility of such 2D systems. We consider biologically relevant systems that exhibit competitive dynamics by assuming $\mu_\alpha > 0$ for $\alpha \in a, b$, and $M_{\alpha\beta} < 0$ for $\alpha, \beta \in a, b$. These systems exhibit two stable and homogeneous fixed points at $(-\mu_a/M_{aa}, 0)$ and $(0, -\mu_b/M_{bb})$. In this case, the system will also possess a hyperbolic fixed point at (x_a^*, x_b^*) with $x_a > 0$ and $x_b > 0$, which topologically guarantees the existence of a separatrix.

In Section 4.5.1 this separatrix is explicitly calculated for the 2D gLV system Eq. (5.3). This result, in conjunction with SSR, allows for an efficient approximation of the high-dimensional separatrix. Then, Section 4.5.2 explores the steady state and transient dynamics of a nondimensionalized form of the 2D gLV system with clinically-inspired modifications.

4.5.1 Explicit form of the separatrix

The long-term dynamics of this system are dictated by the basins of attraction of the stable steady states, and these basins are delineated by a separatrix that, for topological reasons, must be the stable manifold of the hyperbolic fixed point (x_a^*, x_b^*) . In Fig. 4.3 these basins are depicted topographically via isoclines of the split Lyapunov function $V(x_a, x_b)$ (lightly shaded contours), which acts as a potential energy landscape [105].

The separatrix $h(x_a)$ may be analytically computed in a power series expansion about the hyperbolic fixed point (x_a^*, x_b^*) ,

$$h(x_a) = \sum_{n=0}^{\infty} \frac{c_n}{n!} (x_a - x_a^*)^n, \quad (4.4)$$

which as an invariant manifold must satisfy [106]

$$\frac{dh(x_a)}{dx_a} = \frac{dx_b}{dt} \bigg/ \frac{dx_a}{dt}, \quad (4.5)$$

resulting in the recursive relations

$$\begin{aligned} c_0 &= x_b^*, \\ c_1 &= \frac{1}{2M_{ab}x_a^*} \left[M_{bb}x_b^* - M_{aa}x_a^* \right. \\ &\quad \left. - \sqrt{(M_{bb}x_b^* - M_{aa}x_a^*)^2 + 4M_{ab}M_{ba}x_a^*x_b^*} \right], \\ c_2 &= \frac{2c_1(M_{ba} + M_{bb}c_1 - M_{aa} - M_{ab}c_1)}{2x_a^*M_{aa} + 3x_a^*M_{ab}c_1 - M_{bb}x_b^*}, \quad \text{and} \\ c_n &= \frac{1}{(nx_a^*M_{aa} + (n+1)x_a^*M_{ab}c_1 - M_{bb}x_b^*)} \\ &\quad \times \left\{ nc_{n-1}(M_{ba} + M_{bb}c_1 - (n-1)(M_{aa} + M_{ab}c_1)) \right. \\ &\quad \left. + n! \sum_{\ell=2}^{n-1} \left[\frac{c_\ell}{\ell! (n-\ell)!} (M_{bb}c_{n-\ell} - (n-\ell)M_{ab}c_{n-\ell} \right. \right. \\ &\quad \left. \left. - x_a^*M_{ab}c_{n-\ell+1}) \right] \right\}, \quad \text{for } n > 2, \end{aligned} \quad (4.6)$$

as derived in Eqs. (S27-S38). This calculation allows the *a priori* classification of the fate of a given initial condition, without need for simulation. We note that this algebraic calculation of the separatrix is considerably faster than numerical methods that rely on relatively costly quadrature computations. Further, in conjunction with SSR, this analytic form offers an efficient approximation to the in-plane separatrix of high-dimensional systems.

4.5.2 Dynamical landscape of the 2D gLV equations

Next, we analyze a two-state implementation of the gLV equations that parallels the clinically-inspired scenario of antibiotic-induced CDI. In this scenario, antibiotics deplete a health-prone initial condition, requiring administration of FMT in order to recover, as in Fig. 4.3. FMT is implemented in the 2D gLV model by adding a transplant of size s composed of the healthy steady state $(0, 1)$ to an evolving microbial state at a specified time following administration of antibiotics.

We consider a nondimensionalized form of the gLV equations Eq. (5.3) and designate nondimensionalized variables with a tilde. Therapeutic interventions of antibiotics and FMT are included in this model in a manner consistent with previous approaches [7, 9]. In all, this clinically-inspired two-state gLV model is given by

$$\begin{aligned}
 \frac{d\tilde{x}_a}{dt} &= \tilde{x}_a(1 - \tilde{x}_a - \tilde{M}_{ab}\tilde{x}_b) \\
 &\quad + \tilde{x}_a\varepsilon_a u(t) + w_a\delta(t - t^*), \quad \text{and} \\
 \frac{d\tilde{x}_b}{dt} &= \tilde{x}_b(\mu_b - \tilde{M}_{ba}\tilde{x}_a - \tilde{x}_b) \\
 &\quad + \tilde{x}_b\varepsilon_b u(t) + w_b\delta(t - t^*),
 \end{aligned} \tag{4.7}$$

which includes optional antibiotic administration $u(t)$ operating with efficacy $\vec{\varepsilon}$, and optional FMT with transplant \vec{w} administered instantaneously at time t^* .

In the absence of antibiotics and FMT, the dynamical system Eq. (4.7) exhibits three nontrivial steady states at $(1, 0)$, $(0, \mu_b)$, and $(\tilde{x}_a^*, \tilde{x}_b^*) \equiv (\frac{1 - \tilde{M}_{ab}\tilde{\mu}_b}{1 - \tilde{M}_{ab}\tilde{M}_{ba}}, \frac{\tilde{\mu}_b - \tilde{M}_{ba}}{1 - \tilde{M}_{ab}\tilde{M}_{ba}})$. To simplify the presentation of our results in the main text we assume $\mu_b = 1$, though this assumption is relaxed in the Supplementary Information.

As before, suppose the variable \tilde{x}_a corresponds to a diseased state, and \tilde{x}_b corresponds to a healthy state. Also assume the transplant \vec{w} consists of exclusively healthy microbes so that $w_a = 0$. Figs. 4.3, 4.4, and 4.5 are generated with parameter values $\tilde{M}_{ab} = 1.167$




Figure 4.3: The success or failure of fecal microbiota transplantation (FMT) depends on the timing of its administration in a two-state gLV system (Eq. (4.7)). We consider a clinically-inspired scenario that parallels antibiotic-induced CDI. First, a health-prone initial condition (IC) is depleted by antibiotics (RX, orange). If FMT (brown) is administered shortly after the antibiotics, the treatment steers the composition to a healthy state (FMT success). If FMT administration is delayed, the microbial trajectory instead attains the diseased state (FMT failure). The basins of attraction of the healthy and diseased steady states are delineated by the separatrix Eq. 4.6, and isoclines of the potential energy landscape (light contours) are given by the split Lyapunov function Eq. (S47) [105].




Figure 4.4: The FMT transplant size needed to revert an antibiotic-depleted state back to health grows as FMT administration is delayed. The minimum FMT transplant size required to cure five distinct disease-prone microbial trajectories, each evolving according to Eq. (4.7), are calculated and plotted. As trajectories attain the diseased steady state, the required transplant size approaches s^* . The required transplant size changes at two different rates, α and β , with the crossover point between these two rates at size s_c indicated by a hollow square. The transplant size dynamics ds/dt as well as the rates α and β are derived in Eq. (4.9) and the surrounding text.

and $\tilde{M}_{ba} = 1.093$, which give typical results.

Altering the fate of an initial condition requires crossing the separatrix by some external means, which is achieved through FMT. Fig. 4.3 shows two microbial time courses in which long-term outcomes are determined by the timing of FMT administration. At each point along a microbial trajectory in the diseased basin of attraction, the minimum FMT size s required to transfer the microbial state into the healthy basin of attraction is calculated. We use this metric to quantify our notions of “FMT efficacy.”

In clinical practice FMT administration varies in transplant size, transplant composition, and how many transplants are performed. Further, it is unclear how these factors

influence the success of FMT [107]. For the purposes of this chapter, we consider a hypothetical FMT treatment of size s_t (i.e. a horizontal cut across Fig. 4.4) and describe how its success depends on the timing of its administration.

Fig. 4.4 presents the minimum FMT size s as a function of time (main panel) for several trajectories that originate in the diseased basin of attraction (inset), including the main trajectory of Fig. 4.3. This minimum required FMT size increases with time—often dramatically—and there are two discernible rates of increase, denoted α and β in Fig. 4.4. These two rates are related to the fast and slow manifolds of the ecological system, which in turn govern the minimum required transplant size dynamics over time.

To reflect the importance of the separatrix in dictating the microbial dynamics, we change coordinates to the eigenvectors (u, v) of the hyperbolic steady state, shown in Fig. 4.5 (inset). In these coordinates the v -axis corresponds to the separatrix, and u is proportional to the minimum FMT size required for a successful transplant s , such that $s = u/(\hat{u} \cdot \hat{x}_b)$, where (\hat{u}, \hat{v}) and (\hat{x}_a, \hat{x}_b) are the unit vectors associated with their associated coordinates.

In this new (u, v) basis, the 2D gLV equations become

$$\begin{aligned} \frac{du}{dt} &= A_{10}u - A_{11}uv - A_{20}u^2, \quad \text{and} \\ \frac{dv}{dt} &= -B_{01}v - B_{02}v^2 + B_{20}u^2, \end{aligned} \tag{4.8}$$

where each coefficient is a positive algebraic function of the original gLV parameters given analytically in Eqs. (S60-S74) of the Supplementary Information. When $\mu_b \neq 1$, these equations contain additional quadratic terms described in the Supplementary Information that account for the nonlinearity of the separatrix. In the small u and small v limit this model reduces to the linearization about the hyperbolic fixed point. Near this fixed point there is a separation of time scales between u and v ($B_{01}/A_{10} > 1$ always, with median




Figure 4.5: The role of timing in FMT administration. For antibiotic-depleted disease-prone initial conditions in which antibiotics have been cleared ($u(t) = 0$), FMT is most effective when administered immediately ($t_{\text{opt}}^* = 0$, grey) or nearly immediately ($t_{\text{opt}}^* > 0$, colored) following antibiotic administration. The optimal transplant time t_{opt}^* is computed for any initial condition (u_0, v_0) (colorbar) according to Eq. (S88) of the Supplementary Information, which can reduce to Eq. (4.10). Two representative microbial trajectories are plotted in (u, v) (main panel) and (x_a, x_b) (inset) coordinates. For $v_0 > 0$ four possible FMT transplants are shown, including the optimal one that occurs at $t_{\text{opt}}^* = 2.1$. For $v_0 < 0$ it is always best to administer FMT immediately following antibiotic administration.

of 5.9 and IQR of [2.7, 9.1] over random parameter draws), indicating that there are inherent fast and slow manifolds in this system.

This coordinate change also reveals the role of timing in FMT administration, since the minimum required transplant size s is precisely governed by Eq. (4.8), by proxy of u . To demonstrate this analytically, we consider an initial condition (u_0, v_0) that is located near the fast manifold in a system with clear separation of time scales, so that (i) $B_{20}u_0^2$ is negligible, (ii) $A_{10} \ll B_{01}$, and (iii) $B_{02}v_0^2 \ll B_{01}v_0$ (though this assumption is relaxed in Eq. (S87)). In this case, the dynamics in the fast \hat{v} direction are approximately $v(t) \approx v_0 e^{-B_{01}t}$, and the required transplant size dynamics reduce to

$$\frac{ds}{dt} = s (A_{10} - A_{20}(\hat{u} \cdot \hat{x}_b)s - A_{11}v_0 e^{-B_{01}t}). \quad (4.9)$$

Thus, the required transplant size rates α and β in Fig. 4.4 are approximately $\alpha = \left. \frac{ds(0)}{dt} \right|_{s=s_c}$, and $\beta = \left. \frac{ds(\infty)}{dt} \right|_{s=s_c}$, where s_c is the transplant size required at the crossover point between these rates (e.g. as shown in Fig. 4.4).

For an initial condition with $v_0 < 0$, which occurs in Fig. 4.3 when a nearly healthy state is depleted by antibiotics, $\alpha > \beta$. In this case the required transplant size monotonically increases until it attains s^* at the infected steady state, so it is best to administer FMT as soon as possible. Alternatively, when $v_0 > 0$, $\alpha < \beta$. When $A_{11}v_0$ is sufficiently large α becomes negative, which indicates there is a nonzero transplant time at which the required transplant size is minimized (corresponding to $\frac{ds}{dt} = 0$). The concave-up trajectories in Fig. 4.4 exhibit this optimal transplant time. For $v_0 > 0$ and under the same conditions for which Eq. (4.9) was derived, this optimal transplant time t^* is

$$t_{\text{opt}}^* = \frac{1}{B_{01}} \ln \left(\frac{A_{11}v_0}{A_{10} - A_{20}u_0} \right). \quad (4.10)$$




Figure 4.6: Transient dynamics are preserved under SSR. (inset) Microbial trajectories of the CDI model (in-plane projection, dashed) and the associated SSR-reduced model (solid) as in Fig. 4.2 are shown. (main panel) At each time along these trajectories, the minimum FMT size required to make the state health-prone is plotted, for a transplant made up of \vec{y}_b (11D, dashed) or $(0, 1)$ (2D, solid). Phase space is linked to the FMT size dynamics by indicating the time at which z_b begins to decrease with a solid square (2D) or diamond (11D) in both the inset and main figure. Since the required FMT size s is the distance between a state and the separatrix, the similarity between the two time courses of s indicates that SSR preserves transient dynamics.

This nonzero transplant time reflects ecological pressures that temporarily drive the system closer to the separatrix, overpowering the slow unstable manifold. Two trajectories that numerically recapitulate these two cases are shown in Fig. 4.5.

4.6 SSR applied to fecal microbiota transplantation

In Section 4.5, FMT restored a CDI-prone microbial state in a 2D gLV model. In previous work [9], we implemented FMT in the previously mentioned 11D CDI model [7] and observed similar success. Here, the behavior of FMT in the CDI model and in its SSR counterpart are shown to match closely, which indicates that SSR preserves

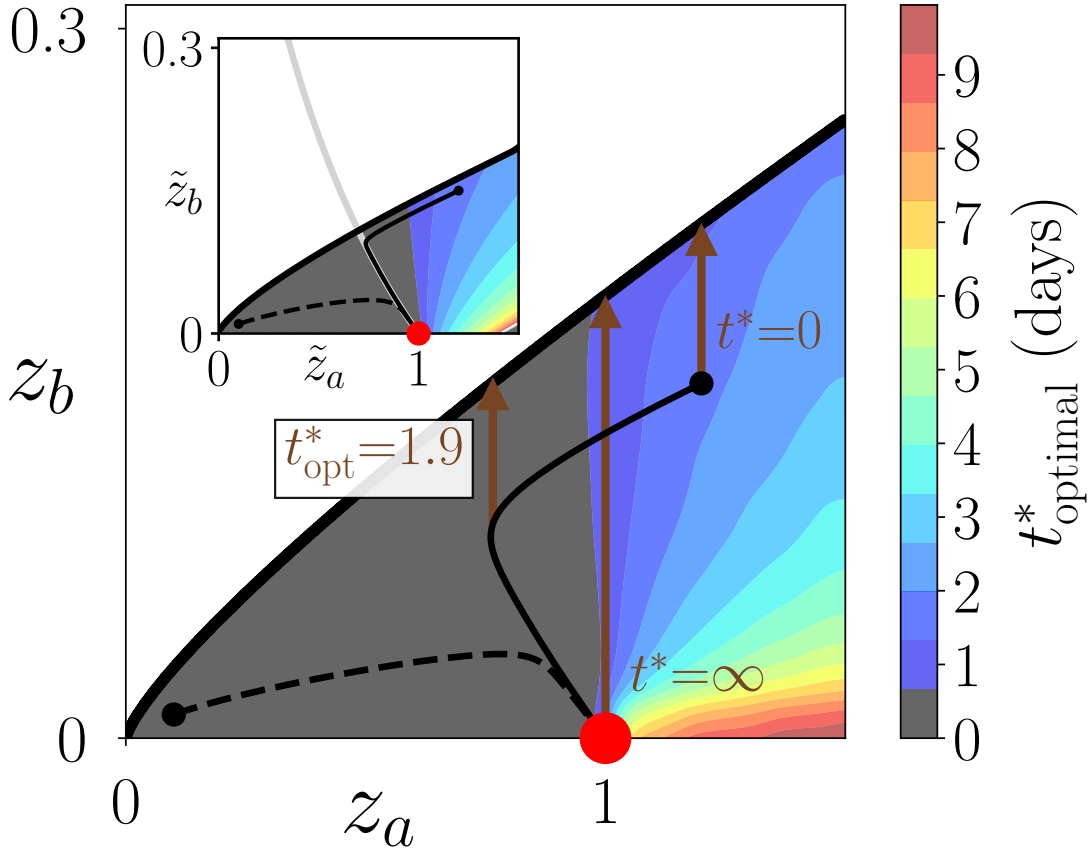


Figure 4.7: Optimal transplant times are preserved under SSR. Optimal transplant times t_{opt}^* of the 11D Stein model (main panel) largely match the predictions of its associated SSR-reduced model (inset). In the high-dimensional model, t_{opt}^* is computed numerically (as in Fig. 4.6) for a grid of points on the plane spanned by \vec{y}_a and \vec{y}_b for disease-prone initial conditions (located underneath the separatrix, which is shown as a thick black line). The spatial and temporal resolutions of this simulation are $\delta z_a = 0.025$, $\delta z_b = 0.01$, and $\delta t_{\text{opt}} = 0.15$, and the resulting data points were smoothed with a Gaussian filter. (inset) We display the optimal transplant times of the corresponding SSR-reduced model, as in Fig. 4.5. The SSR-reduced parameters were nondimensionalized so that t_{opt}^* could be generated with Eq. (S88), and the resulting optimal transplant time predictions were redimensionalized and plotted. The inset and the main panel share the same colorbar.

transient microbial dynamics.

Fig. 4.6 (inset) contains the in-plane projections of the 11D (dashed) and corresponding SSR-reduced 2D (solid) microbial trajectories with initial conditions that lie on the plane spanned by (\hat{y}_a, \hat{y}_b) (11D) and (\hat{z}_a, \hat{z}_b) (2D), as in Fig. 4.2. Fig. 4.6 (main panel) plots the required transplant size s at each state along the two trajectories: the 11D (dashed) transplant is composed of \hat{y}_a , and s is calculated numerically with a bisection method; the 2D (solid) transplant is composed of $\hat{z}_b = (0, 1)$, and s is computed analytically with Eq. (4.6).

In both systems, the microbial trajectories follow a fast stable manifold before switching to a slow manifold of some hyperbolic fixed point. As in the 2D case, the flow along these fast and slow manifolds underpins how the required transplant size s changes over time. In Fig. 6, the transition between the fast and slow manifolds occurs at 8.37 days in 11D (solid diamond, main panel and inset) and at 8.31 days in 2D (solid square).

As in the 2D case, the transition between these manifolds may result in a nonzero optimal transplant time t_{opt}^* . The main panel of Fig. 4.7 displays these optimal transplant times over a range of initial conditions, in which t_{opt}^* is generated with the same numerical bisection method as previously mentioned. Many of the high-dimensional initial conditions exhibit a non-zero optimal transplant time, mirroring the results of Fig. 4.5. Further, the high-dimensional optimal transplant times closely match those predicted by SSR, which are displayed in the inset of Fig. 4.7, and which were analytically calculated with Eq. (S88).

Since the SSR-reduced system largely preserves the high-dimensional transplant time dynamics, and since in the 2D nondimensionalized system t_{opt}^* can be examined analytically, the high-dimensional optimal transplant times may be approximated in terms of the high-dimensional interaction parameters. First, for systems well-approximated by SSR, a nonzero optimal transplant time can only exist when $v_0 > 0$ — this tends to occur

when the size of the initial condition is larger than that of the steady state \vec{y}_b , and when its composition is similar to that of \vec{y}_b . For this class of initial conditions, t_{opt} will be smaller when the eigenvalues of the semistable fixed point (A_{10} and B_{01}) are larger, or in terms of the SSR-reduced parameters, when M_{ab}/M_{bb} and M_{ba}/M_{aa} are larger.

The similarities between the transient dynamics of the high-dimensional and 2D systems, as well as the correspondence in optimal transplant timings, suggest that the theoretical analyses of Section IV may inform more complex and highly-resolved systems.

4.7 Discussion

4.7.1 Compression of complex ecological systems

SSR differs from other model reduction techniques [108, 109] since it preserves core observable ecological features of the original model, namely steady states and their stabilities. The behavior of the model on the transition between two of these steady states is approximated by SSR. Though the implementations of SSR demonstrated in this chapter were accurate, in general the accuracy of SSR is not obvious *a priori*; therefore, in future work it is important to carefully examine the circumstances under which SSR is effective. When SSR is accurate, properties of the steady states the original model may also be extracted from this approximation—for example, the size of the basins of attraction in the approximate system can inform the robustness of a given state in the original system, and the separatrix of the reduced model can approximate the slow manifold on which dynamics evolve in the original model. The speed-up gained by leveraging the analytic tractability of these approximate systems highlights the utility of SSR.

Beyond applications to existing gLV models, SSR-based methods could create two-state gLV systems from raw microbial data by choosing basis vectors during the fitting

process that correspond to experimentally observed steady states [110]. The resulting models would describe interactions between steady states rather than between individual species, and would consist of fewer variables and parameters that have improved explanatory power. This perspective— which effectively changes the basis vectors of a gLV model from species to steady states— may inform the transitions between steady states in ecological models.

4.7.2 Simplification of gLV-based FMT frameworks

Bacteriotherapy is a promising frontier of medicine that relies on the notion that the microbiome’s composition can both influence and be influenced by disease. Then, the deliberate alteration of a dysbiotic microbiome, by FMT for example, might be a viable treatment option for a range of diseases [111, 112]. Since FMT does not contribute to antimicrobial resistance, it is an emerging alternative to antibiotics [113, 114]. Clinical studies continue to regularly identify new diseases that are treated by FMT [115, 116, 117].

In this chapter we examined a bistable two-state gLV model from a clinical perspective, in which interventions such as FMT or antibiotics altered the outcome of a microbial trajectory. The tractability of this two-state system allowed for an explicit understanding of how the efficacy of FMT is influenced by the timing of its administration following antibiotic treatment. In this model, delaying the administration of FMT in disease-prone microbiomes could lead to its failure. Modifying the time course of a treatment has innovated treatment strategies in cancer immunotherapy [118] and HIV vaccination [119], and the results of this two-state ecological model suggest that treatment timing may be relevant for bacteriotherapy as well.

Indeed, some circumstantial evidence exists that supports the predictions of the two-

state model, in which FMT efficacy is improved when administered shortly after antibiotics. Kang et al. [120] used a promising variant of FMT to induce seemingly long-term changes in the microbiome and symptoms of children with autism spectrum disorders. This FMT variant “Microbiota Transfer Therapy” first prescribed the antibiotic vancomycin for two weeks, then bowel cleaning, then a large FMT dose of Standardized Human Gut Microbiota, and finally two months of daily maintenance FMT doses. In their study, they intended for the efficacy of FMT to be improved by first clearing out the microbiome with antibiotics, which is consistent with the results of the 2D gLV system. However, future experiments are needed to quantitatively test the extent to which antibiotic-depleted states are receptive to FMT-like therapies.

4.8 Conclusion

Broadly, SSR realizes a progression towards the simplification of dynamical systems: while linearization approximates a dynamical system about a single steady state, SSR approximates a dynamical system about *two* steady states. We have shown that SSR produces the best possible in-plane 2D gLV approximation to high-dimensional gLV dynamics. Further, we have demonstrated the extent to which the 2D model captures the basins of attraction and transient dynamics of an experimentally derived model. In addition to the computational efficiency of this technique, which employs analytic results rather than expensive simulations, SSR builds an intuition for the high-dimensional system out of connected 2D cross-sections.

By approximating this complex and classic ecological model with analytically tractable ecological subspaces, SSR anchors a high-dimensional system to well-characterized 2D systems. Consequently, this technique offers to unravel the complicated landscapes that accompany complex systems and their behaviors.

4.9 Derivation of steady-state reduction

Consider an N -dimensional gLV system given by Eq. (4.1) that exhibits steady states \vec{y}_a and \vec{y}_b , with dynamics given by $\frac{d\vec{y}}{dt} = \sum_{i=1}^N \frac{dy_i}{dt} \hat{y}_i$. As in the main text, define variables x_a and x_b in the direction of the unit vectors $\hat{x}_a \equiv \vec{y}_a / \|\vec{y}_a\|_2$, and $\hat{x}_b \equiv \vec{y}_b / \|\vec{y}_b\|_2$, where $\|\cdot\|_k$ is the k -norm. Further consider the *in-plane* 2D gLV dynamics that exist on the plane spanned by \hat{x}_a and \hat{x}_b . Here, we prove that the parameters prescribed by steady-state reduction, given in Eq. (4.3), minimize the 2-norm of the deviation $\vec{\epsilon}$ between the high-dimensional and in-plane dynamics at every point on the plane.

Consider coefficients $\mathbf{c} = \{c_1, \dots, c_6\}$ that parameterize the 2D gLV equations,

$$\begin{aligned} \frac{dx_a}{dt} &= x_a (c_1 + c_2 x_a + c_3 x_b), \quad \text{and} \\ \frac{dx_b}{dt} &= x_b (c_4 + c_5 x_a + c_6 x_b), \end{aligned} \tag{4.11}$$

so that the in-plane dynamics are $\frac{d\vec{x}}{dt} = \frac{dx_a}{dt} \hat{x}_a + \frac{dx_b}{dt} \hat{x}_b$. Any point on this plane can be written $\vec{y} = \vec{y}_a x_a + \vec{y}_b x_b$.

The deviation between the high-dimensional and in-plane dynamics $\vec{\epsilon}$ is

$$\vec{\epsilon}(x_a, x_b) = \frac{d\vec{x}}{dt} - \frac{d\vec{y}}{dt}, \tag{4.12}$$

which is defined at every point on the plane (x_a, x_b) . We will show that the parameters prescribed by SSR minimize the 2-norm of this deviation $\|\vec{\epsilon}\|_2$ at point on the plane.

The deviation $\vec{\epsilon}$ can be decomposed into the N -dimensional unit vectors \hat{y}_i , so that

$\vec{\epsilon} = \sum_i \hat{y}_i \epsilon_i$, where the components ϵ_i are given by

$$\begin{aligned} \epsilon_i &= y_{ai} x_a \left((c_1 - \rho_i) + \left(c_2 - \sum_{j=1}^N K_{ij} y_{aj} \right) x_a + \left(c_3 - \sum_{j=1}^N K_{ij} y_{bj} \right) x_b \right) \\ &\quad + y_{bi} x_b \left((c_4 - \rho_i) + \left(c_5 - \sum_{j=1}^N K_{ij} y_{aj} \right) x_a + \left(c_6 - \sum_{j=1}^N K_{ij} y_{bj} \right) x_b \right) \\ &\equiv \epsilon_{10,i} x_a + \epsilon_{20,i} x_a^2 + \epsilon_{11,i} x_a x_b + \epsilon_{01,i} x_b + \epsilon_{02,i} x_b^2, \end{aligned} \quad (4.13)$$

where components $\epsilon_{jk,i}$ are defined to correspond to contributions by $x_a^j x_b^k$ terms. Here, y_{ai} corresponds to the i th component of the unit vector $\hat{x}_a \equiv \vec{y}_a / \|\vec{y}_a\|_2$. In the same way, the deviation vector may be decomposed according to

$$\vec{\epsilon} = \vec{\epsilon}_{10} x_a + \vec{\epsilon}_{20} x_a^2 + \vec{\epsilon}_{11} x_a x_b + \vec{\epsilon}_{01} x_b + \vec{\epsilon}_{02} x_b^2. \quad (4.14)$$

Minimizing this deviation at each point (x_a, x_b) is equivalent to minimizing each orthogonal contribution $\vec{\epsilon}_{jk}$. Each contribution is a function of one or two parameters ($\vec{\epsilon}_{10}(c_1)$, $\vec{\epsilon}_{20}(c_2)$, $\vec{\epsilon}_{01}(c_4)$, $\vec{\epsilon}_{02}(c_6)$, and $\vec{\epsilon}_{11}(c_3, c_5)$), which simplifies the minimization process.

We now find the set of optimal coefficients $\mathbf{c}^* = \{c_1^*, \dots, c_6^*\}$ that minimize the 2-norm of each contribution $\|\vec{\epsilon}_{jk}\|_2$. For convenience, we equivalently minimize the square of this 2-norm. The Hadamard square represents the element-wise square of a vector, defined as $\vec{v} \circ^2 = [v_1^2, v_2^2, \dots, v_N^2]^T$.

The coefficient $\|\vec{\epsilon}_{10}\|_2^2$ is given by

$$\|\vec{\epsilon}_{10}\|_2^2 = \sum_{i=1}^N y_{ai}^2 (c_1 - \rho_i)^2. \quad (4.15)$$

When minimized with respect to c_1 , this quantity obeys

$$\frac{d\|\vec{\epsilon}_{10}\|_2^2}{dc_1} = \sum_{i=1}^N 2y_{ai}^2(c_1 - \rho_i) = 0, \quad (4.16)$$

which is satisfied for

$$c_1^* = \frac{\sum_{i=1}^N y_{ai}^2 \rho_i}{\sum_{i=1}^N y_{ai}^2} = \frac{\vec{y}_a^{\circ 2} \cdot \vec{\rho}}{\|\vec{y}_a\|_2^2}. \quad (4.17)$$

In a similar way, $\|\vec{\epsilon}_{20}\|_2^2$, $\|\vec{\epsilon}_{01}\|_2^2$, and $\|\vec{\epsilon}_{02}\|_2^2$ are minimized when

$$c_2^* = \frac{\sum_{i=1}^N \left(y_{ai}^2 \sum_{j=1}^N K_{ij} y_{aj} \right)}{\sum_{i=1}^N y_{ai}^2} = \frac{(\vec{y}_a^{\circ 2})^T K \vec{y}_a}{\|\vec{y}_a\|_2^3}, \quad (4.18)$$

$$c_4^* = \frac{\sum_{i=1}^N y_{bi}^2 \rho_i}{\sum_{i=1}^N y_{bi}^2} = \frac{\vec{y}_b^{\circ 2} \cdot \vec{\rho}}{\|\vec{y}_b\|_2^2}, \quad (4.19)$$

and

$$c_6^* = \frac{\sum_{i=1}^N \left(y_{bi}^2 \sum_{j=1}^N K_{ij} y_{bj} \right)}{\sum_{i=1}^N y_{bi}^2} = \frac{(\vec{y}_b^{\circ 2})^T K \vec{y}_b}{\|\vec{y}_b\|_2^3}. \quad (4.20)$$

Lastly, the squared norm of the cross-term $\|\vec{\epsilon}_{11}\|_2$ is given by

$$\begin{aligned} \|\vec{\epsilon}_{11}\|_2^2 = \sum_{i=1}^N \left[y_{ai} \left(c_3 - \sum_{j=1}^N K_{ij} y_{bj} \right) \right. \\ \left. + y_{bi} \left(c_5 - \sum_{j=1}^N K_{ij} y_{aj} \right) \right]^2. \end{aligned} \quad (4.21)$$

Minimizing with respect to c_3 and c_5 results in

$$\begin{aligned} \frac{d\|\vec{\epsilon}_{11}\|_2^2}{dc_3} &= \sum_{i=1}^N 2 \left[y_{ai}^2 \left(c_3 - \sum_{j=1}^N K_{ij} y_{bj} \right) \right. \\ &\quad \left. + y_{ai} y_{bi} \left(c_5 - \sum_{j=1}^N K_{ij} y_{aj} \right) \right] \\ &= 0, \end{aligned} \quad (4.22)$$

and

$$\begin{aligned} \frac{d\|\vec{\epsilon}_{11}\|_2^2}{dc_5} &= \sum_{i=1}^N 2 \left[y_{ai} y_{bi} \left(c_3 - \sum_{j=1}^N K_{ij} y_{bj} \right) \right. \\ &\quad \left. + y_{bi}^2 \left(c_5 - \sum_{j=1}^N K_{ij} y_{aj} \right) \right] \\ &= 0. \end{aligned} \quad (4.23)$$

After rearranging terms, these conditions read

$$c_3 \sum_{i=1}^N y_{ai}^2 + c_5 \sum_{i=1}^N y_{ai} y_{bi} = \sum_{i=1}^N \sum_{j=1}^N y_{ai} K_{ij} (y_{ai} y_{bj} + y_{bi} y_{aj}), \quad (4.24)$$

and

$$c_3 \sum_{i=1}^N y_{ai} y_{bi} + c_5 \sum_{i=1}^N y_{bi}^2 = \sum_{i=1}^N \sum_{j=1}^N y_{bi} K_{ij} (y_{ai} y_{bj} + y_{bi} y_{aj}), \quad (4.25)$$

which are satisfied when

$$c_3^* = \frac{\sum_{i,j=1}^N K_{ij} (y_{ai} y_{bj} + y_{bi} y_{aj}) \left(\sum_{k=1}^N y_{ai} y_{bk}^2 - y_{bi} y_{ak} y_{bk} \right)}{\left(\sum_{i=1}^N y_{ai}^2 \right) \left(\sum_{i=1}^N y_{bi}^2 \right) - \left(\sum_{i=1}^N y_{ai} y_{bi} \right)^2}, \quad (4.26)$$

and

$$c_5^* = \frac{\sum_{i,j=1}^N K_{ij}(y_{ai}y_{bj} + y_{bi}y_{aj}) \left(\sum_{k=1}^N y_{bi}y_{ak}^2 - y_{ai}y_{ak}y_{bk} \right)}{\left(\sum_{i=1}^N y_{ai}^2 \right) \left(\sum_{i=1}^N y_{bi}^2 \right) - \left(\sum_{i=1}^N y_{ai}y_{bi} \right)^2}. \quad (4.27)$$

However, when \vec{y}_a and \vec{y}_b are orthogonal, the cross-term deviation $\|\vec{\epsilon}_{11}\|_2^2$ is simplified, and the optimal coefficients c_3^* and c_5^* become

$$c_3^* = \frac{\sum_{i=1}^N \left(y_{ai}^2 \sum_{j=1}^N K_{ij}y_{bj} \right)}{\sum_{i=1}^N y_{ai}^2} = \frac{(\vec{y}_a \circ^2)^T K \vec{y}_b}{\|\vec{y}_a\|_2^2 \|\vec{y}_b\|_2}, \quad (4.28)$$

and

$$c_5^* = \frac{\sum_{i=1}^N \left(y_{bi}^2 \sum_{j=1}^N K_{ij}y_{aj} \right)}{\sum_{i=1}^N y_{bi}^2} = \frac{(\vec{y}_b \circ^2)^T K \vec{y}_a}{\|\vec{y}_b\|_2^2 \|\vec{y}_a\|_2}. \quad (4.29)$$

Since the squared norms of the deviations $\|\epsilon_{jk}\|_2$ are convex, the coefficient set \mathbf{c}^* is a global minimum for $\|\vec{\epsilon}\|_2$. Therefore, we have identified the parameters that minimize the deviation between the in-plane and high-dimensional gLV dynamics for any point on the plane spanned by \vec{y}_a and \vec{y}_b .

Chapter 5

Control of ecological outcomes with direct bacteriotherapies

Having carefully examined the ecological dynamics between a single pair of steady states with SSR in chapter 4, in this chapter we generalize our approach to study the ecological dynamics between multiple pairs of steady states. We compress the dynamical landscape of a multistable high-dimensional gLV system (i.e., a gLV system that possesses several steady states that are each point attractors) into an “attractor network,” which contains information about the ecological dynamics between each possible pair of steady states. The attractor network is a graph, where nodes are steady states and edges correspond to ecological states that are convex combinations of the steady states. Along each edge we indicate where the separatrix between the two steady states is located. Therefore, the attractor network compresses the dynamical landscape of high-dimensional gLV systems into an intuitive schematic that reveals how to transition from one steady state to another.

The attractor network of the CDI model informs how inherent ecological dynamics can be leveraged to minimize the required size of successful direct bacteriotherapies. For example, our results indicate that when attempting to drive a system towards a particular

steady state, it can sometimes be more efficient (i.e. require a smaller net transplant size) to drive the system first towards an intermediate state, and then towards the target state. The attractor network therefore serves as a roadmap of the dynamical landscape, and informs the optimal composition of direct bacteriotherapies.

Finally, the attractor network is derived from the same perspective as SSR— that bistable subsystems provide useful information about the total system— but it can be constructed without using SSR. However, constructing an attractor network without SSR requires the high-dimensional gLV systems to be simulated for initial conditions along each convex combination of steady states, which is computationally expensive. Alternatively, SSR approximates the location of the separatrix on each edge instantly, which is useful since the number of edges of an attractor networks grows as the square of the number of steady states of interest. Therefore, due to the analytic properties of 2D gLV models, SSR holds great utility for efficiently mapping out high-dimensional gLV systems.

5.1 Abstract for “Navigation and control of outcomes in a generalized Lotka-Volterra model of the microbiome”

The generalized Lotka-Volterra (gLV) equations model the microbiome as a collection of interacting ecological species. Here we use a particular experimentally-derived gLV model of *C. difficile* infection (CDI) as a case study to generate methods that are applicable to generic gLV models. We examine how to transition gLV systems between multiple steady states through the application of direct control protocols, which alter the state of the system via the instantaneous addition or subtraction of microbial species.

In the CDI model, the state of the system is more easily altered by direct control protocols that introduce a foreign population composed of many species (e.g. a steady state composition of the system) rather than protocols that are made up of a single microbial species. Then, to better understand how to navigate the dynamical landscape of a gLV system, the geometry of the basins of attraction of point attractors is compressed into an *attractor network*, which decomposes a multistable high-dimensional landscape into web of bistable subsystems. This attractor network is used to identify the most efficient (total intervention volume minimizing) protocol to drive the system from one basin to another. When trying to drive a system towards a target state, sometimes the most efficient control protocol is circuitous and will take the system through intermediate states. Clinically, the efficient control of the microbiome has pertinent applications for bacteriotherapies, which seek to remedy microbiome-affiliated diseases by directly altering the composition of the gut microbiome.

5.2 Introduction

In this chapter we characterize the dynamics of the generalized Lotka-Volterra (gLV) equations, a set of nonlinear coupled differential equations that are traditionally used in theoretical ecology to study interacting populations. In particular, gLV models are examined in the context of the gut microbiome, which consists of an ensemble of microorganisms that inhabit the gastrointestinal tract. Our scope is restricted to ecological dynamics that relax towards attractors; in this case microbial dynamics can be construed as being generated by a psuedo-energy landscape similar to a Lyapunov function, in which minima of the landscape correspond to steady states of the system. Canonically, landscape-based descriptions of biological processes have been used to describe how cell fates are determined in Waddington's epigenetic landscape, and more recently they have

been employed to study the genetic landscape of gene regulatory networks [121]. To better understand these high-dimensional dynamics, the geometric structure of the basins of attraction of a system are compressed into an *attractor network*, which represents the full landscape of a multi-stable gLV system as a web of interconnected bistable landscapes. By mapping the structure of dynamical landscapes, attractor networks inform the control of steady-state outcomes in gLV systems.

The study of the microbiome is motivated by a desire to better understand the ecological dynamics that underlie microbiome behavior, which might advance the ability of clinicians to respond to microbiome-associated disorders and to improve microbiome health. Experimental evidence linking microbiome to host health is reviewed in section 5.5.1, and alternative mathematical models of the microbiome are summarized in section 5.5.2.

As a case study, we consider an experimentally-derived gLV model of *Clostridioides difficile* infection (CDI), and use attractor networks to inform how to navigate between stable microbial communities of the system. The system is manipulated with *direct interventions* that modify an existing microbial state by either introducing a foreign microbial population (referred to as a *transplant*) or by removing existing microbial species. These direct intervention are interpreted as numerical implementations of bacteriotherapies, and when they are successful these interventions drive a microbial state into the basin of attraction of a target state. Broadly, these results examine the ecological mechanisms that underlie the successful administration of bacteriotherapies, and inform how the intrinsic ecological dynamics of the microbiome might be harnessed to improve the efficacy of bacteriotherapies.

5.3 Background

5.3.1 Generalized Lotka-Volterra (gLV) models

In this chapter the generalized Lotka-Volterra (gLV) equations are treated as a mathematical proxy for the microbial dynamics of the gut microbiome. The gLV equations describe the dynamics of a microbial population x_i in a system of N total interacting populations as

$$\frac{dx_i}{dt} = x_i \left(\mu_i + \sum_{j=1}^N M_{ij} x_j \right), \quad i \in 1, \dots, N, \quad (5.1)$$

where the growth rate of species i is given by μ_i , and the pairwise interaction M_{ij} encodes the ecological effects of species j on species i .

A gLV system with N populations can exhibit up to 2^N steady states, where each steady state is specific to a distinct presence/absence combination of the N species. Motivated by clinical bacteriotherapies like fecal microbiota transplantation (FMT) that seek to drive a diseased microbiome towards a healthy composition, we consider how to control the steady-state outcomes of gLV systems with direct interventions that instantaneously supplement or deplete the microbial state of a system. Towards this end, as in previous work [7, 9] the gLV model is extended to include the instantaneous addition of a foreign microbial transplant \mathbf{v} at time t^* , as well as the administration of an antibiotic treatment $u(t)$, so that

$$\frac{dx_i}{dt} = x_i \left(\mu_i + \sum_{j=1}^N M_{ij} x_j \right) + v_i \delta(t - t^*) + u(t) \varepsilon_i x_i, \quad (5.2)$$

where v_i is the i th component of the foreign transplant \mathbf{v} , $\delta(\cdot)$ is the Dirac delta function, and ε_i is the antibiotic susceptibility of population i .

Three types of bacteriotherapy-inspired direct interventions are examined, corre-




Figure 5.1: **Pairwise interactions between bacterial populations interpreted as a microbial food web.** An arrow from population j to population i represents the effect of j on the growth of i , as described by the interaction term M_{ij} in a generalized Lotka-Volterra model Eq (5.1). The width and opacity of an arrow are proportional to $|M_{ij}|$, and positive interactions ($M_{ij} > 0$) are green while inhibitory interactions ($M_{ij} < 0$) are red. The pairwise interactions M_{ij} here were fit from an experimental mouse model of *C. difficile* infection (CDI) [7, 6]. *C. difficile*, the culprit behind CDI, is colored red and located in the center of the food web. This figure and caption are adapted from [9].

sponding to three types of transplant \mathbf{v} : the addition of a single species (interpreted as a probiotic), the addition of a stable external microbial community (interpreted as fecal microbial transplant), and the removal of a particular species (interpreted as a phage therapy). More explicitly, we consider the direct control problem in which an initial condition is in the basin of attraction of some attractor \mathbf{x}_a , and the goal is to identify the smallest intervention \mathbf{v} that is able to drive the system into the basin of attraction of some other target state \mathbf{x}_b .

5.3.2 Experimentally-derived gLV model of *C. difficile* infection

We use a gLV model of *C. difficile* infection (CDI) as a case study for how to transition between basins of attraction. In this model, the growth rates μ and interaction parameters \mathbf{M} were fit by Stein *et al.* to microbial abundance time-series data from a mouse experiment performed by Buffie *et al.* [7, 6]. To reduce the number of parameters of the model, Stein *et al.* assumed that bacteria within a given genus behave similarly, and consolidated the species-level data into genus-level data. In Fig. 5.1 the interaction parameters are displayed as a food web where the circles are microbial populations, and the edges describe the effect of one population on another, where positive interactions ($M_{ij} > 0$) are black and inhibitory interactions ($M_{ji} < 0$) are red. The interactions between populations have no clear hierarchy, indicating that the fitted model describes microbes on the same trophic level competing for shared resources. The fitted antibiotic susceptibilities ε_i were mostly negative, indicating that antibiotics tend to deplete the growth rates microbial populations.

The CDI model produces microbial trajectories that allow for the simulated application of medical interventions. Previous work explored the effects of antibiotics on

microbial dynamics [7, 9] in this CDI model, and found that the CDI model exhibits the clinically- and experimentally-observed outcome that antibiotic-treated microbiomes were vulnerable to CDI.

To demonstrate antibiotic-induced CDI in the gLV model, Fig. 5.2 shows simulated microbial trajectories that result from applying four separate intervention scenarios to an initial condition measured by Buffie *et al.* For clarity, these figures plot the total microbe count on a log scale (in which the total microbe count is the sum of all of the microbes in each microbial population), and at each time each microbial population is linearly colored according to its proportion at that time. First, in Fig. 5.2a the system is not perturbed and the system evolves to a “healthy” steady state, in the sense that CD is unable to invade this steady state. In Fig. 5.2b the system is exposed to a unit dose of antibiotics (i.e., $u(t) = 1$ for $0 \leq t \leq 1$, and $u(t) = 0$ otherwise) which drives the system to an “antibiotic-depleted” state, in the sense that CD is able to invade it. Bearing this out, in Fig. 5.2c the system is exposed to a unit dose of antibiotics and then inoculated with CD on day 10 (i.e. the transplant \mathbf{v} is a unit vector of CD applied at time $t^* = 10$), and this system evolves towards an “CD infected” steady state in which CD is present. Finally, in Fig. 5.2d the system is exposed to a unit dose of antibiotics, inoculated with CD on day 1, and then also supplemented with an external foreign transplant (i.e. the foreign population \mathbf{v} is composed of an experimentally-measured initial condition introduced at a time $t^* = 1$); due to the direct intervention, this system now evolves towards the healthy steady state.

In each panel of Fig. 5.2 the colored mice represent attained steady state microbiome compositions: green represents healthy, yellow represents antibiotic-depleted, and red represents CD-infected. These three steady states of the CDI model resemble the compositions of the experimentally-observed mouse microbiome compositions, including their susceptibility or resilience to CD exposure [6, 7]. In addition to these three steady

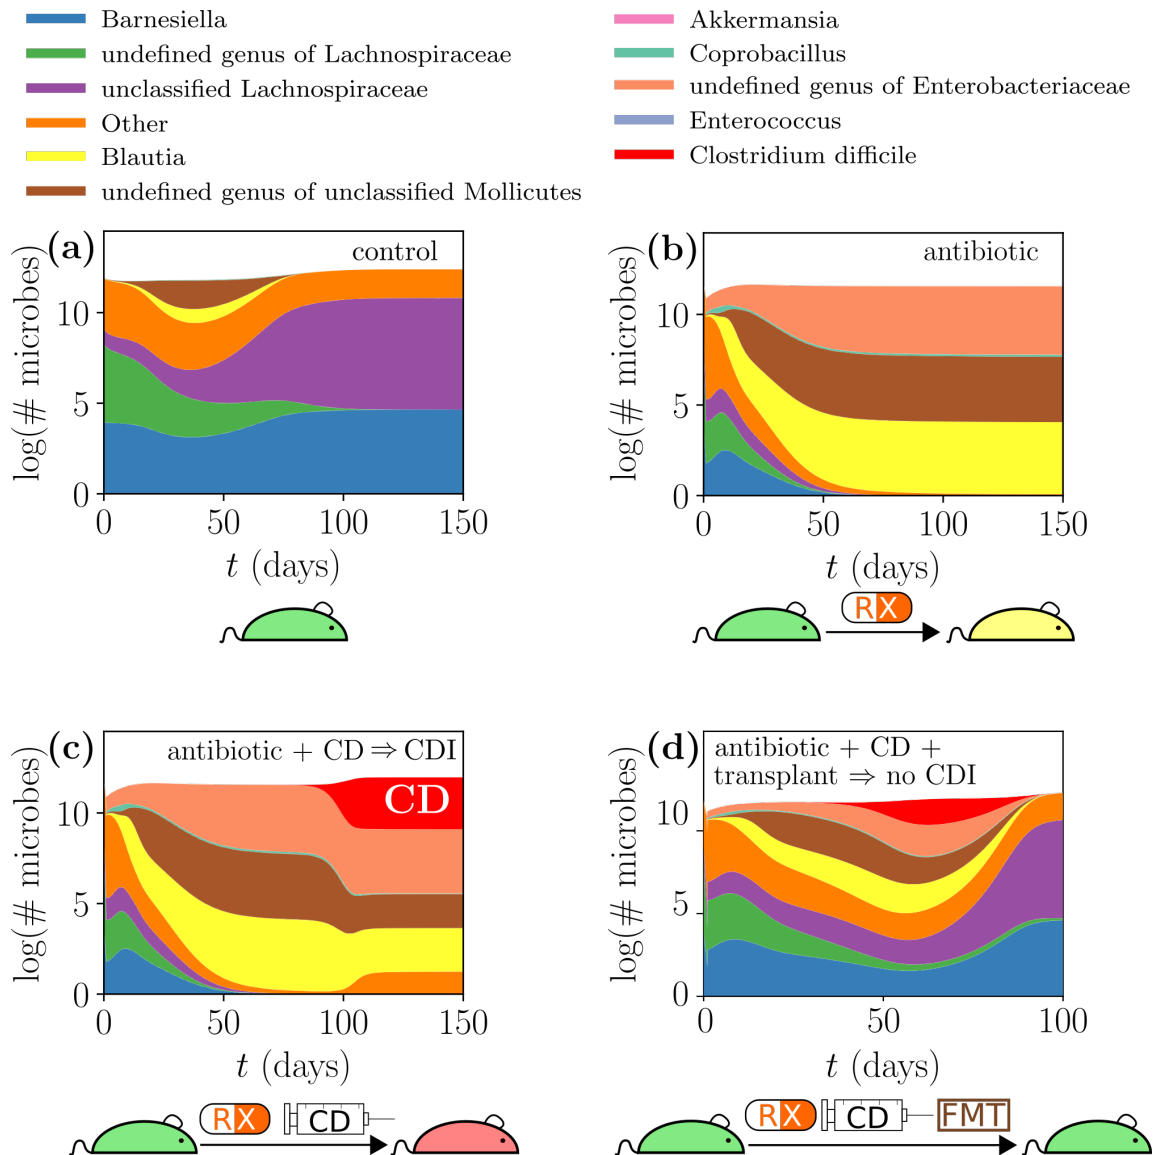


Figure 5.2: **External interventions can alter the steady-state outcome of an initial condition.** All panels originate from the same experimentally-measured initial condition, but different panels correspond to different interventions: (a) no interventions occur; (b) one dose of antibiotics administered at day 0; (c) one dose of antibiotic administered at day 0, and CD inoculation on day 10; and (d) one dose of antibiotic administered at day 0, CD inoculation on day 1, followed by the immediate introduction of a foreign microbial population made-up of a stable microbial community on day 1. Together these panels demonstrate that antibiotic-induced CDI may be remedied by the administration of a direct intervention, as in fecal microbiota transplantation. This figure and caption are adapted from [9].

states, two other steady states can be reached by initializing the system at one of the nine experimentally-measured initial conditions and then applying some type of intervention. We call this set of five steady states the “reachable” steady states of the CDI model, and focus on the ecological dynamics nearby them. The microbial compositions of these five steady states are displayed in Fig. 5.3. In this figure, the previously-mentioned health state is labeled steady state C, the antibiotic-depleted state is labeled steady state E, and the CD-infected state is labeled steady state D.

This experimentally-characterized CDI model provides a clinically-motivated case study in which different steady states can be associated with biologically meaningful microbiome compositions. It is pertinent to be able to efficiently switch basins of attraction in order to attain a “healthy” state, and in the remainder of this chapter we investigate this goal in detail.

5.3.3 Approximation of bistable gLV dynamics

The ecological dynamics of high-dimensional gLV systems are straightforward to simulate, but difficult to investigate analytically. Often the most fundamental features of a model are captured not by the particular high-dimensional vector that describe the state of a system, but rather by an abstract biological outcome associated with the state of a system. For example, in the case of the CDI model, the most important question is whether the system will tend towards a healthy (steady state C) or an antibiotic-depleted state (steady state E), and the precise composition of those steady states is not as important. Thus, in some sense there is a low-dimensional description of the transition between a pair of steady states, even as the actual ecological dynamics flow in a high-dimensional state space.

In earlier work we exploited this abstract outcome-oriented perspective to design

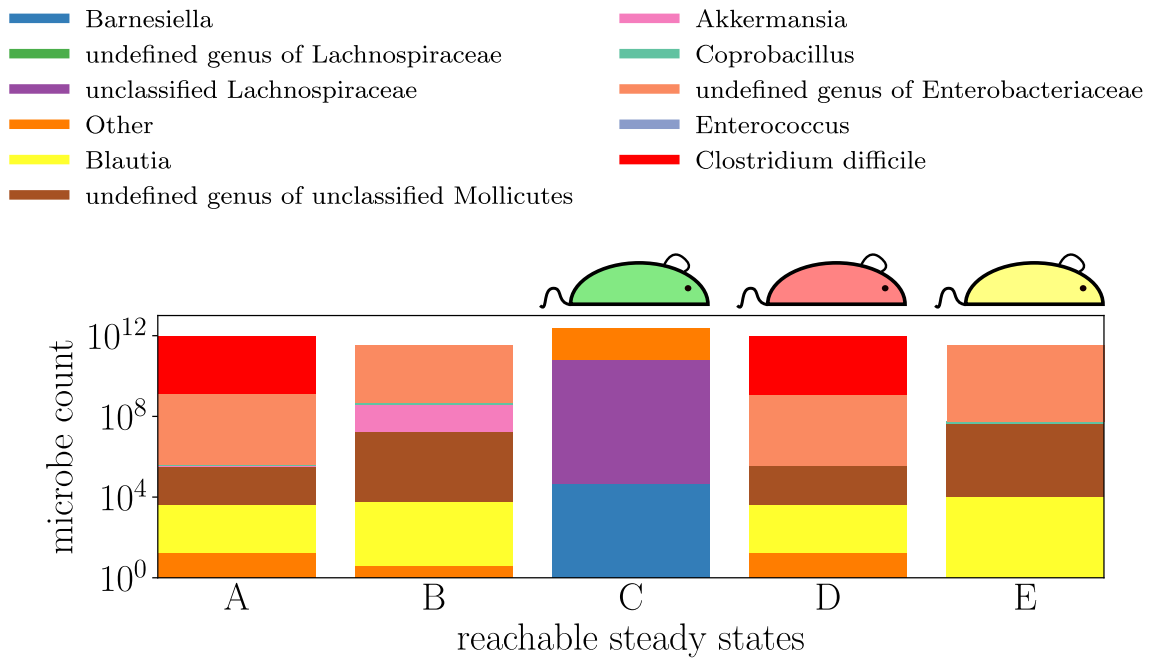


Figure 5.3: **Microbial composition of reachable steady states.** Under the gLV model Eq. (5.1) and for each of the nine experimentally-measured initial conditions, every treatment scenario tested in this chapter resulted in one of the steady states A-E. Together these five steady states encompass a region of phase space that is relevant for systems originating near the nine measured initial conditions. Note that while steady states A and D appear indistinguishable in this plot, their compositions vary slightly. This figure and caption are adapted from [9].




Figure 5.4: **Schematic of steady-state reduction (SSR).** A gLV system of N species (Eq. (5.1)) exhibits two steady states \vec{y}_a and \vec{y}_b , characterized as diseased (red) and healthy (green). SSR identifies the two-dimensional (2D) gLV system defined on the 2D subspace spanned by the two high-dimensional steady states (Eq. (5.3)) that best approximates the high-dimensional system. Specifically, SSR prescribes 2D parameters (Eqs. (5.4) and (5.5)) that minimize the deviation between the N -dimensional gLV dynamics $d\vec{y}/dt$ and the embedded 2D SSR-reduced dynamics $d\vec{x}/dt$. This figure and caption are adapted from [10].

steady-state reduction (SSR) [10]. This method compresses a bistable region of a high-dimensional gLV model into a reduced two-state gLV model whose two unit species represent a pair of steady states of the original model. As depicted in Fig. 5.4, this reduced two-dimensional (2D) gLV model is defined on the 2D subspace spanned by a pair of steady states of the original model, and the subspace itself is embedded within the high-dimensional ecological phase space of the original gLV model. The parameters of the reduced model are weighted combinations of the parameters of the original model, with weights that are related to the composition of the two high-dimensional steady states. Within this subspace, these reduced dynamics constitute the best possible 2D gLV approximation of the high-dimensional gLV dynamics. Though gLV systems in general are capable of displaying periodic or chaotic behaviors, here our attention is restricted to trajectories that approach a fixed point, and in particular to regions of phase space that are nearby the interface of the basins of attraction of a pair of steady states (i.e., nearby the separatrix).

More explicitly, to determine the SSR-reduced 2D gLV system associated with a bistable high-dimensional gLV model, first define variables x_a and x_b in the direction of unit vectors \hat{x}_a and \hat{x}_b that parallel the two steady states according to $\hat{x}_a \equiv \vec{y}_a / \|\vec{y}_a\|_2$, and $\hat{x}_b \equiv \vec{y}_b / \|\vec{y}_b\|_2$, where $\|\cdot\|_2$ is the 2-norm. The 2D gLV dynamics on the subspace spanned by \hat{x}_a and \hat{x}_b are given by

$$\begin{aligned} \frac{dx_a}{dt} &= x_a(\mu_a + M_{aa}x_a + M_{ab}x_b), \quad \text{and} \\ \frac{dx_b}{dt} &= x_b(\mu_b + M_{ba}x_a + M_{bb}x_b). \end{aligned} \tag{5.3}$$

Here, the *in-plane dynamics* on this subspace in vector form are written $\frac{d\vec{x}}{dt} = \frac{dx_a}{dt} \hat{x}_a + \frac{dx_b}{dt} \hat{x}_b$.

The 2D parameters generated by SSR are chosen to minimize the deviation between

the in-plane and high-dimensional gLV dynamics $\epsilon = \|\frac{d\vec{y}}{dt} - \frac{d\vec{x}}{dt}\|_2$ for any point on the subspace spanned by \hat{x}_a and \hat{x}_b . The values of the SSR parameters that minimize ϵ are derived in [10]. When the two steady states \vec{y}_a and \vec{y}_b are orthogonal, the 2D parameters are given by

$$\begin{aligned}\mu_\gamma &= \frac{\bar{\rho} \cdot \vec{y}_\gamma^{\circ 2}}{\|\vec{y}_\gamma\|_2^2}, \text{ and} \\ M_{\gamma\delta} &= \frac{(\vec{y}_\gamma^{\circ 2})^T K \vec{y}_\delta}{\|\vec{y}_\gamma\|_2^2 \|\vec{y}_\delta\|_2},\end{aligned}\tag{5.4}$$

where $\vec{y}^{\circ 2} \equiv \text{diag}(\vec{y})\vec{y}$ is the element-wise square of \vec{y} . When \vec{y}_a and \vec{y}_b are not orthogonal the interspecies interaction terms become slightly more complicated and are given by

$$\begin{aligned}M_{ab} &= \frac{\sum_{i,j=1}^N K_{ij}(y_{ai}y_{bj} + y_{bi}y_{aj})(\sum_{k=1}^N y_{bi}y_{ak}^2 - y_{bi}y_{ak}y_{bk})}{(\sum_{i=1}^N y_{ai}^2)(\sum_{i=1}^N y_{bi}^2) - (\sum_{i=1}^N y_{ai}y_{bi})^2}, \text{ and} \\ M_{ba} &= \frac{\sum_{i,j=1}^N K_{ij}(y_{bi}y_{aj} + y_{ai}y_{bj})(\sum_{k=1}^N y_{ai}y_{bk}^2 - y_{ai}y_{bk}y_{ak})}{(\sum_{i=1}^N y_{bi}^2)(\sum_{i=1}^N y_{ai}^2) - (\sum_{i=1}^N y_{bi}y_{ai})^2},\end{aligned}\tag{5.5}$$

where $\gamma, \delta \in a, b$, and y_{ai} is the i th element of the steady state \vec{y}_a . Under this construction, the high-dimensional steady states \vec{y}_a and \vec{y}_b have in-plane steady state counterparts at $(\|\vec{y}_a\|_2, 0)$ and $(0, \|\vec{y}_b\|_2)$, respectively.

Crucially, this reduced 2D gLV system is analytically tractable: the separatrix can be written analytically [10], and bifurcation analyses readily inform how interaction parameters alter the basins of attraction of a system [12]. SSR associates complex high-dimensional dynamics with an intuitive low-dimensional system. In the context of the CDI model, this SSR-based understanding informs the transition between e.g. the healthy “green” and antibiotic-depleted “yellow” microbial states, and eschews the high-dimensional details. Later, we will decompose multistable gLV systems into a web of bistable subsystems, then use this web to describe the geometry of the basins of attraction of steady states of the model as an *attractor network*. Further, SSR will be leveraged

to analytically compute the location of the separatrix in each subsystem, which is much more efficient than computing the locations of the separatrices numerically.

5.3.4 Transplant size and timing affect the efficacy of direct interventions

It is difficult to discern the influence of transplant size and timing on transplant efficacy in high-dimensional gLV systems, since only numerical methods are available to probe this relationship. However, SSR provides a link between high-dimensional and reduced 2D systems, allowing the mechanisms that underlie direct interventions in high-dimensional systems to be understood in terms of their low-dimensional counterparts.

Fig. 5.5 demonstrates how variability in transplant size and timing affects steady-state outcomes in a clinically-motivated scenario of CDI in a 2D gLV system. Here, an initial condition is depleted by antibiotics (RX) to the extent that it enters a diseased basin of attraction. Immediately, the reduced system makes clear that the location of the separatrix is crucial in determining the outcome of a microbial state: successful direct interventions must cross the separatrix. The transplant size required for success depends on the composition of the transplant (in this figure transplants are composed entirely of \tilde{x}_b). Furthermore, the required transplant size is variable, and depends on the ecological dynamics of the system— in Fig. 5.5 a transplant of the same size and composition is administered at two different times, but the later transplant is unsuccessful.

SSR allows for the relationships between transplant size, composition, and timing to be examined analytically in terms of an intuitive low-dimensional system. Next, having characterized the operation of direct transplants in a two-dimensional system, we examine how direct transplants may be used to transition between steady states in the full context of a multistable gLV system.




Figure 5.5: **The success of direct interventions depends on the size and timing of the transplant.** We consider a clinically-inspired scenario that parallels antibiotic-induced CDI in the 2D gLV system Eq. (5.3). First, a health-prone initial condition (IC) is depleted by antibiotics (RX, orange). If a direct transplant (FMT, brown) is administered shortly after the antibiotics, the treatment steers the composition to a healthy state (FMT success). Also, if the direct transplant is large enough to cross the separatrix, the intervention will be successful. Alternatively, if FMT administration is delayed and the transplant is too small, the microbial trajectory will instead attain the diseased state (FMT failure). The basins of attraction of the healthy and diseased steady states are delineated by the separatrix, and the light contours depict isoclines of the potential energy landscape (given by a split Lyapunov function). This figure and caption are adapted from [10].

5.4 Results

5.4.1 Transplant compositions and their success rates

The ecological state of a system can be permanently altered by deliberately introducing some foreign microbial composition or by selectively depleting the population of a single species. Fig. 5.3 provides proof-of-concept that an ecological trajectory can be shifted from one basin of attraction to another by introducing a foreign microbial population \mathbf{v} , and Fig. 5.5 demonstrates the same phenomenon graphically in a reduced 2D gLV system. In this section, the influence of transplant composition on direct intervention efficacy is examined using the high-dimensional CDI model.

Different types of direct interventions vary in their ability to switch a trajectory between basins of attraction. Three intervention types are considered: the introduction of a single-species “probiotic,” a stable community via “fecal microbiota transplantation” (FMT), or the elimination of a single species via “phage therapy.” These interventions act as *in silico* proxies for medical therapies. The success rates of these interventions are plotted as a function of the intervention magnitudes in Fig. 5.6. As an example, the transplant administered in Fig. 5.2d is considered a “successful” intervention, since it was able to alter a trajectory tending towards steady state D (the CD-infected steady state) and drive it towards steady state C (the healthy steady state).

Fecal microbiota transplantation is implemented in the model by setting the transplant \mathbf{v} proportional to a steady state of the gLV system. In particular, transplants \mathbf{v} are set proportional to one of the five reachable steady states depicted in Fig. 5.3, or they are set proportional to one of the nine experimentally-measured initial conditions from the mouse experiment performed by Buffie *et al.* [6]. Probiotics are realized by setting the transplant \mathbf{v} proportional to a single microbial species. Lastly, phage therapies are described by making the transplant \mathbf{v} negative, and setting it proportional to a single

microbial species (so that it depletes a particular species).

In Fig. 5.6 these interventions are applied to initial conditions located at the reachable steady states of the CDI model depicted in Fig. 5.3. Interventions are considered successful when they shift the basin of attraction of an initial condition, and for each type of intervention several transplants are considered. For the single species “probiotic,” transplants are solely composed of one of the 11 bacterial species of the model. The steady state “FMT” populations are composed of one of the five reachable steady states. There are 11 “phage” interventions that each deplete a single bacterial species. Lastly, the nine “experimental IC” foreign populations are proportional to the experimental initial conditions measured by Buffie *et al.* [6]. In Fig. 5.6 the success rates of each intervention are plotted according to their magnitude (i.e. the one-norm $\|\mathbf{v}\|_1$ of each foreign population). For scale, the five reachable steady states of the CDI model vary in size between 3×10^{11} and 24×10^{11} microbes, which informs the range of direct intervention sizes considered here.

As shown in Fig. 5.6, the success rates of multi-species interventions (steady states and experimental ICs) are significantly higher than the single-species interventions (single-species and phage therapies). In particular, phage therapies are complete ineffective at altering the basin of attraction that an initial condition is in. For each type of intervention, the success rates of each candidate transplant within each intervention type are plotted in dashed lines. The bulk success rates of each intervention are plotted in bold, and are computed by averaging the success rates of the individual introduced foreign populations within each intervention type.

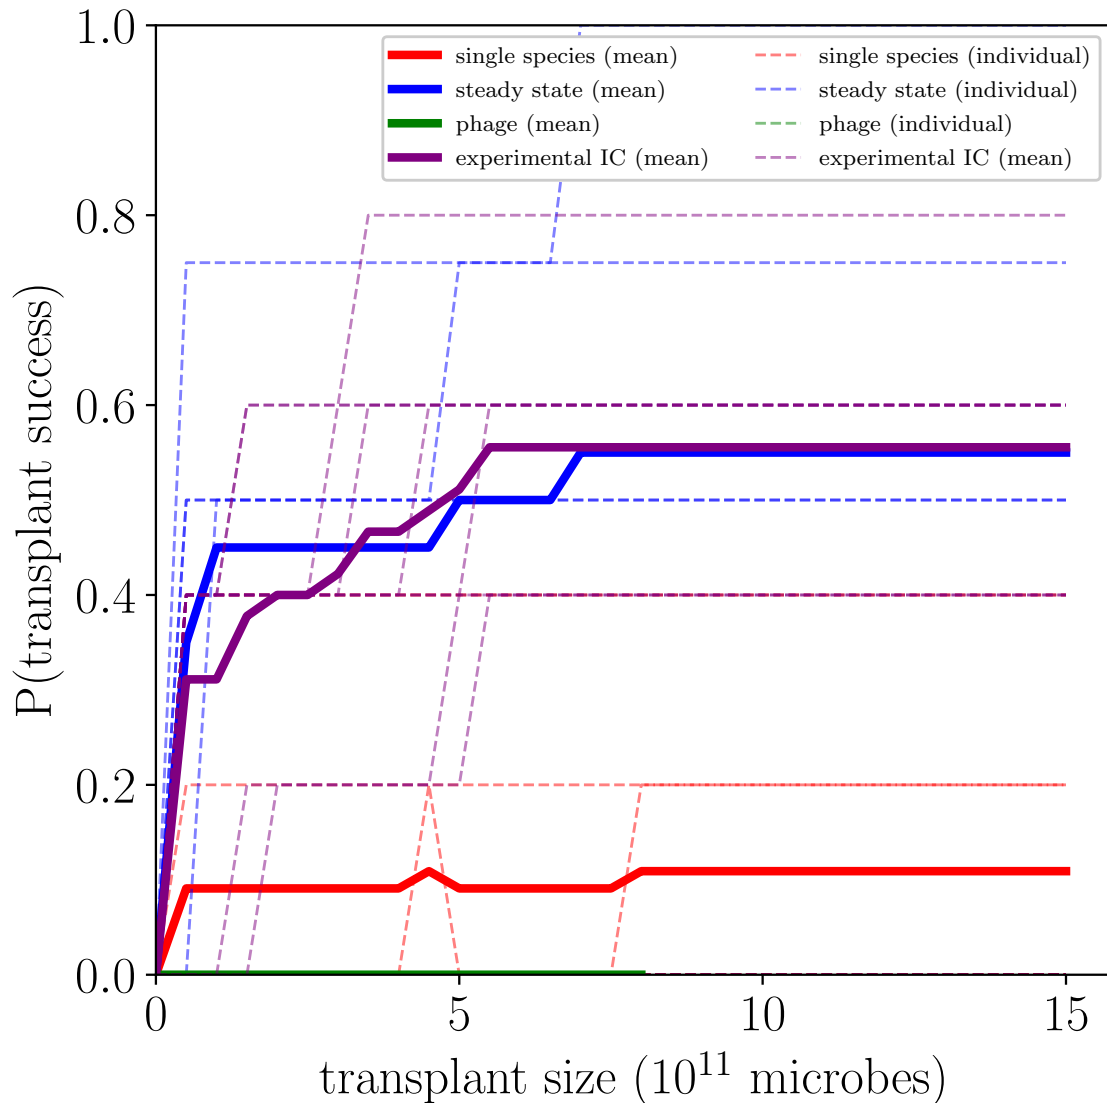


Figure 5.6: **Success rates of different microbial interventions at altering steady state outcomes.** Initial conditions were located one of the five steady states of the CDI model (displayed in Fig. 5.3) that were reached from experimentally-measured initial conditions. Then these initial conditions were subjected to the introduction of a foreign population \mathbf{v} over a range of sizes $\|\mathbf{v}\|_1$ (as shown on the x-axis). If the introduced foreign population drove the initial condition into a different basin of attraction, the intervention was considered “successful.” For each of the four types of intervention, several candidate transplants \mathbf{v} were implemented, and the success rates of each candidate intervention were plotted (dashed lines). A bulk success rate for each type of intervention (bold lines) was produced by averaging the success rates of the candidate transplants for each intervention type. Details about the candidate transplant compositions used for each type of intervention are described in the main text.

5.4.2 The dynamical landscape of a gLV system

The ecological dynamics of gLV systems are dictated by complex feedbacks between populations. Conceptually, in the absence of periodic or chaotic dynamics these ecological dynamics can be interpreted as flowing downhill a dynamical landscape (e.g. a Lyapunov function) towards a point attractor. A visualization of a dynamical landscape for a two-dimensional state space is displayed schematically in Fig. 5.7a. Here, the stable steady states A (yellow), B (blue), and C (green) are located at the minima of the landscape. Basins of attraction are displayed topographically and are the same color as their associated steady state.

In the CDI model, as demonstrated in Fig. 5.2 three distinct steady states can be reached from the same initial condition by administering different interventions (antibiotic administration and CD exposure). In earlier work we showed that the transition between these steady states is sudden as a function of the magnitude of the intervention (for example, an antibiotic dose of 0.70 units leads to steady state C, while a dose of 0.72 units leads to steady state E) [9]. Thus, the basins of attraction of these three steady states are touching, but the structure of this high-dimensional dynamical landscape is difficult to visualize. To rectify this, attractor networks can be used to compress information about the basins of attraction of the high-dimensional phase space into a visually-digestible form. Attractor networks were originally introduced by Wang *et al.* to study the controllability of gene regulatory networks associated with cancer [122].

In the schematic Fig. 5.7, panel (B) displays the attractor network for the dynamical landscape in panel (A). In the attractor network, nodes are steady states of the high-dimensional system, edges are convex combinations of pairs of steady states, and the colors along edges correspond to the basins of attraction along each convex combination.

Attractor networks are especially valuable for mapping the landscape nearby a few


(A) Dynamical Landscape **(B) Attractor Network**

Figure 5.7: **Schematic of a dynamical landscape and its corresponding attractor network.** (A) Ecological dynamics that are not periodic or chaotic may be represented as particles flowing on a dynamical landscape. Here, this artificial landscape exhibits three stable steady states at A (yellow), B (blue), and C (green). The basins of attraction are colored according to the color of their associated steady state, and isoclines of the landscape are also plotted. The black dotted line displays the value of the energy landscape along convex combinations of steady state pairs. (B) To compress this dynamical landscape into an attractor network, a graph is made in which nodes are steady states and edges are convex combinations of those steady states. The edge color for a particular convex combination corresponds to the basin of attraction of that ecological state. The small black circles in the attractor network correspond to the separatrixes of the dynamical landscape. The transplants \tilde{v}_1 , \tilde{v}_2 , and \tilde{v}_3 , whose compositions are given in the main text, demonstrate how direct interventions can alter the basin of attraction a state is in.

steady states in a high-dimensional system. For example, Fig. 5.8 presents an attractor network for the five reachable steady states of the CDI model (described in Fig. 5.3). The attractor network preserves geometric information about the basins of attraction of these five steady states, without requiring awkward visualization of an 11-dimensional state space.

The realization of antibiotic-induced CDI, as demonstrated in Fig. 5.2, indicated that (i) antibiotic administration shifted a microbial trajectory from the healthy steady state C towards the antibiotic-depleted steady state E, (ii) antibiotic exposure coupled with CD inoculation caused the microbial trajectory to flow towards the CD-infected steady state D, and (iii) trajectories in the basin of attraction of the CD-infected steady state D could be driven to the healthy steady state C through the introduction of a foreign population. The attractor network provides a compressed description of how an ecological state responds to external interventions, and complements and strengthens the numerical proof-of-concept analysis of antibiotic-induced CDI in Fig. 5.2.

5.4.3 Efficient navigation of an attractor network

The attractor network acts as a “roadmap” for a dynamical landscape, places microbial trajectories in the context of other nearby attractors, and can guide the administration of interventions in a gLV model. In discussing how to navigate gLV systems, we will assume that we are able to drive the state of the system along convex combinations of pairs of steady states.

Explicitly, represent steady states A, B, and C in Fig. 5.7 by the vectors \vec{x}_A , \vec{x}_B , and \vec{x}_C , respectively. Then, starting from steady state i , a direct intervention in the direction of steady states j corresponds to a transplant $\tilde{\vec{v}} = p(\vec{x}_j - \vec{x}_i)$ in the gLV equations Eq. 5.2, where p varies between 0 and 1 and describes the severity of the intervention.




Figure 5.8: **Attractor network of the five reachable steady states in the CDI model.** The small black circles on each edge represent the numerically-generated separatrices of each convex combination of steady states, which denote the boundary of the basins of attraction for a pair of steady states. The relationships between steady states C (healthy), D (CD-infected), and E (antibiotic-depleted) are explored numerically in Fig. 5.2. Representing the basins of attraction of these reachable steady states as an attractor network allows for an intuitive low-dimensional description of the dynamical landscape of the high-dimensional CDI system.

For example, in the schematic attractor network Fig. 5.7b, starting from steady state A an introduced transplant $\tilde{v}_1 = 0.3(\vec{x}_B - \vec{x}_A)$ will drive the state of the system into the basin of attraction of steady state B.

Explicitly, we seek to minimize the size of the intervention required to drive a system from an initial steady state into the basin of attraction of a target steady state $\|\tilde{v}\|$. Using the same schematic attractor network from Fig. 5.7b, driving an initial condition at steady state B into the basin of attraction of steady state C will require a transplant $\tilde{v}_2 = 0.3(\vec{x}_C - \vec{x}_B)$. Then, assuming that the distances between each pair of steady states are one unit, applying these two interventions in a sequential manner will cause an initial condition at steady state A to transition into steady state C with a total intervention size of 0.6. Alternatively, to drive steady state A to steady state C directly requires a transplant $\tilde{v}_3 = 0.7(\vec{x}_C - \vec{x}_A)$, with a total intervention size of 0.7. Thus, constructing an attractor network demonstrates that inherent ecological dynamics can be leveraged to efficiently control the state of an ecological system.

In the CDI model, the attractor network in Fig. 5.8 reveals that scenarios exist in which the most efficient control method goes along indirect paths. In Fig. 5.9 the microbial time courses associated with two control protocols— one that goes along a direct path, and the other that goes along a circuitous path— are compared.

Explicitly, let the steady state i of the CDI model correspond to the vector \vec{y}_i . To transition from the antibiotic-depleted steady state E directly to the healthy steady state C requires a transplant $\vec{v}_1 = 0.162(\vec{y}_C - \vec{y}_E)$, with total transplant size $\|\vec{v}_1\| = 2.576$, where $\|\cdot\|$ is the 2-norm and each transplant is in units of 10^{11} microbes. A similar control protocol is plotted in Fig. 5.9a, which successfully drives the state of the system into the basin of attraction of steady state C (the plotted transplant \vec{w}_1 is chosen to be slightly larger than the required transplant \vec{v}_1 for demonstrative purposes).

However, in this CDI model it is most efficient to apply two sequential interventions

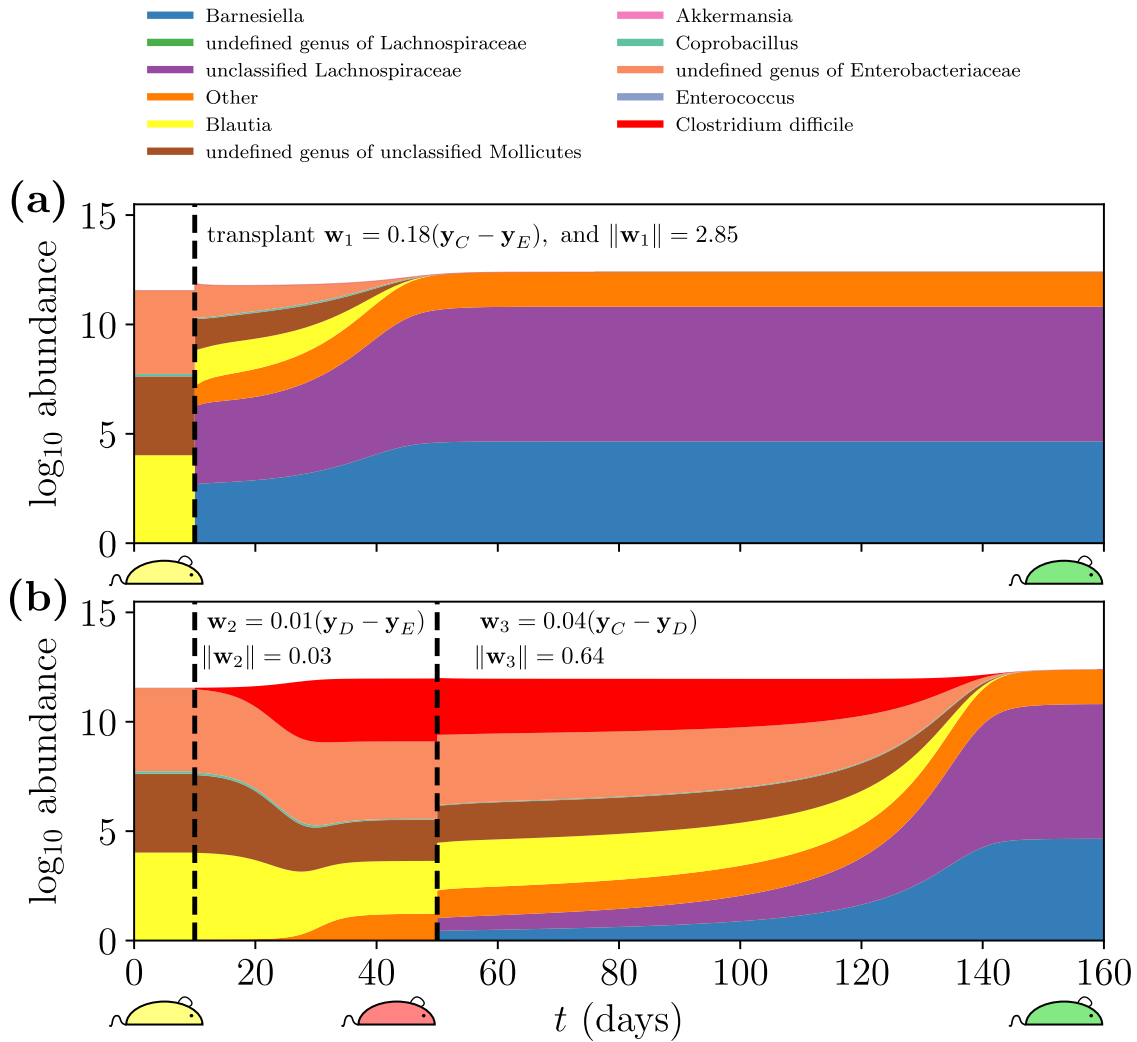


Figure 5.9: **Direct control protocols drive a microbial composition at steady state E (antibiotic-depleted, yellow mouse) towards steady state C (healthy, green mouse).** (a) The microbial trajectory associated with a control protocol that drives the system directly from steady state E to steady state C. This control protocol introduces a transplant $\vec{w}_1 = 0.18(\vec{y}_C - \vec{y}_E)$ on day 10. The introduced transplant is of size $\|\vec{w}_1\| = 2.85 \times 10^{11}$ microbes. (b) A circuitous control protocol drives the system at steady state E first towards steady state D (CD-infected, red mouse), and then applies a second transplant to drive the system to steady state D. By administering two smaller transplants sequentially, this control protocol requires fewer total microbes. The first transplant $\vec{w}_2 = 0.001(\vec{y}_D - \vec{y}_E)$ has a size of $\|\vec{w}_2\| = 0.03 \times 10^{11}$ microbes, and the second transplant $\vec{w}_3 = 0.04(\vec{y}_C - \vec{y}_D)$ has a size of $\|\vec{w}_3\| = 0.64 \times 10^{11}$ microbes. By navigating the dynamical landscape in a roundabout manner, the protocol in (b) requires total transplant size that is 23% the size of the protocol in (a).

\vec{v}_2 and \vec{v}_3 : the first transplant drives the system towards the CD-infected steady state D, $\vec{v}_2 = 0.001(\vec{y}_D - \vec{y}_E)$; and the second transplant drives the system towards the healthy steady state C, $\vec{v}_3 = 0.022(\vec{y}_C - \vec{y}_D)$. Here, since steady state E is unstable in the direction of steady state D, an infinitesimal transplant is all that is needed (though for practical purposes we use a value of 0.001). The sizes of these two sequential transplants are $\|v_2\| = 0.0001$ and $\|v_3\| = 0.360$. A similar sequential control protocol is demonstrated in Fig. 5.9b; once again, for demonstrative purposes the transplants \vec{w}_2 and \vec{w}_3 are chosen to be slightly larger than the required transplants \vec{v}_2 and \vec{v}_3 .

Taken together, the indirect control protocol $E \rightarrow D \rightarrow C$ (total transplant size 0.36×10^{11} microbes) requires a smaller total direct intervention than the direct path $E \rightarrow C$ (total transplant size 2.57×10^{11} microbes). In general, taking advantage of this “ecological inertia” may reduce the magnitude of the intervention required to drive the system towards a target state.

5.4.4 Efficient construction of attractor networks with steady-state reduction

Finally, attractor networks can be approximated in constant time (in algorithmic complexity terms) with the dimensionality-reduction technique SSR. The attractor network for the CDI model in Fig. 5.8 was generated numerically by simulating a set of microbial trajectories, each originating at initial conditions along convex combinations of steady state pairs, and then tracking the steady-state outcome of each simulation. This procedure is computationally expensive even when bisection-type algorithms are implemented to identify the location of the separatrix.

As derived in previous work, SSR allows for the separatrix of the approximate SSR-reduced 2D gLV system to be computed analytically in a Taylor series about a semi-




Figure 5.10: **Attractor network generated exactly (numerically) and approximately (with SSR).** This figure is identical to Fig. 5.8, but in addition to the small black circles corresponding to the numerically generated separatrix, the small black squares correspond to the separatrix locations as predicted by SSR. SSR erroneously predicts that any convex combination of steady states B and D will evolve towards steady state D, when in fact any convex combination of these steady states will evolve towards steady state A. Despite this error, SSR effectively and efficiently approximates the geometry of the dynamical landscape of the CDI gLV system.

stable fixed point (x_a^*, x_b^*) of the reduced system [10]. For example, the separatrix of the 2D gLV system in Fig. 5.5 was created in this fashion. Though the current attractor network is relatively small, consisting of five nodes and $\binom{5}{2}$ edges, scaling up the size of a numerically-generated attractor network quickly becomes computationally infeasible. On the other hand, larger attractor networks can be efficiently approximated with SSR: as a demonstration of the accuracy with which SSR captures the location of the separatrices, in Fig. 5.10 the squares on each edge represent the separatrix predicted by applying SSR to the corresponding steady state pair.

The SSR-generated attractor network incorrectly predicts that convex combinations of steady states B and D will flow towards steady state D, when in fact they will flow towards steady state A. Still, the strong agreement between the exact and SSR-generated separatrices indicates the potential of SSR to efficiently generate large attractor networks to first-order that intuitively describe relevant ecological dynamics.

5.5 Discussion

5.5.1 Microbiome composition is associated with host health

Microbiome research has been the recent beneficiary of advances in experimental 16S pyrosequencing techniques that have revealed similarities in the microbiome compositions of people suffering various diseases, which differ from the microbiome compositions of healthy individuals [123]. These disease-associated microbiome compositions are called *dysbiotic*, and are observed in people that suffer cardiovascular disease, ulcerative colitis, and irritable bowel disease [124, 123].

In this chapter we studied a mathematical model of *C. difficile* infection (CDI). CDI occurs when the bacteria *C. difficile* (CD) colonizes the gut and becomes sufficiently

abundant to induce colon inflammation and diarrhea via the production of toxins TcdA and TcdB [125]. This disease is an explicit example of how impaired microbiome compositions cause disease.

Traditional treatments for CDI seek to eliminate the presence of CD by administering antibiotics like vancomycin or metronidazole [126]. However, CDI recurrence rates that range from 30-65% and concerns about antibiotic resistance have heralded fecal microbiota transplantation (FMT) as an alternative treatment for CDI [127, 128]. FMT is a bacteriotherapy that attempts to alter the composition of a dysbiotic microbiome (e.g., the microbiome of a person suffering CDI) by engrafting a foreign microbial population provided by a healthy donor into it. When successful, FMT causes the restoration of a healthy microbiome, which suppresses the growth of CD as well as the production of CDI-associated toxins [129]. The success rate of FMT for treating CDI approaches 90% [127].

FMT has been proposed as a treatment for other gastrointestinal diseases including irritable bowel syndrome (IBS), inflammatory bowel disease, ulcerative colitis, and Crohn's disease. Recent clinical trials have returned promising but inconclusive evidence regarding the success of FMT in treating these conditions [2, 3, 4, 5]. For other conditions with distinct microbial signatures— for example, autism spectrum disorder (ASD) or metabolic syndrome [130, 131] — a more intrinsic question regarding the efficacy of bacteriotherapies remains: can symptoms of these conditions be ameliorated by altering the gut microbiome composition? For ASD, preliminary evidence supports this hypothesis. Patients with ASD were given Microbiota Transfer Therapy, a treatment that consists of an initial course of antibiotics followed by regular FMT treatments for ten weeks, and this treatment caused lasting shifts to their microbiome compositions and reductions in their ASD symptoms [132].

At this time, the ability of FMT to treat microbiome-associated diseases is variable.

FMT appears to be an effective treatment for CDI and ASD, but recent clinical studies have been unable to conclusively show that FMT is an effective treatment for IBS [133]. Part of this discrepancy might be due to antibiotics, which are typically administered to CDI patients before FMT is attempted, and which precede the Microbiota Transfer Therapy treatment. More broadly, the factors that contribute to the success or failure of FMT— that is, whether a diseased microbiome is able to be altered or not— are not yet fully known. To shed light on these factors, this chapter used mathematical models to examine the ecological dynamics that underlie FMT.

5.5.2 Advances in modeling the microbiome

Models of the microbiome can span levels of organization, stochasticity, neutrality, and complexity. Recent advances in multi-omics sequencing have buoyed microbiome modeling by measuring the abundances of the microbial genomes, mRNA transcripts, proteins, and metabolites contained in the microbiome [134].

For microbial ecosystems with well-characterized metabolic pathways, flux balance analysis (FBA) predicts the production and consumption of metabolites by microbes at steady state using inferred stoichiometric matrices [134]. In dynamic flux balance analysis (dFBA), FBA is generalized to track metabolite abundances over time, at the expense of requiring additional kinetic parameters in the model [134]. In both of these metabolite-based analyses, fine-grained details about metabolite kinetics need to be either measured or fit. This required level of detail makes FBA and dFBA ill-suited for analyzing complex microbial ecosystems like the human microbiome which, for example, consists of ~ 1000 microbial species and ~ 100000 metabolites [100, 135].

To account for the complexity present in real-world microbial systems, coarse-graining approaches have been employed to model aggregate microbial dynamics and processes.

For example, in wastewater treatment activated sludge models (ASMs) track the abundance of only a few nutrients (e.g. organic carbon, nitrogen, and phosphorus) and the population dynamics of only a few aggregate microbial populations (e.g. microbes that consume carbon, nitrogen, and phosphorus) [60].

Though less detailed than FBA and dFBA models, ASMs require significantly fewer parameters and are able to provide relevant information regarding the design and function of wastewater treatment processes.

Generalized Lotka-Volterra (gLV) models, which were studied in this chapter, assume that the production and consumption of nutrients can be implicitly captured by pairwise interactions between microbial populations. Thus, gLV models only track microbial population abundances. Often, gLV models assume that microbial species within a genus are indistinguishable, and study microbial dynamics coarse-grained at the genus level. Part of the convenience of gLV models is that they may be parameterized with time-series microbial abundance data, which is readily available due to advances in genomic sequencing. Several algorithms have been developed to infer growth rate and interaction parameters of gLV models from time-series abundance data (e.g. LIMITS [98, 104]). gLV models have been used to investigate microbial interactions in cheese [136], to probe how antibiotic perturbations alter the gut microbiome in the context of CDI [7], and to inform treatment of polymicrobial urinary tract infections [137]. When gLV models accurately approximate the underlying microbial systems, simulated interventions (e.g. antibiotic administration or fecal microbiota transplantation) can guide clinical efforts to alter the composition of diseased microbiomes [7, 59, 9].

Each of these previously discussed models make the simplifying assumption that microbial processes and the resulting microbial dynamics are deterministic. In ecology these deterministic approaches are typically used to model niche processes. While microbial dynamics can be relatively consistent in some systems (e.g. in soil microbial communi-

ties [138]), variability— reflecting neutral processes— is generally evident in microbial communities (e.g. in human [139] or fly [27] microbiomes). Mathematical models have been developed to account for this variability [140]. For example, Sloan *et al.* proposed a model to predict microbial abundances based solely on birth-death processes and immigration following Hubbell’s neutral theory [141], while the Ricker model has been used to introduce stochasticity into gLV models [98].

5.6 Conclusion

The ability to control steady state outcomes of ecological systems has a broad practical appeal. These control protocols will rely on a foundation of ecological theory that is still under exploration. Here we introduce a technique to map the dynamical landscape of gLV systems with attractor networks. Although our analysis was solely concerned with gLV systems, fundamental ecological behaviors are demonstrated. For example, when attempting to control the steady-state outcome of a gLV system, the relevant objective is driving the system into the target basin of attraction rather than exactly driving the system to the target steady state. Additionally, when attempting to drive an ecosystem towards a target state, it might be more efficient to use a multi-step control protocol. With the steady progression of ecological theory it is feasible that precision bacteriotherapies, based upon ecological models of the microbiome, will become a regularly administered medicine for the microbiome.

Chapter 6

Control of ecological outcomes with indirect bacteriotherapies

Finally, we take advantage of the analytic tractability of 2D gLV models to create indirect bacteriotherapies. While the previously studied direct bacteriotherapies switch which basin of attraction the system is in by modifying the state of the system, indirect bacteriotherapies alter the size of the basins of attraction by modifying parameters of the high-dimensional gLV model. Interpreted biologically, these indirect bacteriotherapies correspond to altering the gut microbiome environment (e.g. microbiome acidity or macronutrient availability) to drive the system towards a target state.

Without a technique like SSR, the complexity of high-dimensional gLV models impedes the identification of indirect bacteriotherapies: it is *a priori* unclear which of the N^2 interaction parameters should be modified, and so searching for right parameter modification must occur through trial and error. For gLV systems with large numbers of species, this process is computationally expensive. Fortunately, SSR can compress the high-dimensional microbial interactions that are most relevant for the transition between two steady states into a reduced set of four interaction parameters in a 2D gLV model. In

the reduced model it is straightforward to identify which of these four parameters should be modified to alter the 2D system's steady state outcome, and these 2D parameters correspond to a change in the parameters of the high-dimensional system. Therefore, using SSR as a guide, indirect bacteriotherapies in high-dimensional gLV systems may be immediately identified.

6.1 Abstract of “Control of ecological outcomes through deliberate parameter changes in a model of the gut microbiome”

The generalized Lotka-Volterra (gLV) equations are a mathematical proxy for ecological dynamics. We focus on a gLV model of the gut microbiome, in which the evolution of the gut microbial state is determined in part by pairwise inter-species interaction parameters that encode environmentally-mediated resource competition between microbes. We develop an *in silico* method that controls the steady-state outcome of the system by adjusting these interaction parameters. This approach is confined to a bistable region of the gLV model. In this method, a dimensionality reduction technique called steady-state reduction (SSR) is first used to generate a two-dimensional (2D) gLV model that approximates the high-dimensional dynamics on the 2D subspace spanned by the two steady states. Then a bifurcation analysis of the 2D model analytically determines parameter modifications that drive an initial condition to a target steady state. This parameter modification of the reduced 2D model guides parameter modifications of the original high-dimensional model, resulting in a change of steady-state outcome in the high-dimensional model. This control method, called SPARC (SSR-guided parameter change), bypasses the computational challenge of directly determining parameter modifi-

cations in the original high-dimensional system. SPARC could guide the development of indirect bacteriotherapies, which seek to change microbial compositions by deliberately modifying gut environmental variables such as gut acidity or macronutrient availability.

6.2 Introduction

A shared goal in environmental management, ecology, and medicine is to drive an ecosystem towards a target community structure. For example, ocean and lake ecosystems benefit from the suppression of algal blooms, the control of invasive fish species helps preserve the biodiversity of local fish populations, and certain microbial compositions of the gut microbiome that resist pathogenic infections improve the health of the host [142, 143, 7]. It is common to control these ecosystems by directly altering the ecological composition of the community: unwanted algae can be removed by clay, invasive fish species can be killed by biocides, and gut pathogens can be killed by antibiotics [142, 144, 145].

In contrast to these direct methods that modify the ecological state of the system, *indirect* methods can control ecological outcomes by modifying environmental variables which effectively change the dynamical landscape of the system [9]. For example, indirect control methods commonly applied to the previously mentioned systems include reducing nutrient concentrations in water to inhibit algal blooms, lowering the water level to disrupt the spawning of invasive fish, and introducing prebiotics to promote biodiversity in the gut microbiome [142, 146, 147].

In this chapter, we create an *in silico* technique that drives an ecological model towards a target outcome by manipulating parameters that correspond to coarse-grained interactions between populations. Specifically, we seek a finite-time modification of the dynamical landscape that drives an arbitrary initial condition towards a target state.

Although the intervention is temporary, the change in the ecological outcome can be permanent.

This control method is demonstrated in the context of a data-driven model of the gut microbiome [6]. Ecological dynamics are simulated using the generalized Lotka-Volterra (gLV) equations, a commonly used model in theoretical ecology [148]. In these equations, species-species interaction parameters represent environmentally-mediated competition for resources. These systems are often modeled by high-dimensional gLV equations in order to capture the dynamics of the large number of microbial species that inhabit the gut microbiome. In these models, the many inter-species feedbacks lead to complex dynamics. Accordingly, it is difficult to achieve a target steady-state outcome by naively modifying parameter values, as such an approach requires exhaustively searching a large parameter space.

To address these challenges, we focus on a bistable region of the ecological phase space that includes one target steady state and another alternative steady state within the gLV system. Then, a dimensionality-reduction technique called steady-state reduction (SSR) is used to approximate the bistable region of interest and to create a low-dimensional system with a compressed set of interaction parameters [10]. A bifurcation analysis of this 2D system determines a parameter modification that produces a targeted change in steady-state outcome. Lastly, the low-dimensional interaction parameter change is associated with a parameter change in the high-dimensional model, which drives the original system to the target state.

This control method, referred to as SPARC (SSR-guided parameter change), is applied to an 11-dimensional gLV model fit to time-series data from a mouse microbiome experiment [6, 7]. In this experimentally-derived gLV model, SPARC successfully alters the steady-state outcomes of initial conditions by modifying interaction parameters of the model. SPARC as an *in silico* approach is effective when applied to generic gLV

systems, but its applicability to real-world systems is dependent on the fidelity of the underlying gLV model. More generally, this method offers a systematic understanding of how environmental factors and species-species interactions can be manipulated to control ecological outcomes.

6.3 Materials and methods

6.3.1 The generalized Lotka-Volterra equations

The generalized Lotka-Volterra (gLV) equations are a traditional model in theoretical ecology. Due to their flexibility, gLV models have been used to describe a wide variety of system dynamics, including the market values of firms in the stock market, wolf predation of multiple prey species, and the infection dynamics of RNA viruses [149, 150, 151]. In context of the microbiome, the gLV equations have been used to model the population dynamics of gut microbial communities [152], and are given by

$$\frac{d}{dt}y_i(t) = y_i(t)\left(\rho_i + \sum_{j=1}^N K_{ij}y_j(t)\right), \quad (6.1)$$

where $y_i(t)$ denotes the abundance of microbes of species i at a given time t , ρ_i is the growth rate of species i , and K_{ij} is the interaction coefficient between two populations i and j . The interaction parameters K_{ij} form the $N \times N$ interaction matrix K , where N is the number of species. The growth rate parameters ρ_i are constrained by $\rho_i > 0$. The interaction parameters K_{ij} capture prototypical ecological interactions such as competition, symbiosis, and amensalism [152]. Specifically, the parameter K_{ij} represents the effect of species j on species i , which is mediated by environmental factors such as available nutrients. Thus, if environmental factors are changed, the parameters K_{ij} will change as well.

In general, gLV systems can exhibit periodic and chaotic behaviors [153, 154], and the criteria that predict the stability of ecosystems based on their structure have been prominently studied in theoretical ecology [155, 156, 157]. Here, we focus on ecological dynamics that relax to point attractors. In this regime, the gLV dynamics of concern can be represented by a pseudo-energy landscape (e.g. a Lyapunov function), which is a scalar field in ecological state space that behaves analogously to a physical energy landscape.

We wish to determine a coordinated modification of these interaction parameters K_{ij} that drive the system to a target state. The growth rates ρ_i and interactions K_{ij} determine the dynamical landscape on which the microbial system evolves. A modification of the interaction parameter matrix K reshapes the dynamical landscape of the gLV system. This reshaping process is visualized schematically in Fig. 6.1.

All simulations in this chapter were run with the quadrature method `odeint` from the Python module `scipy.integrate`. A Python module that implements SPARC and recreates Fig. 6.3 is provided in the Supplementary Code¹.

6.3.2 A gLV model fit to experimental data

In a mouse experiment, Buffie *et al.* demonstrated that mice that are administered the antibiotic clindamycin become susceptible to *Clostridioides difficile* infection (CDI) [6]. Stein *et al.* fit a gLV model, referred to as the CDI model, to the time-series microbial abundance data from this mouse experiment [7]. For modeling purposes, microbial species are coarse-grained at the genus level, resulting in 11 microbial populations, each described by a population y_i in the gLV model. This gLV model captures the CDI-resistant and CDI-susceptible steady states that are observed in the experiment [7, 9].

¹Supplementary Code used to implement SPARC and generate Fig. 6.3 is available at http://github.com/erijones/ssr_module.

The dynamical structure of this CDI model is characterized by the composition and stability of its steady states. Two steady states in this model, including the experimentally-observed CDI-resistant state, are locally stable (i.e., the eigenvalues of the Jacobian matrix evaluated at these steady state compositions are exclusively negative). Additionally, the CDI model features six steady states whose Jacobian matrices have one non-negative eigenvalue (referred to as having one “unstable direction”); it also features 23 steady states, including the CDI-susceptible state, whose Jacobian matrices have two non-negative eigenvalues (i.e., with two unstable directions). The CDI-susceptible state is composed of 5 coarse-grained species and the CDI-resistant state is composed of 3 coarse-grained species. In these steady states, the abundance of all other species is zero. The detailed compositions of these steady states are given in the Supplementary Information.

In this chapter, we examine the transition between the CDI-susceptible state and the CDI-resistant state, and apply SPARC to the bistable region formed by these states. First, we demonstrate SPARC in the “infection” scenario in which the CDI-susceptible steady state is treated as the target state and the CDI-resistant steady state is designated the alternative state. We consider an initial condition on the plane spanned by the target state \vec{y}_a and the alternative state \vec{y}_b that tends towards \vec{y}_b in the absence of any intervention. The goal of SPARC is to find a modification to the interaction matrix ΔK that alters the evolution of this initial condition and drives it towards the target state. In the infection scenario, this parameter change represents a disruption of the microbial dynamics that can drive the system towards a state susceptible to CDI.

We later consider the “recovery” scenario in which the target state is the CDI-resistant state and the alternative state is the CDI-susceptible state. In this scenario too, SPARC alters the steady-state behavior of an initial condition so that it flows towards the target CDI-resistant state. The parameter change generated by SPARC in this scenario informs

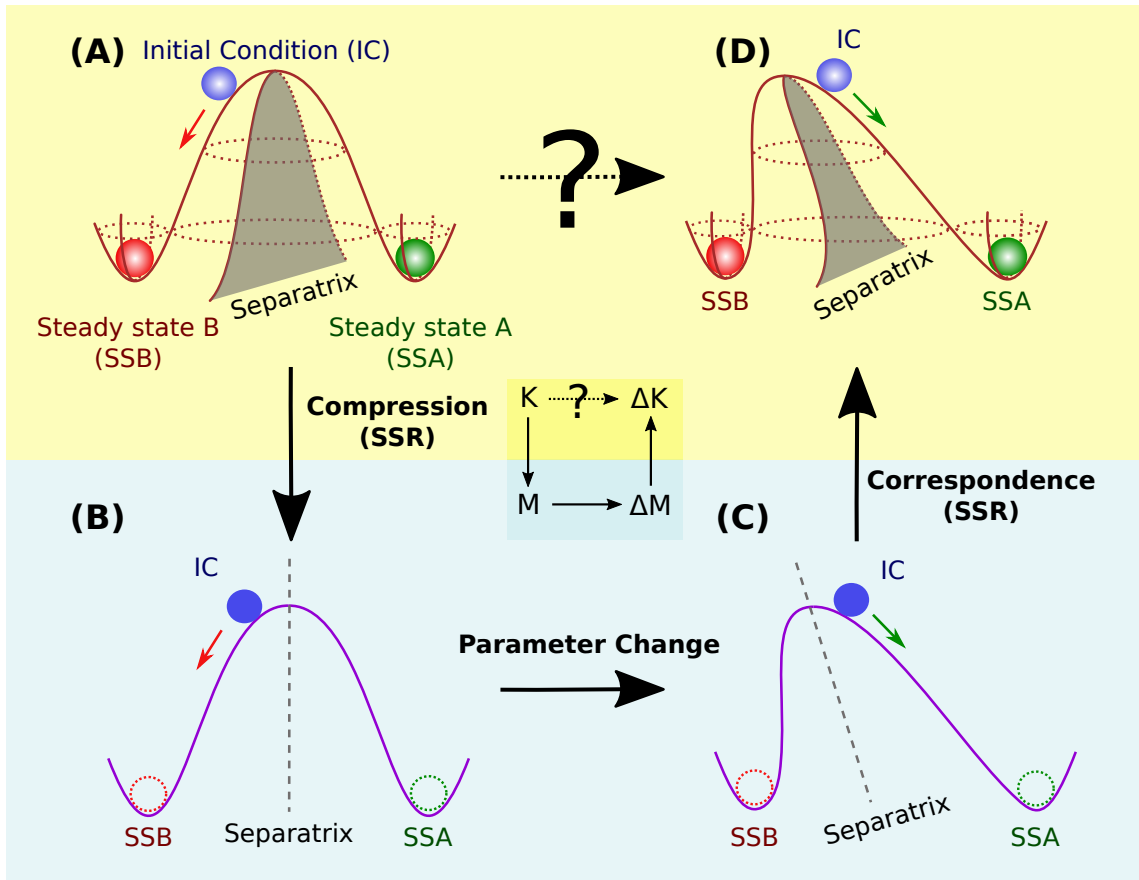


Figure 6.1: **A schematic overview of how SPARC (SSR-guided parameter change) controls steady-state outcomes.** (A) A bistable region in a high-dimensional gLV model, with two steady states and an initial condition tending towards the alternative steady state (shown in red), is represented as a pseudo-energy (Lyapunov) landscape. This landscape is parameterized by the interaction matrix K of the high-dimensional gLV system. (B) The high-dimensional landscape is compressed into a reduced 2-dimensional landscape, generated by the dimensionality-reduction technique steady-state reduction (SSR) as described in Eq. (6.4). This 2D landscape is parameterized by a 2×2 interaction matrix M . (C) Guided by a bifurcation analysis of this reduced 2D system, a modification of the interaction matrix ΔM changes the Lyapunov landscape in a targeted way. After this change, the initial condition tends towards the healthy steady state (shown in green) in the low-dimensional system. (D) A high-dimensional parameter modification ΔK , informed by the 2D parameter change ΔM via the SSR formulae, changes the high-dimensional Lyapunov landscape. It is computationally difficult to identify this parameter change directly from the original model (A to D), but using SSR and the bifurcation analysis of the 2D model, this change is straightforward (A to B to C to D).

the intervention needed to recover from the CDI-susceptible state in this model. These results demonstrate that, for this pair of steady states in the CDI model, SPARC is able to drive microbial dynamics in the direction of either steady state.

6.3.3 SSR-guided parameter change (SPARC)

We develop a multi-step control framework to determine a parameter change that drives a given initial condition towards a target state. A bistable landscape of interest in a high-dimensional gLV model is first reduced into a 2D gLV model using steady-state reduction (SSR) [10]. This control framework is called SPARC (SSR-guided parameter change), and summarized in Fig. 6.1.

Steady-state reduction

Steady-state reduction (SSR), developed by Jones and Carlson, is a mathematical technique that compresses a high-dimensional gLV system into a 2D gLV system, as shown in Fig. 6.1A and B [10]. In a high-dimensional gLV model of N species, there are N^2 interaction parameters. Due to the complexity of the feedbacks of the ecological system, it is analytically intractable and computationally expensive to numerically determine how modifications of interaction parameters affect the asymptotic behavior of arbitrary initial conditions.

To understand the dynamics in the high-dimensional phase space, we consider bistable systems and focus on the subspace spanned by the two steady states \vec{y}_a and \vec{y}_b . The SSR technique views steady states \vec{y}_a and \vec{y}_b of the high-dimensional model as idealized composite states and constructs a new set of 2D gLV equations in which the basis vectors correspond to the high-dimensional steady states. This 2D gLV system approximates the

slow manifold that connects \vec{y}_a and \vec{y}_b , and is the best possible gLV approximation of the high-dimensional dynamics on the subspace spanned by \vec{y}_a and \vec{y}_b [10]. Explicitly, the approximate 2D gLV system has the form

$$\begin{aligned}\frac{dx_a}{dt} &= x_a(\mu_a + M_{aa}x_a + M_{ab}x_b), \text{ and} \\ \frac{dx_b}{dt} &= x_b(\mu_b + M_{ba}x_a + M_{bb}x_b),\end{aligned}\tag{6.2}$$

where x_a corresponds to the high-dimensional gLV system's component in the direction $\hat{x}_a = \frac{\vec{y}_a}{\|\vec{y}_a\|_2}$, x_b corresponds to the direction $\hat{x}_b = \frac{\vec{y}_b}{\|\vec{y}_b\|_2}$, and $\|\vec{v}\|_2$ is the 2-norm of \vec{v} . The parameters μ_a and μ_b represent the growth rates of x_a and x_b , and the M_{ij} interaction parameters form a 2D interaction matrix M . SSR yields the reduced interspecies interaction parameters M_{ab} and M_{ba} , which are given by

$$\begin{aligned}M_{ab} &= \frac{\sum_{i,j=1}^N K_{ij}(y_{ai}y_{bj} + y_{bi}y_{aj}) \left(y_{ai} - y_{bi} \sum_{k=1}^N y_{ak}y_{bk} \right)}{1 - \left(\sum_{i=1}^N y_{ai}y_{bi} \right)^2}, \text{ and} \\ M_{ba} &= \frac{\sum_{i,j=1}^N K_{ij}(y_{bi}y_{aj} + y_{ai}y_{bj}) \left(y_{bi} - y_{ai} \sum_{k=1}^N y_{bk}y_{ak} \right)}{1 - \left(\sum_{i=1}^N y_{bi}y_{ai} \right)^2},\end{aligned}\tag{6.3}$$

where y_{ai} and y_{bi} are the i th components of the unit vectors $\hat{y}_a \equiv \vec{y}_a/\|\vec{y}_a\|_2$ and $\hat{y}_b \equiv \vec{y}_b/\|\vec{y}_b\|_2$, respectively. The other 2D parameters μ_a , μ_b , M_{aa} , and M_{ab} are given by

$$\begin{aligned}\mu_\gamma &= \frac{\vec{\rho} \cdot \vec{y}_\gamma^{\circ 2}}{\|\vec{y}_\gamma\|_2^2}, \text{ and} \\ M_{\gamma\delta} &= \frac{(\vec{y}_\gamma^{\circ 2})^T K \vec{y}_\delta}{\|\vec{y}_\gamma\|_2^2 \|\vec{y}_\delta\|_2},\end{aligned}\tag{6.4}$$

where $\gamma, \delta \in a, b$. When the high-dimensional steady states \vec{y}_a and \vec{y}_b are orthogonal, the interspecies interaction parameters M_{ab} and M_{ba} in Eq. (6.3) reduce to the interaction parameters in Eq. (6.4). In these formulae, $\vec{y}^{\circ 2} \equiv \text{diag}(\vec{y})\vec{y}$ is the element-wise square of \vec{y} . Note that SSR maps the high-dimensional steady states \vec{y}_a and \vec{y}_b to the points $(\|\vec{y}_a\|_2, 0)$ and $(0, \|\vec{y}_b\|_2)$, which are the steady states of the 2D model. Additionally, if the high-dimensional steady states are stable, SSR guarantees that their low-dimensional counterparts are stable as well. The fidelity of the SSR method is demonstrated in Fig. 6.3, where it is applied to an experimentally-derived gLV system. Additional examples are provided in the Supplementary Information.

Bifurcation analysis

After the high-dimensional gLV model is reduced to a 2D model, the next step is to find a parameter change in the 2D model that changes the steady-state behavior of the system, as shown in Fig. 6.1B and C. Simplifying the high-dimensional system using SSR results in a 2D gLV model with two growth rate parameters, μ_a and μ_b and four interaction parameters, M_{aa} , M_{ab} , M_{ba} , and M_{bb} .

When the steady states of the original high-dimensional bistable system are stable, SSR guarantees two stable steady states at $(1, 0)$ and $(0, 1)$. In addition to these two steady states, the system possesses a trivial unstable steady state at $(0, 0)$, and another hyperbolic fixed point with nonzero x_a and x_b components. The separatrix, which delineates the basins of attraction of the $(1, 0)$ and $(0, 1)$ steady states, is topologically required to pass through this hyperbolic fixed point.

When nondimensionalized, the 2D gLV equations Eq. (6.2) become

$$\begin{aligned}\frac{d\tilde{x}_a}{dT} &= \tilde{x}_a(1 - \tilde{x}_a - \tilde{M}_{ab}\tilde{x}_b), \text{ and} \\ \frac{d\tilde{x}_b}{dT} &= \tilde{x}_b(\tilde{\mu}_b - \tilde{M}_{ba}\tilde{x}_a - \tilde{x}_b),\end{aligned}\tag{6.5}$$

where $\tilde{x}_a = -\frac{M_{aa}}{\mu_a}x_a$, $\tilde{x}_b = -\frac{M_{bb}}{\mu_b}x_b$, $T = \mu_a t$, $\tilde{M}_{ab} = M_{ab}/M_{bb}$, $\tilde{M}_{ba} = M_{ba}/M_{aa}$, and $\tilde{\mu}_b = \mu_b/\mu_a$. In terms of these nondimensionalized parameters, the two steady states are now at $(1, 0)$ and $(0, \tilde{\mu}_b)$. The coordinate of the hyperbolic fixed point is given by

$$\left(\frac{\tilde{M}_{ab}\tilde{\mu}_b - 1}{\tilde{M}_{ab}\tilde{M}_{ba} - 1}, \frac{\tilde{M}_{ba} - \tilde{\mu}_b}{\tilde{M}_{ab}\tilde{M}_{ba} - 1} \right).\tag{6.6}$$

Since the separatrix passes through this steady state, adjusting the parameters \tilde{M}_{ab} and \tilde{M}_{ba} alters its position and stability, as shown in Fig. 6.2.

A necessary condition for the steady states $(1, 0)$ and $(0, \tilde{\mu}_b)$ to be stable is that $\tilde{M}_{ab}\tilde{M}_{ba} - 1 > 0$ [10]. Thus, when \tilde{M}_{ab} is made smaller than $1/\tilde{\mu}_b$ with \tilde{M}_{ba} fixed, the \tilde{x}_a coordinate of the unstable steady state becomes negative. Equivalently, in Fig. 6.2 this corresponds to system moving from the top-right configuration to the top-left configuration. A linearized stability analysis finds that the topological structure of the 2D phase space also changes after this parameter change is made. As shown in the top-left panel of Fig. 6.2, the steady state at $(0, \tilde{\mu}_b)$ becomes unstable once \tilde{M}_{ab} is smaller than $1/\tilde{\mu}_b$, which forces initial conditions in the top-right quadrant of the phase space towards the stable state at $(1, 0)$. Similarly, once \tilde{M}_{ba} is smaller than $\tilde{\mu}_b$, the \tilde{x}_b coordinate of the hyperbolic steady state becomes negative. In Fig. 6.2 this corresponds to crossing from the top-right to the bottom-right, at which point the steady state at $(1, 0)$ becomes unstable. The bifurcation diagram in Fig. 6.2 provides a guide for how the steady-state structure of the 2D gLV equations depends on the interaction parameters.

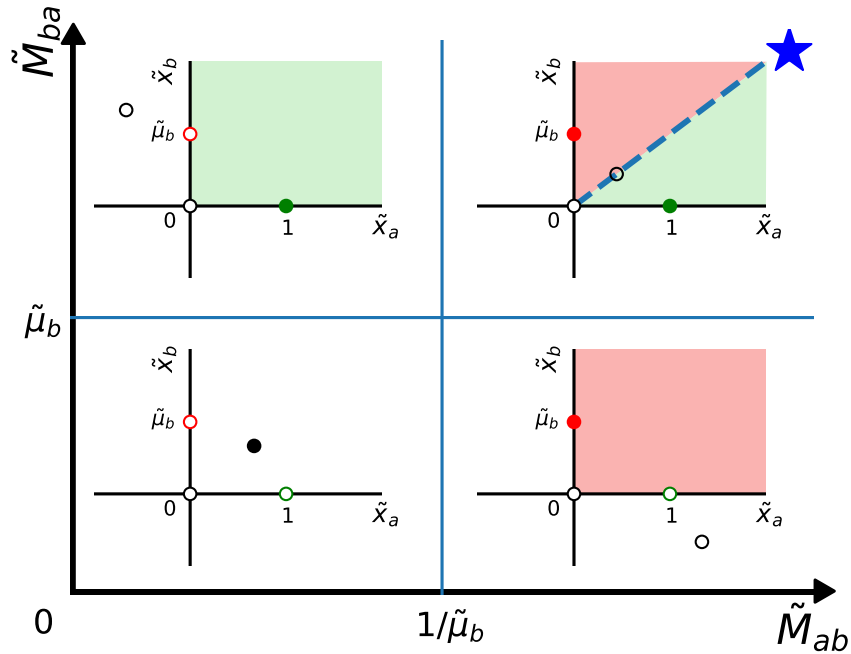


Figure 6.2: **A bifurcation diagram of nondimensionalized 2D gLV systems.** This diagram shows phase space representations of different topological classes of 2D gLV dynamical landscapes, and their dependence on the nondimensionalized parameter values \tilde{M}_{ab} , \tilde{M}_{ba} , and $\tilde{\mu}_b$ of Eq. (6.5). The lines at $\tilde{M}_{ab} = 1/\tilde{\mu}_b$ and $\tilde{M}_{ba} = \tilde{\mu}_b$ split the parameter space into four quadrants that each correspond to a different topological configuration of phase space. The graph inside each quadrant shows a representative phase space configuration of the nondimensionalized gLV system, where \tilde{x}_a and \tilde{x}_b are the rescaled populations in Eq. (6.5). The hollow dots represent unstable steady states, and the filled dots represent stable steady states. The basins of attraction of the steady states $(1, 0)$ and $(0, \tilde{\mu}_b)$ are shaded in green and red, respectively. The upper-right quadrant, labeled with a blue star, represents the parameter regime in which bistable 2D landscapes occur. An alternative visualization of this bistable landscape is schematized in Fig. 6.1 as a pseudo-energy landscape. The reduced 2D gLV models, generated by applying SSR to bistable regions in high-dimensional gLV models, reside in this upper-right quadrant. In this bistable quadrant, the separatrix passes through the hyperbolic steady state with non-negative coordinates. The steady states at $(1, 0)$ and $(0, \tilde{\mu}_b)$ undergo transcritical bifurcations in response to changes in \tilde{M}_{ab} and \tilde{M}_{ba} , yielding the diagrams in adjacent panels. The lower-left quadrant is included for completeness.

This bifurcation analysis indicates how to move the separatrix in a particular direction. Numerical methods determine the minimal change of parameters \tilde{M}_{ab} or \tilde{M}_{ba} that switch the asymptotic steady-state behavior of a given initial condition. In simulations where the target steady state is located at $(1, 0)$, the value of \tilde{M}_{ab} is decreased incrementally, spanning from its original value through $1/\tilde{\mu}_b$. In terms of the dimensionalized 2D gLV system, this corresponds to keeping M_{bb} constant while M_{ab} is modified until the separatrix is shifted to a position where the initial condition switches from one basin of attraction to the other.

Correspondence between 2D and high-dimensional gLV models

Changes in the 2D interaction parameters that drive an initial condition to a target state are associated with changes in the high-dimensional interaction parameters, since the 2D reduced parameters are functions of the high-dimensional parameters via the SSR formulae. This is schematically shown in the transition from Fig. 6.1C to Fig. 6.1D. More explicitly, Eq. (6.4) can be re-written as

$$M_{\gamma\delta} = \sum_{i,j} \alpha_{ij}^{\gamma\delta}(\vec{y}_a, \vec{y}_b) K_{ij}, \quad (6.7)$$

where $\gamma, \delta \in \{a, b\}$, and \vec{y}_a and \vec{y}_b are the two steady states of interest. In this chapter, since the target state is placed at $(1, 0)$, it is most important to modify the parameter $\tilde{M}_{ab} = M_{ab}/M_{bb}$. For simplicity we only consider modifications to M_{ab} , and therefore are primarily concerned with the coefficients α_{ij}^{ab} , hereafter referred to as α_{ij} . Thus, from this correspondence a modification in the 2D interaction matrix M may be reproduced in the high-dimensional system by modifying the high-dimensional interaction matrix K . This choice is degenerate — there is more than one way to change the high-dimensional interaction matrix K that corresponds to the same 2D parameter modification. Note

that the smallest possible high-dimensional parameter change ΔK_{ij} is associated with the largest coefficient α_{ij} .

6.4 Results

SPARC (SSR-guided parameter change) controls the steady-state outcome of a high-dimensional gLV system by deliberately changing the geometry of its dynamical landscape. SPARC (i) approximates a bistable landscape of a high-dimensional gLV system by its 2D SSR-generated counterpart, (ii) identifies a 2D interaction parameter change that switches the asymptotic behavior of an initial condition on this bistable landscape, and (iii) associates the 2D parameter change with a parameter modification in the high-dimensional gLV system. This parameter modification shifts the high-dimensional landscape so that an otherwise disease-prone initial condition will instead tend towards the target state.

Note that since the steady states of the high-dimensional model are dependent on the interaction matrix K , a small change in this matrix will slightly modify the coordinates of the steady states. Thus, to allow the system to evolve back to the original steady states, this parameter modification must be turned off after some time. To initially demonstrate SPARC, the parameter modification is turned off once the system stabilizes at the shifted steady state (Fig. 6.3). When SPARC is applied to the CDI model, the parameter modification is small enough that the changes in steady state locations are negligible. Later, when considering the “recovery” scenario, the parameter modification is turned off before the system stabilizes at any steady state; in this case there is a critical duration that the parameter modification must be active for in order for the intervention to be successful (Fig. 6.5).

In this section, SPARC is first applied to the CDI model fit by Stein *et al.* to data from

a *Clostridioides difficile* infection (CDI) experiment in mice [7]. Then, the robustness of SPARC is examined by applying it to synthetic gLV models.

6.4.1 Steady-state reduction (SSR) produces a 2D approximation to bistable dynamics in a high-dimensional gut microbiome model

First, bistable dynamics in the CDI model are approximated by reduced dynamics on a 2D subspace generated by steady-state reduction (SSR). We focus on two steady states of this gLV model that correspond to experimentally observed CDI-resistant and CDI-susceptible microbiome compositions. For the initial demonstration of SPARC, we consider the “infection” scenario in which the CDI-susceptible state is defined as the target state and the CDI-resistant state is defined as the alternative state.

The target state and the alternative state are represented by the high-dimensional vectors \vec{y}_a and \vec{y}_b , respectively. The microbial dynamics that result from the initial condition $(0.5\vec{y}_a + 0.5\vec{y}_b)$ tend towards the alternative steady state \vec{y}_b . To visualize these dynamics, the trajectory is projected onto a plane spanned by the steady states \vec{y}_a and \vec{y}_b , as displayed in Fig. 6.3A. In this figure, the axes are rescaled so that the steady state \vec{y}_a is located at point $(1, 0)$ and the steady state \vec{y}_b is located at point $(0, 1)$. The separatrix shown in Fig. 6.3A is numerically generated from trajectory simulations. Notice that on this subspace, the initial condition is above the separatrix, and hence the initial condition evolves towards the alternative steady state at $(0, 1)$.

This 11D bistable landscape is approximated by a reduced 2D gLV model generated by SSR, according to Eq. (6.4). The SSR-generated parameter values and their nondimensionalized counterparts are provided in the Supplementary Information. The dynamics of the reduced 2D trajectory were initial condition $(0.5, 0.5)$ are displayed in

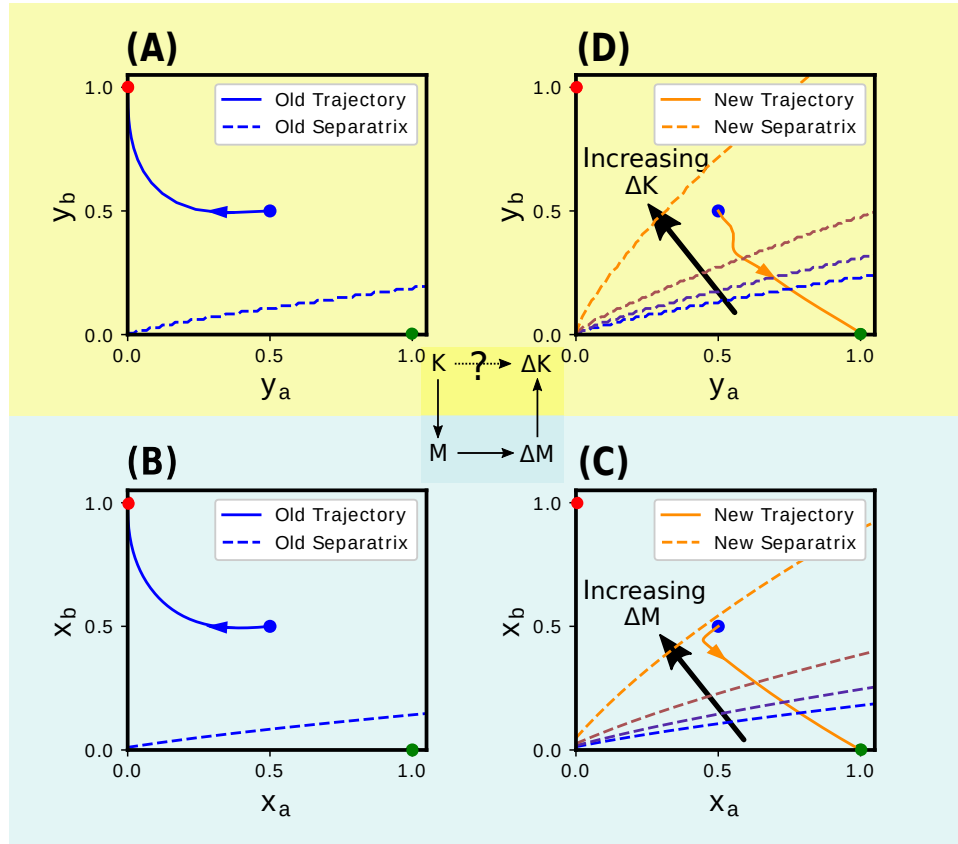


Figure 6.3: **A realization of SPARC, as described in Fig. 1, applied to the infection scenario of the CDI model.** (A) The phase space of the CDI model [7] is projected onto the 2D plane spanned by the target steady state \vec{y}_a and the alternative steady state \vec{y}_b . The target and alternative steady states at $(\|\vec{y}_a\|_2, 0)$ and $(0, \|\vec{y}_b\|_2)$ are rescaled in this plot to the points $(1, 0)$ and $(0, 1)$. The in-plane separatrix, generated numerically, delineates the basins of attraction. (B) Steady-state reduction (SSR) generates an approximate 2D phase space. Notice that the 2D separatrix and trajectory qualitatively resemble those in (A). (C) The 2D separatrix moves as the 2D interaction matrix M is modified. Four separatrices corresponding to four changes with increasing magnitude in the interaction matrix ΔM are shown. The matrix element and direction of this change are guided by the bifurcation analysis in Fig. 6.2. A sufficiently large parameter change alters the steady-state outcome of the initial condition $(0.5, 0.5)$. (D) Changes in the low-dimensional interaction parameter ΔM are associated with changes in high-dimensional parameter ΔK by SSR formulae. The resulting shift in the high-dimensional separatrix is qualitatively similar to that of the low-dimensional system. In particular, the initial condition $(0.5\vec{y}_a + 0.5\vec{y}_b)$ now evolves towards the target steady state. SPARC successfully alters the steady-state outcome without having to search a 121-dimensional parameter space.

Fig. 6.3B, and are similar to the projection of the 11D dynamics in Fig. 6.3A. Note that the position of the separatrix, which is generated analytically in the 2D model [10], is well-approximated by SSR. In the Supplementary Information it is further demonstrated that this reduced 2D model accurately approximates the high-dimensional trajectories that originate from other initial conditions.

It is difficult to identify the interspecies feedbacks that induce bistability in a high-dimensional system: in general, it is unclear how the separatrix changes as a function of the system parameters. On the other hand, in the reduced 2D gLV system, there are well-defined conditions for bistability, namely

$$\begin{aligned}\tilde{M}_{ab}\tilde{\mu}_b &= (M_{ab}/M_{bb})(\mu_b/\mu_a) > 1, \quad \text{and} \\ \tilde{M}_{ba} &= M_{ba}/M_{aa} > \tilde{\mu}_b = \mu_b/\mu_a.\end{aligned}\tag{6.8}$$

Since these low-dimensional parameters M_{ab} and M_{ba} are linear combinations of the high-dimensional parameters K_{ij} , the conditions for bistability can be decomposed into their relative contributions from the high-dimensional interspecies feedbacks K_{ij} .

Specifically, consider the numerators of these inequalities, $M_{ab} = \sum_{ij} \alpha_{ij}^{ab} K_{ij}$ and $M_{ba} = \sum_{ij} \alpha_{ij}^{ba} K_{ij}$ (as in Eq. (6.7)). Then, the relative contributions to M_{ab} by each of the $\alpha_{ij}^{ab} K_{ij}$ terms may be compared (and likewise for M_{ba}). When \vec{y}_a corresponds to the CDI-susceptible state and \vec{y}_b corresponds to the CDI-resistant state, the contributions to M_{ab} are dominated by the inhibition of *Barnesiella* on both *Blautia* and undefined genus of *Enterobacteriaceae* (i.e., the contributions $\alpha_{9,1}^{ab} K_{9,1}$ and $\alpha_{5,1}^{ab} K_{5,1}$). Contributions to M_{ba} are dominated by the inhibition of undefined genus of *Enterobacteriaceae* and *Blautia* on unclassified *Lachnospiraceae* and *Barnesiella* (i.e., the contributions $\alpha_{3,9}^{ba} K_{3,9}$, $\alpha_{3,5}^{ba} K_{3,5}$, $\alpha_{1,9}^{ba} K_{1,9}$, and $\alpha_{1,5}^{ba} K_{1,5}$). Additional details about these contributions are provided in the Supplementary Information. Thus, the bistability between steady states \vec{y}_a

and \vec{y}_b is largely driven by feedbacks between a pair of species present in \vec{y}_a (undefined genus of Enterobacteriaceae and Blautia) and a pair of species present in \vec{y}_b (unclassified Lachnospiraceae and Barnesiella).

6.4.2 Bifurcation analysis guides interaction parameter changes that modify steady-state outcomes in reduced 2D gLV systems

Next, the bifurcation analysis of 2D gLV systems depicted in Fig. 6.2 indicates how to drive an initial condition (0.5, 0.5) towards the target steady state (1, 0). This requires enlarging the basin of attraction of the steady state (1, 0), which is equivalent to rotating the separatrix counter-clockwise. The SSR-generated 2D system is bistable, and thus belongs to the topological class in the upper-right quadrant of Fig. 6.2. Accordingly, the parameter M_{ab} is decreased. When $M_{ab} = M_{bb}\mu_a/\mu_b$, the alternative steady state at (0, 1) becomes unstable, guaranteeing the initial condition (0.5, 0.5) will tend towards the target state at (1, 0). However, to identify the minimal intervention that drives the system towards the target state, we consider intermediate steps between the original value of M_{ab} and the bifurcation point $M_{bb}\mu_a/\mu_b$.

Four incremental parameter changes are plotted in Fig. 6.3C. On the fourth step, the separatrix is sufficiently modified so that the initial condition tends towards the target healthy steady state. The original 2D interaction matrix M , the parameter change to M_{ab} , and the resulting interaction matrix $M + \Delta M$ are visualized in Fig. 6.4E-G. The trajectory plots in the bottom-left and the bottom-right corners of Fig. 6.4 illustrate the behavior of the 2D gLV system parameterized by M and $M + \Delta M$, respectively. Therefore, SPARC can identify and modify interaction parameters to switch the steady state behavior of this 2D model.

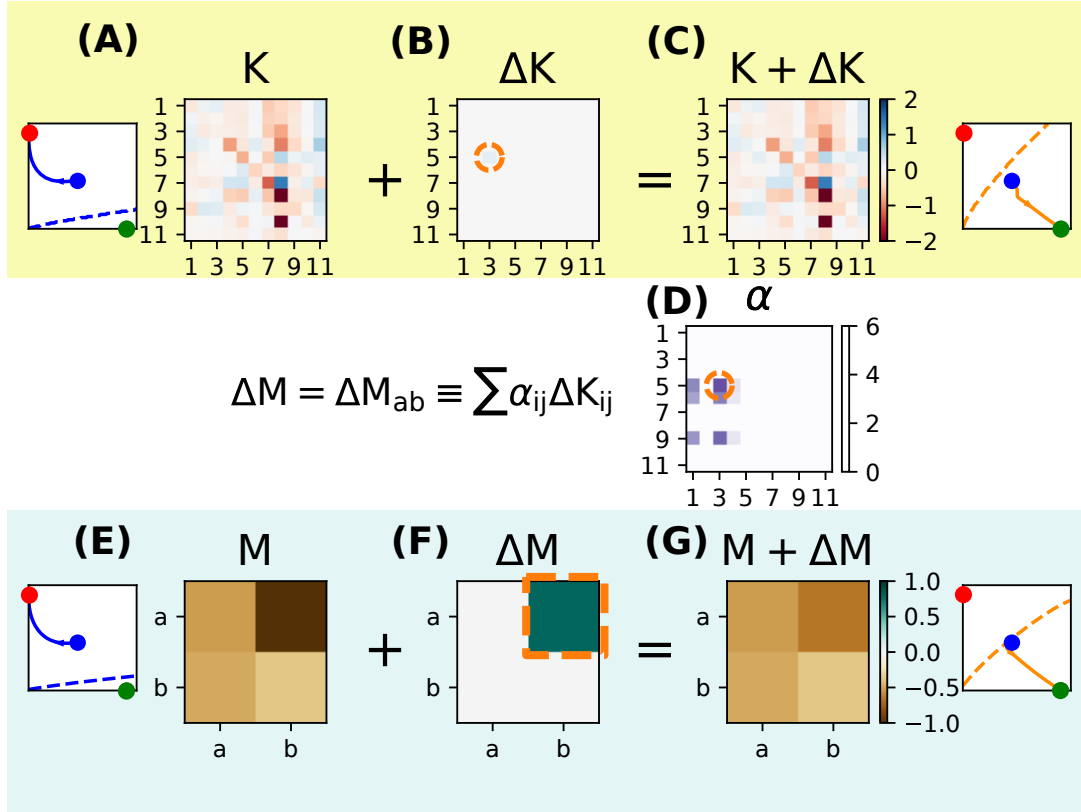


Figure 6.4: **Visualization of the SSR-guided changes to gLV interaction parameters.** As described in SPARC, the original high-dimensional interaction matrix K (A), SSR-guided parameter change ΔK (B), and the resulting interaction matrix $K + \Delta K$ (C) are displayed. The steady-state reduced parameter matrix M (E), bifurcation analysis guided parameter change ΔM (F), and the resulting 2D interaction matrix $M + \Delta M$ (G) are also displayed. The low-dimensional parameter change ΔM , is related to high-dimensional parameter changes through the SSR formulae Eq. (6.4). The α_{ij} coefficients represent the weights of the elements of the high-dimensional interaction matrix K in the steady-state reduced interaction matrix M , as in Eq. (6.7) and these coefficients are visualized in panel (D). To minimize the size of the high-dimensional parameter change, the interaction parameter K_{ij} that corresponds to largest coefficients α_{ij} is chosen to be modified. In this case, the coefficient $\alpha_{5,3}$ is the largest, which determines the choice of ΔK . The phase space diagrams in each corner illustrate the trajectory of the initial condition $(0.5\vec{y}_a + 0.5\vec{y}_b)$ or $(0.5, 0.5)$, for each of the adjacent interaction matrices.

6.4.3 SSR maps low-dimensional bifurcation behavior to the high-dimensional system

Finally, having determined the low-dimensional parameter modification that alters the separatrix in the reduced 2D model (as shown in Fig. 6.3C), corresponding high-dimensional parameters that alter the system outcome in the original model can be identified. Due to the degeneracy associated with mapping from the low-dimensional to high-dimensional parameters, as is clear in the SSR formulae given by Eq. (6.4), there are numerous modifications to the high-dimensional interaction matrix K that correspond to the same change in the 2D interaction matrix, as shown in Fig. 6.4D. In the CDI model, if the parameter change is confined to only one element of K , there are a total of 121 choices. In order to make the smallest change in the interaction matrix K , the coefficient K_{ij} corresponding to the largest α_{ij} value is chosen, as described in Eq. (6.7). Specifically, the parameter change $\Delta K_{5,3} = 0.1744$ is used.

In Fig. 6.4D the magnitudes of the α_{ij} coefficients are plotted, and the largest coefficient is highlighted with a dashed box. In the bottom row of Fig. 6.4, the original K matrix (panel A), the required modification ΔK corresponding to that α_{ij} coefficient (panel B), and the resulting modified interaction matrix $K + \Delta K$ (panel C) are displayed. The trajectories in the upper-left and upper-right corners indicate the behavior of the systems parameterized by K and $K + \Delta K$, respectively.

Fig. 6.3D displays the results of a representative 11D interaction matrix change ΔK that drives the initial condition to the target state \vec{y}_a . As in Fig. 6.3C, four incremental parameter changes that each modify the separatrix are plotted. The largest of these four parameter changes rotates the 11D separatrix counter-clockwise so that the initial condition $(0.5\vec{y}_a + 0.5\vec{y}_b)$ tends towards the healthy steady state \vec{y}_a . Although small discrepancies exist between Fig. 6.3C and Fig. 6.3D due to the SSR approximation,

SPARC successfully alters the steady-state outcome of a high-dimensional gLV system by deliberately changing its interaction parameters.

6.4.4 SPARC generates a finite-time intervention that drives a disease-prone initial condition towards a healthy state in the CDI model

Next, we consider the recovery scenario in which the “healthy” CDI-resistant state is the target state \vec{y}_a and the “diseased” CDI-susceptible state is the alternative state \vec{y}_b . The initial condition at $(0.1\vec{y}_a + 0.9\vec{y}_b)$ is chosen to demonstrate that SPARC can be effective even when the initial condition is closer to the alternative state than to the target state. As in the previous case, SPARC is applied to change the steady-state outcome of this initial condition, which is shown in Fig. 6.5. For clarity, the shifted separatrices in Fig. 6.5C and D are not displayed.

Without any parameter modification, the bistable region is exactly the reflection of the previous case, as shown in Fig. 6.5A and B. However, the parameter modification generated by SPARC shifts the separatrix in the opposite direction. In this case, the separatrix is already close to the alternative steady state at $(0, 1)$. The 2D parameter modification makes $\tilde{M}_{ab} < 1/\tilde{\mu}_b$, resulting in the steady state at $(0, 1)$ becoming unstable, as shown in Fig. 6.2 (top-right and top-left panels). Therefore, although the initial condition is nearby the alternative steady state, after modifying the low-dimensional parameters it tends towards the target state at $(1, 0)$.

The successful 2D parameter change is projected to the high-dimensional model. Notably, the applied parameter change causes the steady state \vec{y}_b in the high-dimensional model to become unstable. Thus SPARC is capable of altering the stability properties of high-dimensional steady states, which enables the control of initial conditions even when

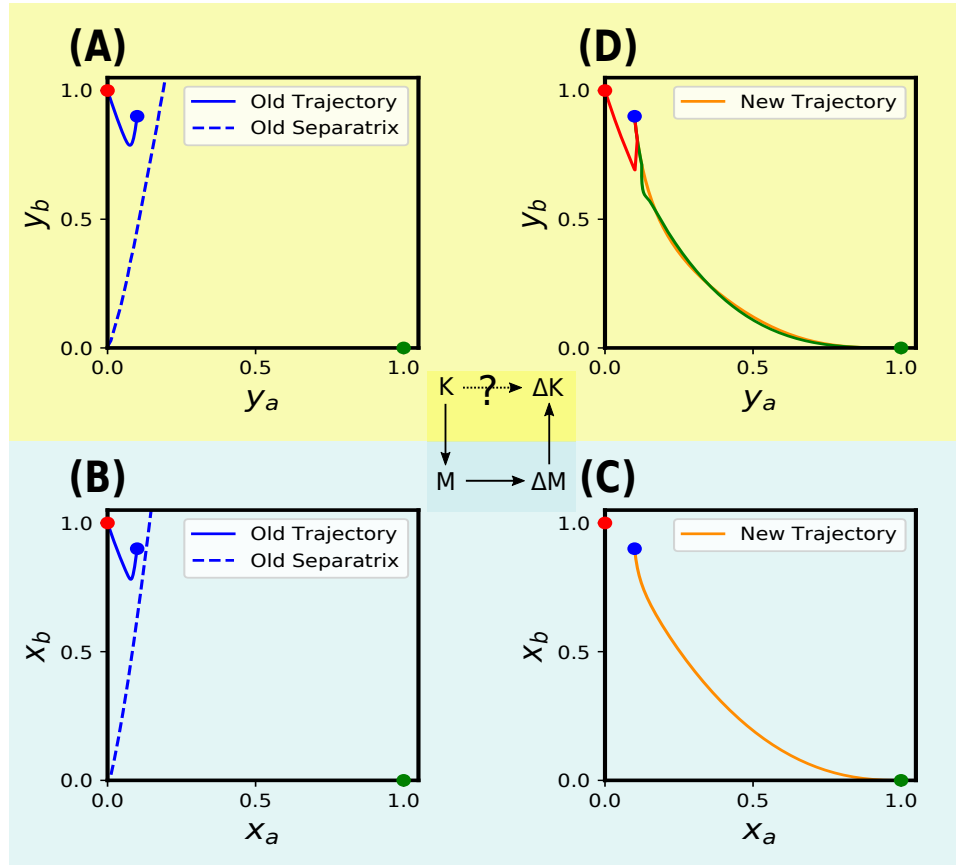


Figure 6.5: **A realization of SPARC applied to the CDI model in the recovery scenario.** Here, the target state \vec{y}_a is the CDI-resistant state and the alternative state \vec{y}_b is the CDI-susceptible state. The target and alternative steady states at $(\|\vec{y}_a\|_2, 0)$ and $(0, \|\vec{y}_b\|_2)$ are rescaled in this plot to the points $(1, 0)$ and $(0, 1)$. (A) With the positions of the steady states switched, the 2D projection of the high-dimensional bistable region shown in Fig. 6.3 is redrawn. Here the separatrix is close to the alternative state \vec{y}_b . The initial condition at $0.1\vec{y}_a + 0.9\vec{y}_b$ tends towards the alternative state \vec{y}_b . (B) The SSR formulae are applied to generate a 2D approximate model. (C) After a parameter change ΔM_{ab} , the steady state at $(0, 1)$ becomes unstable and the initial condition now tends towards the target state. (D) The parameter change in the 2D model is associated with a parameter change in the original CDI model. The yellow line plots the trajectory when the parameter modification is turned off after the system stabilizes, as in the case of Fig. 6.3. The red line shows the trajectory when the parameter change is turned off before the critical time, and the green line shows the trajectory when the parameter change is turned off after the critical time.

they are located at or nearby an alternative steady state.

Fig. 6.5D also shows the effect of the duration of the parameter modification. For SPARC to succeed, the parameter modification needs to be active long enough for the microbial state to escape its original basin of attraction. The red trajectory in Fig. 6.5 demonstrates that the system returns back to the alternative steady state if the parameter change is applied for too short of a duration. The green trajectory illustrates that the system will evolve towards the target state as long as the parameter change is active beyond a critical duration. This critical duration varies from case to case and was determined here numerically by trial-and-error. The orange trajectory occurs when the parameter change is active until the system stabilizes at the shifted steady state, as in Fig. 6.3.

6.4.5 SPARC successfully changes steady state outcomes in synthetic gLV models

“Permuted” synthetic models

To verify that SPARC is generalizable, it is applied to 100 synthetic parameter sets generated by permuting the interaction parameters of the CDI model. In these synthetic parameter sets, the growth rates ρ_i are kept the same as in the CDI model. The diagonal entries of the interaction matrix K are all negative (as shown in Fig. 6.4A), which is biologically reasonable since positive diagonal entries imply unphysical infinite growth. To ensure the synthetic data sets preserve this property, the diagonal and off-diagonal entries of the K matrix are permuted independently. All 100 parameter sets are generated in this way. This permutation process is demonstrated in Fig. 6.6A and B.

In the next step, bistable regions for each synthetic system must be identified in order for SPARC to be applicable. Steady state analysis shows that, for a randomly

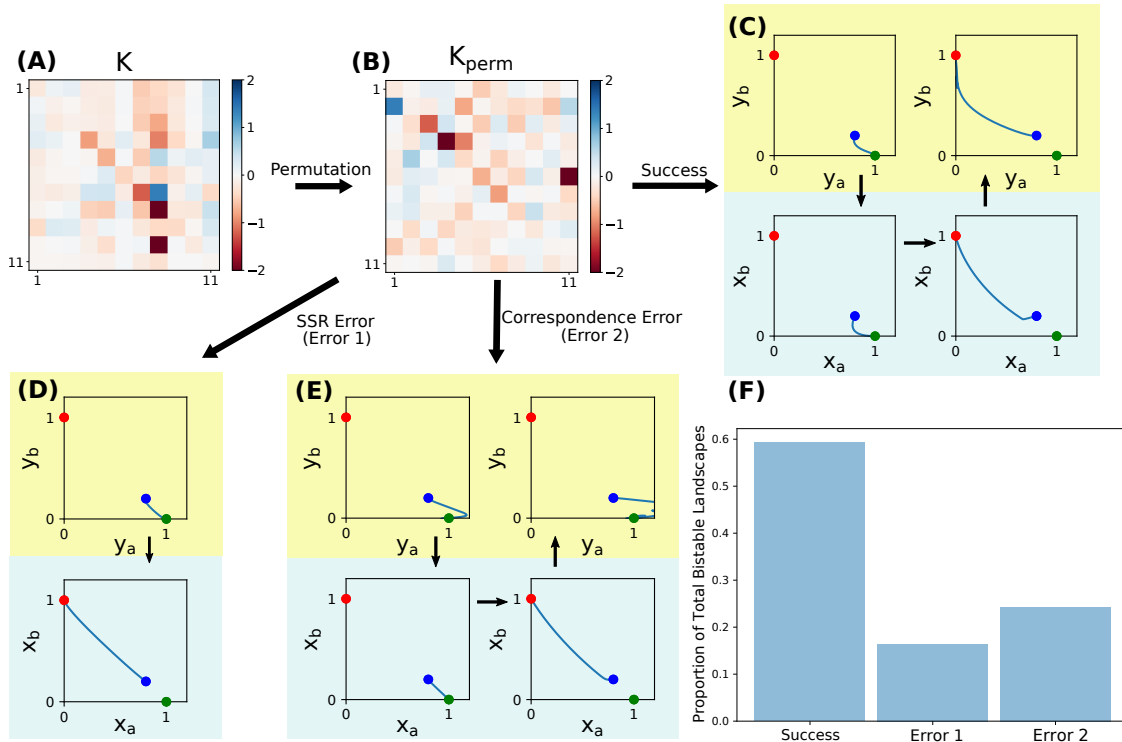


Figure 6.6: **SPARC is effective at modifying steady-state outcomes in synthetic gLV models.** (A, B) The interaction matrix K from the CDI model is randomly permuted to generate 100 synthetic parameter sets. From these 100 synthetic gLV systems, 140 bistable regions are identified. SPARC is applied to these synthetic models. (C) SPARC is considered successful if the parameter modification changes the trajectory of the initial condition so that it tends towards the target state (green), rather than the alternative state (red). (D, E) Two types of errors in SPARC are possible. SPARC can fail during the steady-state reduction process if the outcome of the high-dimensional system does not agree with the steady-state outcome of the reduced system (SSR Error, panel D). It can also fail if the high-dimensional parameter change ΔK does not appropriately alter the steady-state outcome (Correspondence Error, panel E). (F) SPARC successfully modifies 57% (77/136) of the synthetically-generated bistable landscapes. These numbers represent a baseline error rate of SPARC that may be further improved through manual intervention.

permuted parameter set, stable steady states are small in number. From 100 permuted gLV parameter sets, there are on average 0.8 completely stable steady states and 5.3 steady states with at most one unstable direction (i.e., steady states whose Jacobian matrices have at most one non-negative eigenvalue) per parameter set.

To ensure there are enough steady states to form bistable landscapes, we compute all 2^N steady states of each synthetic parameter set, then identify all steady states whose Jacobian has 0 or 1 positive eigenvalues in each parameter set, and use numerical simulations to test whether each steady state pair forms a bistable landscape. Specifically, for a steady state pair \vec{y}_a and \vec{y}_b , trajectories with initial conditions $(0.95\vec{y}_a + 0.05\vec{y}_b)$ and $(0.05\vec{y}_a + 0.95\vec{y}_b)$ are simulated to test whether they tend towards their nearest steady state. In addition, if initial conditions at $(0.8\vec{y}_a + 0.2\vec{y}_b)$ or $(0.2\vec{y}_a + 0.8\vec{y}_b)$ tend towards some other third steady state, the steady state pair is excluded. Out of the 100 synthetic parameter sets, a total of 136 bistable landscapes were identified.

In this context, SPARC is considered successful if it identifies high-dimensional interaction parameter changes that alter the steady-state outcome in a bistable system, as in Fig. 6.6C. This success relies on the correspondence between the 11D and 2D landscapes generated by SSR, the bifurcation analysis of the 2D system, and the correspondence between 2D and 11D parameters governed by the SSR formulae. Therefore, if an initial condition in both the unperturbed 11D and 2D models tends towards the same steady state, and the same initial condition in both the perturbed 11D and 2D models tend towards the other steady state in the bistable landscape, SPARC is considered successful.

To examine the fidelity of SPARC on synthetic parameter sets, it is applied to an ensemble of synthetically generated models. The two steady states of the synthetic bistable system are arbitrarily labeled as \vec{y}_a and \vec{y}_b . In Fig. 6.3 the initial condition was located at $(0.5\vec{y}_a + 0.5\vec{y}_b)$, but for these synthetic parameter sets the initial condition is located at $(0.2\vec{y}_a + 0.8\vec{y}_b)$. Since SSR is more accurate near the two steady states, this

choice of initial condition improves the success rate of SPARC. SPARC can fail at two steps, corresponding to the arrows A to B and C to D in the schematic Fig. 6.1. The first type of error occurs when SSR fails to preserve the steady state behavior of the gLV model; this error is demonstrated in Fig. 6.6D, where the high-dimensional initial condition tends towards steady state \vec{y}_b but the initial condition of the SSR-reduced model tends towards steady state \vec{x}_a . The second type of error occurs when associating the low-dimensional parameter change with a high-dimensional parameter change; this error is demonstrated in Fig. 6.6E, where the modified low-dimensional trajectory correctly tends towards steady state \vec{x}_a , but its corresponding high-dimensional trajectory erroneously tends towards steady state \vec{y}_b . Since the choice of a high-dimensional parameter change is degenerate, modifications to four interaction parameters K_{ij} corresponding to the four largest α_{ij} coefficients are tested. Small changes in the interaction matrix K will slightly change the location of the steady states, so this perturbation is turned off after the system has relaxed to the shifted steady states to allow the system to return to its original steady states.

Out of the 136 bistable landscapes generated from 100 synthetic parameter sets, SPARC successfully identified parameter modifications that led to the targeted transition between steady-state outcomes 57% (77/136) of the time. Details about specific errors rates occurred are provided in Fig. 6.6F: 17% (23/136) occurred during the SSR compression step, and 26% (36/136) occurred during the mapping from the 2D parameters to high-dimensional parameters. Manual intervention (e.g., trying different sizes of the prescribed parameter change) can improve this success rate. Therefore, SPARC is effective at altering steady-state behavior in generic gLV systems.

“Noisy” synthetic models

Finally, since inferring parameter values in gLV systems is an intrinsically noisy procedure, it is valuable to understand whether parameter changes generated by SPARC are robust to noise in the fitted parameters. We consider the “infection” scenario in which the CDI-susceptible state is the target state \vec{y}_a and the CDI-resistant state is the alternative state \vec{y}_b . Implementing the parameter change prescribed by SPARC ($\Delta K_{5,3} = 0.1744$) successfully drives an initial condition at $(0.5\vec{y}_a + 0.5\vec{y}_b)$ towards the target state \vec{y}_a . To test the robustness of the SPARC method, this parameter change is applied to synthetic “noisy” systems that are generated by independently scaling each interaction parameter K_{ij} by a number randomly drawn from a uniform distribution between $1 - \beta$ and $1 + \beta$. This parameter β is called the “noise.” In the Supplementary Information the following analysis is also performed using a parameter change 20% larger than the original one ($\Delta K_{5,3} = 0.2092$). This increment compensates for the deviation between the original and the SSR-generated separatrices.

The steady states of these synthetic systems are functions of the interaction parameters, and therefore differ from the steady states of the original CDI model. In gLV systems the presence/absence combination of species uniquely identifies a steady state, so it is straightforward to identify the two steady states in these noisy systems, called $\tilde{\vec{y}}_a$ and $\tilde{\vec{y}}_b$, that correspond to the target and alternative steady states \vec{y}_a and \vec{y}_b of the original CDI model. Many of these newly-generated steady states are biologically unreasonable: for a noise of $\beta = 0.025$, nearly half of the noisy steady states $\tilde{\vec{y}}_a$ contain negative entries. Additional details regarding the deviations of the noisy steady states as a function of the noise β are provided in the Supplementary Information.

We only consider noisy synthetic systems (i) that do not contain any negative entries in the steady states $\tilde{\vec{y}}_a$ and $\tilde{\vec{y}}_b$, and (ii) in which an initial condition at $(0.5\vec{y}_a + 0.5\vec{y}_b)$ flows

towards the alternative state \tilde{y}_b in the absence of any intervention (note that the initial condition is based on steady states of the original CDI model). Then, the parameter change $\Delta K_{5,3} = 0.1744$ is applied to the noisy models for an initial condition $(0.5\tilde{y}_a + 0.5\tilde{y}_b)$; if the system flows towards the target state \tilde{y}_a the parameter change is considered successful, and if it does not it is considered an error. The error rate of SPARC as a function of the noise β is plotted in the Supplementary Information. For each noise value, 1000 synthetic systems are created to generate statistics for the error rates. Using the original parameter change $\Delta K_{5,3} = 0.1744$ works well for very small noise values ($\beta < 0.005$), rapidly increases to an error rate of 40% with a noise of $\beta = 0.02$, and eventually approaches 80% for a noise of $\beta = 0.5$. With the incremented parameter change of $\Delta K_{5,3} = 0.2092$, the CDI model works nearly perfectly for synthetic systems with small noise values ($\beta < 0.02$). Then, as the noise increases the error rate worsens: a noise of $\beta = 0.1$ corresponds to a 30% error rate, and a noise of $\beta = 0.5$ approaches an 80% error rate. These analyses indicate that interventions generated by SPARC are effective for gLV systems whose parameters are known precisely, but are less effective when parameters are relatively unconstrained. Taken another way, these results place a limit on the required accuracy of parameter estimation, beyond which point two measured systems will differ enough in their parameter values that they diverge in their behavior.

6.5 Discussion

6.5.1 SPARC is efficient and flexible

SPARC generates a 2D gLV model to guide high-dimensional parameter modifications that alter the system outcome. Without such a guide, this parameter change must be instead selected through trial-and-error. A study about T-cell cancer networks used this

exhaustive trial-and-error method to find parameter perturbations that drive the system between attractors, but it was computationally expensive to search their parameter space [122]. In gLV systems, the number of computations needed for this trial-and-error method grows as $\mathcal{O}(N^2)$, where N is the number of species in the gLV model. As the number of species N becomes large, the exhaustive method becomes computationally intractable.

Rather than exploring the N^2 -dimensional parameter space of K , SPARC allows exploration of a 2-dimensional subspace of M associated with the bistable dynamics of interest of the high-dimensional model. In the SSR-generated 2D model, parameter modifications are analytically tractable using bifurcation analysis, which determines the sign of the parameter change according to the direction of the required separatrix shift. After the 2D model parameter change is determined, SSR formulae provide a direct correspondence between the 2D and the high-dimensional parameter modifications that produce the same steady-state outcome. For example, since bistability is well-defined in the 2D gLV system, SSR reveals the interspecies feedbacks most responsible for bistability in the high-dimensional system.

Furthermore, SPARC is flexible enough to drive the dynamical system bidirectionally between steady states, as demonstrated in the infection and recovery scenarios. When one steady state in the bistable region is desirable, as in the clinically-motivated recovery scenario considered here, SPARC identifies both which parameter changes to avoid and which to perform in order to achieve the target outcome. Both types of parameter changes are informative when trying to prevent the system from tending toward an undesirable steady state.

Finally, we note that the applicability of the parameter changes recommended by SPARC is sensitive to the accuracy of the fitted gLV interaction parameters. For example, in the CDI system (as demonstrated by the noisy synthetic models), the SPARC parameter change becomes less effective as the noise in the interaction parameters in-

creases. This analysis quantifies the tolerable level of uncertainty in fitted interaction parameters before they result in fundamentally different classes of model behavior.

6.5.2 Perturbing ecological interactions indirectly controls steady-state outcomes

Direct control methods modify the steady-state outcome of the gut microbiome by changing the state of the microbial system while retaining the same dynamical landscape. Implementations of this direct control method include bacteriotherapies such as Fecal Microbial Transplantation (FMT), which has been shown to be an effective treatment for *Clostridioides difficile* infection (CDI). FMT introduces a foreign microbial transplant that alters a host's microbiome composition, thereby ameliorating symptoms of CDI [103]. As realized in the gLV model, this amounts to an instantaneous shift in the microbial composition that moves the microbial state from one basin of attraction to another.

In contrast to this direct control method, SPARC indirectly controls the steady-state outcomes of a high-dimensional gLV model by modifying its dynamical landscape. Instead of adding foreign microbes, SPARC perturbs the interaction parameters of the gLV model, which we interpret as changing the environment in which the microbes live. Fig. 6.3 illustrates how this parameter-altering control method changes the steady-state outcome of a simulated gut microbial system.

SPARC could be applied to other ecological systems in order to attain a target community structure. In marine ecosystems, the target community structures may correspond to ecological states without harmful algal blooms or invasive fish species. In these cases, environmental factors such as the abundance of chemical fertilizers or pesticides, the pH, and the velocity of stream flows influence the state of the ecosystem [158, 142]. Previously, algal blooms and population dynamics of invasive fish species have been modeled

with gLV systems [159, 160]. Therefore, SPARC could provide a systematic framework that guides environmental interventions to remove harmful algae or invasive fish species.

SPARC identifies a single entry in a high-dimensional interaction matrix that can be altered to change the system behavior. However, it might not be possible in practice to identify environmental factors that, when modified, change only one entry of the interaction matrix. Importantly, the parameter entry generated by SPARC is not unique, as shown in Eq. (6.7). As a result, it is possible to find a linear combination of changes in the environmental factors that maximize the parameter changes in the most effective entries (i.e., entries with the largest α_{ij} values) and minimize other changes, especially the most effective entries in the opposite direction. This more complex parameter change can then be simulated to assess its effectiveness.

6.5.3 SPARC provides a lens for understanding the effect of the environment on microbial composition

Having demonstrated the effectiveness of SPARC *in silico*, it would be valuable to verify this method in an experimental model system of the microbiome. SPARC relies on changing interactions between microbial species in the gut microbiome, which could be achieved by deliberately changing environmental factors in a controlled experimental setting. Therefore, any realization of this method would require an experimental microbiome model of limited microbial diversity that allows the manipulation of oxygen levels, nutrient availability, or other factors. One such experimental model might be the intestine-on-a-chip system, which simulates the human gut microbiome in a manipulable *in vitro* environment [161, 162, 163]. By fitting gLV models to time-series data from the intestine-on-a-chip, it may be possible to isolate the effect of environmental perturbations and identify the corresponding interaction matrix change ΔK underlying SPARC.

In real microbial systems, changes in environmental factors potentially affect the interactions between many species, thus changing multiple interaction parameters at a time. For example, Lin *et al.* found that four dominant bacterial genera with carbon assimilation pathways gain ecological advantages when there is a lack of dissolved carbon in the environment [164]. Therefore, environmental changes such as the removal of dissolved carbon will alter the effective microbe-microbe interactions between these species. In cases such as these, SPARC could systematically specify how environmental changes alter the dynamical landscape.

In future applications, the environmental degrees of freedom will be as myriad as diet, designer probiotics, or designer prebiotics. The combinatorial complexity of these contributions will require a systematic framework, such as SPARC, in order to understand how to drive the system towards a target state. Once environmental interventions are associated with changes in species-species interaction parameters in gLV models, SPARC could help predict how environmental changes affect gut microbiome compositions.

6.6 Conclusion

SPARC controls the steady-state outcome of bistable regions in gLV systems by altering ecological interaction parameters. This method circumvents the computational task of performing numerical trials to exhaustively search a high-dimensional parameter space. Instead, SPARC uses a recently-developed dimensionality-reduction technique to reduce the problem to searching a 2-dimensional parameter subspace. Consequently, we are able to efficiently and systematically identify a minimal parameter change that results in desired system behavior.

SPARC provides a novel alternative to canonical control methods that modify the system state directly. SPARC instead focuses on how environmental factors and microbial

interactions dictate microbial dynamics. Eventually, indirect and direct methods could be used in conjunction to provide a comprehensive framework for the control of ecological systems.

Chapter 7

Conclusion

7.1 Overview

In my doctoral research I employed novel theoretical approaches and tractable experimental systems to study simplified instances of complex microbial dynamics in the microbiome. This research was motivated by a desire to inform the mechanism of action and development of bacteriotherapies. This thesis applied an engineering perspective—that a system is fully understood only when it can be controlled—to an ecological system, and presented theory that pertains to the development of direct and indirect bacteriotherapies.

To manage the complexity of high-dimensional ecological models, we introduced steady-state reduction, a dimensionality-reduction technique that maps bistable dynamics in a high-dimensional generalized Lotka-Volterra (gLV) system to an approximate reduced system. The correspondence between the high-dimensional system and this analytically-tractable reduced system enables us to identify the dynamics of the high-dimensional system that are most pertinent for the system’s eventual steady state outcome. This compressed system significantly reduces the space of control parameters that

need to be searched by potential interventions, and therefore facilitates the development of direct and indirect bacteriotherapies.

7.2 Future work: A statistical mechanical treatment of the microbiome

Our results have not yet been tested experimentally, but in the future I am excited to validate these theories in the fruit fly microbiome. Before these results are applied to real experimental systems, a few additional considerations are needed to refine our theoretical implementation of direct and indirect bacteriotherapies. First, biological systems by nature are non-equilibrium and full of noise, but gLV systems are deterministic (and we have assumed that ecological dynamics equilibrate towards point attractors). Thus, statistical mechanics must be incorporated into our models to properly describe the stochasticity of microbial dynamics. Second, driving an ecological system between states requires an understanding of its out-of-equilibrium dynamics, and these dynamics— especially in thermodynamic systems— are difficult to characterize.

As a postdoctoral researcher, I plan to address these challenges by applying recently developed methods from out-of-equilibrium thermodynamics to ecological systems. The microbiome is a natural candidate for this statistical mechanical treatment, since it is composed of many stochastic microscopic units that individually are governed by relatively simple rules, but as a whole exhibit robust collective behaviors [165, 166].

The first hurdle will be to draw a correspondence between thermodynamic systems (which have a well-defined conserved energy) and ecological systems (which do not). Ecological dynamics are typically separated into two regimes. The first regime is niche theory, in which species deterministically colonize a specific habitat, leading to the prin-

principle of competitive exclusion in which only the fittest species survives. The alternative regime is neutral theory, which supposes that the colonization of species on the same trophic level is entirely stochastic, and that competition does not play a factor. Naturally, ecologists agree that the true ecological dynamics lie somewhere in-between these two extremes [167, 98]. In statistical mechanics, systems are biased but not forced to occupy a low-energy state—this behavior is similar to ecological systems, which are biased (niche) but not forced (neutral) to form stable ecological networks of species.

As an example of these two regimes, in my research with experimentalist Will Ludington examining variability in the fruit fly microbiome, we find that colonization of the fruit fly microbiome is probabilistic, and that there exist clear signatures of bacterial species that are “strong” and “weak” colonizers of the microbiome [27]. Thus, even this simple model system has evident roles for both stochastic variability and deterministic competition in ecological systems.

It therefore seems natural to apply methods from statistical mechanics to ecology, but the complexity of ecological systems— which are much more complicated than the uniform bouncing molecules that form the basis for thermodynamics— has constrained previous attempts, which typically introduce a statistical flair by studying ecological models with randomly drawn parameters. These existing approaches specify the breadth of behaviors that an ecological model can exhibit, and they often permit insightful analytic analyses, but they do not incorporate the competing roles of niche and neutral behaviors [168, 169]. By interpreting ecological dynamics as a thermodynamic system, our proposed method will inherently reflect these deterministic and stochastic components.

The population dynamics of microbes within the microbiome are mediated by a complex network of biochemical interactions (e.g., cellular metabolism) performed by microbial and host cells [101, 102]. Ecological models seek to consolidate these myriad biochemical mechanisms into nonspecific coefficients that characterize the interactions

between microbial populations. When these ecological models relax towards equilibrium states, their dynamics may be construed as following an ecological dynamical landscape “downhill” (in dynamical systems theory, this landscape is called a Lyapunov function). In traditional statistical mechanics, the state of a system is dependent on an energy landscape, with the states of lower energy being preferentially abundant. We will seek to find a correspondence between ecology and nonequilibrium statistical mechanics by formally associating these two landscapes.

Once the problem has been posed in the language of out-of-equilibrium thermodynamics, I can employ the non-equilibrium control theory developed by my postdoctoral research advisor David Sivak. This control theory designs optimal protocols that drive a microscopic system between states, for example informing the function and thermodynamic efficiency of the molecular machine ATP synthase, and describing how torque can be applied to ATP synthase to be maximally efficient [170]. By applying these methods to ecological systems via the correspondence between Lyapunov and energy landscapes, they could inform how to drive a ecosystem between states with an ecological control protocol (for example through culling, quarantining, or transplantation of species populations). Since David’s research deals with abstract control parameters, his methods should be agnostic to whether the controls are microscopic torques or the introduction of a new species. Even though his research has traditionally functioned at time and length scales disparate from ecological dynamics, we expect that parallel approaches will be fruitful when applied to an ecological domain.

7.3 Future work: Experimental validation in the fruit fly microbiome

Finally, to bridge the gap between theory and experiment, we will test our approaches in the fruit fly experimental system with Will Ludington. Success criteria for our approach will be based on whether we develop a mathematical formalism that explicitly captures niche/neutral ecological dynamics in a statistical mechanical framework, and experimentally observe that our constructed theory properly balances the niche and neutral components of ecological systems in the fruit fly model system. If we do not observe these niche/neutral behaviors, it will indicate that our approach is too coarse-grained: for example, it might signal that we need to consider nutrient concentrations and fly gut geometry in addition to microbial abundances. In this case, we will proceed by running experiments in a turbidostat (a simpler though less clinically-relevant experimental system in which microbes grow in a fixed concentration gradient) rather than in the living fruit fly, which should provide a sufficiently well-characterized experimental system for our research approach to succeed.

7.4 Closing remarks

It is a prosperous time for microbial ecology. Ecological models are more valuable than ever in their ability to explain the wealth of microbial data that is now available. More and more clinical evidence linking microbiome composition to host physiology is established each week. And many clinicians are captivated by the therapeutic potential of bacteriotherapies, and are actively pursuing clinical studies to test them on all sorts of diseases that are associated with impaired microbiome compositions.

Eventually, “medicine for the microbiome” might become widespread and provide

people with personalized recommendations that improve their health. Before this can occur, ecological models of the microbiome and mechanistic models of bacteriotherapies must be created. Crucially, these theories will depend on the union of engineering and ecology, well-characterized experimental systems, and years of incremental progress. The contents of this thesis contribute to this formidable goal, and will serve as a foundation for future progress.

Bibliography

- [1] L. J. Brandt, O. C. Aroniadis, M. Mellow, A. Kanatzar, C. Kelly, T. Park, N. Stollman, F. Rohlke, and C. Surawicz, *Long-term follow-up of colonoscopic fecal microbiota transplant for recurrent Clostridium difficile infection*, *Am J Gastroenterol* **107** (2012), no. 7 1079–1087.
- [2] M. Coryell, M. McAlpine, N. V. Pinkham, T. R. McDermott, and S. T. Walk, *The gut microbiome is required for full protection against acute arsenic toxicity in mouse models*, *Nature Communications* **9** (2018), no. 1 5424.
- [3] E. A. Franzosa, A. Sirota-Madi, J. Avila-Pacheco, N. Fornelos, H. J. Haiser, S. Reinker, T. Vatanen, A. B. Hall, H. Mallick, L. J. McIver, *et. al.*, *Gut microbiome structure and metabolic activity in inflammatory bowel disease*, *Nature Microbiology* (2018) 1.
- [4] L. M. De Leon, J. B. Watson, and C. R. Kelly, *Transient flare of ulcerative colitis after fecal microbiota transplantation for recurrent Clostridium difficile infection*, *Clinical Gastroenterology and Hepatology* **11** (2013), no. 8 1036–1038.
- [5] J. Anderson, R. Edney, and K. Whelan, *Systematic review: faecal microbiota transplantation in the management of inflammatory bowel disease*, *Alimentary Pharmacology & Therapeutics* **36** (2012), no. 6 503–516.
- [6] C. G. Buffie, I. Jarchum, M. Equinda, L. Lipuma, A. Gobourne, A. Viale, C. Ubeda, J. Xavier, and E. G. Pamer, *Profound alterations of intestinal microbiota following a single dose of clindamycin results in sustained susceptibility to Clostridium difficile-induced colitis*, *Infect Immun* **80** (2012), no. 1 62–73.
- [7] R. R. Stein, V. Bucci, N. C. Toussaint, C. G. Buffie, G. Räscht, E. G. Pamer, C. Sander, and J. B. Xavier, *Ecological modeling from time-series inference: Insight into dynamics and stability of intestinal microbiota*, *PLoS Comput Biol* **9** (2013), no. 12 1–11.
- [8] A. L. Gould, V. Zhang, L. Lamberti, E. W. Jones, B. Obadia, N. Korasidis, A. Gavryushkin, J. M. Carlson, N. Beerewinkel, and W. B. Ludington, *Microbiome interactions shape host fitness*, *Proceedings of the National Academy of Sciences* **115** (2018), no. 51 E11951–E11960.

- [9] E. W. Jones and J. M. Carlson, *In silico analysis of antibiotic-induced Clostridium difficile infection: Remediation techniques and biological adaptations*, *PLoS Computational Biology* **14** (2018), no. 2 1–24.
- [10] E. W. Jones and J. M. Carlson, *Steady-state reduction of generalized Lotka-Volterra systems in the microbiome*, *Phys. Rev. E* **99** (2019) 032403.
- [11] E. W. Jones, P. Shankin-Clarke, and J. M. Carlson, *Navigation and control of outcomes in a generalized Lotka-Volterra model of the microbiome*, [arXiv:2003.1295](https://arxiv.org/abs/2003.1295).
- [12] Z. Wang, E. W. Jones, J. M. Mueller, and J. M. Carlson, *Control of ecological outcomes through deliberate parameter changes in a model of the gut microbiome*, *Phys. Rev. E* **101** (2020) 052402.
- [13] R. I. Clark, A. Salazar, R. Yamada, S. Fitz-Gibbon, M. Morselli, J. Alcaraz, A. Rana, M. Rera, M. Pellegrini, W. J. William, *et. al.*, *Distinct shifts in microbiota composition during Drosophila aging impair intestinal function and drive mortality*, *Cell Reports* **12** (2015), no. 10 1656–1667.
- [14] P. Smith, D. Willemsen, M. Popkes, F. Metge, E. Gandiwa, M. Reichard, and D. R. Valenzano, *Regulation of life span by the gut microbiota in the short-lived African turquoise killifish*, *eLife* **6** (2017) e27014.
- [15] L. M. Travers, F. Garcia-Gonzalez, and L. W. Simmons, *Live fast die young life history in females: evolutionary trade-off between early life mating and lifespan in female Drosophila melanogaster*, *Scientific Reports* **5** (2015) 15469.
- [16] D. A. Roff and D. Fairbairn, *The evolution of trade-offs: where are we?*, *Journal of Evolutionary Biology* **20** (2007), no. 2 433–447.
- [17] S. C. Stearns, *Trade-offs in life-history evolution*, *Functional Ecology* **3** (1989), no. 3 259–268.
- [18] A. J. Van Noordwijk and G. de Jong, *Acquisition and allocation of resources: their influence on variation in life history tactics*, *The American Naturalist* **128** (1986), no. 1 137–142.
- [19] I. Cho and M. J. Blaser, *The human microbiome: at the interface of health and disease*, *Nature Reviews Genetics* **13** (2012), no. 4 260–270.
- [20] P. H. Degan, M. E. Taga, and A. L. Goodman, *Vitamin B12 as a modulator of gut microbial ecology*, *Cell Metabolism* **20** (2014), no. 5 769–778.
- [21] E. C. Seth and M. E. Taga, *Nutrient cross-feeding in the microbial world*, *Frontiers in Microbiology* **5** (2014) 350.

- [22] A. S. Rolig, R. Parthasarathy, A. R. Burns, B. J. Bohannon, and K. Guillemin, *Individual members of the microbiota disproportionately modulate host innate immune responses*, *Cell Host & Microbe* **18** (2015), no. 5 613–620.
- [23] J.-H. Ryu, S.-H. Kim, H.-Y. Lee, J. Y. Bai, Y.-D. Nam, J.-W. Bae, D. G. Lee, S. C. Shin, E.-M. Ha, and W.-J. Lee, *Innate immune homeostasis by the homeobox gene caudal and commensal-gut mutualism in Drosophila*, *Science* **319** (2008), no. 5864 777–782.
- [24] C. N. A. Wong, P. Ng, and A. E. Douglas, *Low-diversity bacterial community in the gut of the fruitfly Drosophila melanogaster*, *Environmental Microbiology* **13** (2011), no. 7 1889–1900.
- [25] J. A. Chandler, J. M. Lang, S. Bhatnagar, J. A. Eisen, and A. Kopp, *Bacterial communities of diverse Drosophila species: ecological context of a host–microbe model system*, *PLoS Genetics* **7** (2011), no. 9.
- [26] J. E. Blum, C. N. Fischer, J. Miles, and J. Handelsman, *Frequent replenishment sustains the beneficial microbiome of Drosophila melanogaster*, *MBio* **4** (2013), no. 6 e00860–13.
- [27] B. Obadia, Z. T. Güvener, V. Zhang, J. A. Ceja-Navarro, E. L. Brodie, W. J. William, and W. B. Ludington, *Probabilistic invasion underlies natural gut microbiome stability*, *Current Biology* **27** (2017), no. 13 1999–2006.
- [28] I. S. Pais, R. S. Valente, M. Sporniak, and L. Teixeira, *Drosophila melanogaster establishes a species-specific mutualistic interaction with stable gut-colonizing bacteria*, *PLoS Biol.* **16** (2018), no. 7 e2005710.
- [29] F. Staubach, J. F. Baines, S. Künzel, E. M. Bik, and D. A. Petrov, *Host species and environmental effects on bacterial communities associated with Drosophila in the laboratory and in the natural environment*, *PLoS One* **8** (2013), no. 8.
- [30] K. L. Adair, M. Wilson, A. Bost, and A. E. Douglas, *Microbial community assembly in wild populations of the fruit fly Drosophila melanogaster*, *The ISME Journal* **12** (2018), no. 4 959–972.
- [31] P. D. Newell and A. E. Douglas, *Interspecies interactions determine the impact of the gut microbiota on nutrient allocation in Drosophila melanogaster*, *Appl. Environ. Microbiol.* **80** (2014), no. 2 788–796.
- [32] L. E. O’Brien, S. S. Soliman, X. Li, and D. Bilder, *Altered modes of stem cell division drive adaptive intestinal growth*, *Cell* **147** (2011), no. 3 603–614.

- [33] A. K. Chippindale, A. M. Leroi, S. B. Kim, and M. R. Rose, *Phenotypic plasticity and selection in Drosophila life-history evolution. i. nutrition and the cost of reproduction*, in *Methuselah flies: A case study in the evolution of aging*, pp. 122–144. World Scientific, 2004.
- [34] P. H. Leslie, *On the use of matrices in certain population mathematics*, *Biometrika* **33** (1945), no. 3 183–212.
- [35] R. C. Team, *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org> (2017).
- [36] A. C. Wong, J. M. Chaston, and A. E. Douglas, *The inconstant gut microbiota of Drosophila species revealed by 16S rRNA gene analysis*, *The ISME Journal* **7** (2013), no. 10 1922–1932.
- [37] B. Obadia, E. S. Keebaugh, R. Yamada, W. B. Ludington, and W. J. William, *Diet influences host–microbiota associations in Drosophila*, *Proceedings of the National Academy of Sciences* **115** (2018), no. 20 E4547–E4548.
- [38] C. N. Fischer, E. P. Trautman, J. M. Crawford, E. V. Stabb, J. Handelsman, and N. A. Broderick, *Metabolite exchange between microbiome members produces compounds that influence Drosophila behavior*, *eLife* **6** (2017) e18855.
- [39] R. T. Paine, *Food-web analysis through field measurement of per capita interaction strength*, *Nature* **355** (1992), no. 6355 73–75.
- [40] J. Bascompte, P. Jordano, and J. M. Olesen, *Asymmetric coevolutionary networks facilitate biodiversity maintenance*, *Science* **312** (2006), no. 5772 431–433.
- [41] G. Gellner and K. S. McCann, *Consistent role of weak and strong interactions in high-and low-diversity trophic food webs*, *Nature Communications* **7** (2016), no. 1 1–7.
- [42] J. Grilli, G. Barabás, M. J. Michalska-Smith, and S. Allesina, *Higher-order interactions stabilize dynamics in competitive network models*, *Nature* **548** (2017), no. 7666 210–213.
- [43] H. Inamine, S. P. Ellner, P. D. Newell, Y. Luo, N. Buchon, and A. E. Douglas, *Spatiotemporally heterogeneous population dynamics of gut bacteria inferred from fecal time series data*, *MBio* **9** (2018), no. 1 e01453–17.
- [44] D. Fast, B. Kostiuik, E. Foley, and S. Pukatzki, *Commensal pathogen competition impacts host viability*, *Proceedings of the National Academy of Sciences* **115** (2018), no. 27 7099–7104.

- [45] R. Yamada, S. A. Deshpande, K. D. Bruce, E. M. Mak, and W. J. William, *Microbes promote amino acid harvest to rescue undernutrition in Drosophila*, *Cell Reports* **10** (2015), no. 6 865–872.
- [46] W. W. Ja, G. B. Carvalho, M. Madrigal, R. W. Roberts, and S. Benzer, *The Drosophila g protein-coupled receptor, methuselah, exhibits a promiscuous response to peptides*, *Protein Science* **18** (2009), no. 11 2203–2208.
- [47] T. Brummel, A. Ching, L. Seroude, A. F. Simon, and S. Benzer, *Drosophila lifespan enhancement by exogenous bacteria*, *Proceedings of the National Academy of Sciences* **101** (2004), no. 35 12974–12979.
- [48] A. W. Walters, R. C. Hughes, T. B. Call, C. J. Walker, H. Wilcox, S. C. Petersen, S. M. Rudman, P. D. Newell, A. E. Douglas, P. S. Schmidt, *et. al.*, *The microbiota influences the Drosophila melanogaster life history strategy*, *Molecular Ecology* **29** (2020), no. 3 639–653.
- [49] J. L. Round and S. K. Mazmanian, *The gut microbiome shapes intestinal immune responses during health and disease*, *Nat Rev Immunol* **9** (2009), no. 5 313–323.
- [50] R. A. Britton and V. B. Young, *Interaction between the intestinal microbiota and host in Clostridium difficile colonization resistance*, *Trends Microbiol* **20** (2012), no. 7 313–319.
- [51] P. McKenney and E. Pamer, *From hype to hope: The gut microbiota in enteric infectious disease*, *Cell* **163** (2015), no. 6 1326 – 1332.
- [52] I. Cho and M. J. Blaser, *The human microbiome: at the interface of health and disease*, *Nat Rev Genet* **13** (2012), no. 4 260–270.
- [53] C. M. Theriot and V. B. Young, *Interactions between the gastrointestinal microbiome and Clostridium difficile*, *Annu Rev Microbiol* **69** (2015) 445–461.
- [54] A. M. Schubert, M. A. M. Rogers, C. Ring, J. Mogle, J. P. Petrosino, V. B. Young, D. M. Aronoff, and P. D. Schloss, *Microbiome data distinguish patients with Clostridium difficile infection and non-C. difficile-associated diarrhea from healthy controls*, *mBio* **5** (2014), no. 3 e01021–14.
- [55] M. Kircher and J. Kelso, *High-throughput DNA sequencing – concepts and limitations*, *BioEssays* **32** (2010), no. 6 524–536.
- [56] J. G. Caporaso, J. Kuczynski, J. Stombaugh, K. Bittinger, F. D. Bushman, E. K. Costello, N. Fierer, A. G. Peña, J. K. Goodrich, J. I. Gordon, G. A. Huttley, S. T. Kelley, D. Knights, J. E. Koenig, R. E. Ley, C. A. Lozupone, D. McDonald, B. D. Muegge, M. Pirrung, J. Reeder, J. R. Sevinsky, P. J. Turnbaugh, W. A. Walters, J. Widmann, T. Yatsunenko, J. Zaneveld, and R. Knight, *Qiime allows analysis of high-throughput community sequencing data*, *Nat Methods* **7** (2010), no. 5 335–336.

- [57] Y. Takeuchi, *Global Dynamical Properties of Lotka-Volterra Systems*. World Scientific, 1996.
- [58] M. Arumugam, J. Raes, E. Pelletier, D. Le Paslier, T. Yamada, D. R. Mende, G. R. Fernandes, J. Tap, T. Bruls, J.-M. Batto, M. Bertalan, N. Borruel, F. Casellas, L. Fernandez, L. Gautier, T. Hansen, M. Hattori, T. Hayashi, M. Kleerebezem, K. Kurokawa, M. Leclerc, F. Levenez, C. Manichanh, H. B. Nielsen, T. Nielsen, N. Pons, J. Poulain, J. Qin, T. Sicheritz-Ponten, S. Tims, D. Torrents, E. Ugarte, E. G. Zoetendal, J. Wang, F. Guarner, O. Pedersen, W. M. de Vos, S. Brunak, J. Dore, J. Weissenbach, S. D. Ehrlich, and P. Bork, *Enterotypes of the human gut microbiome*, *Nature* **473** (2011), no. 7346 174–180.
- [59] C. G. Buffie, V. Bucci, R. R. Stein, P. T. McKenney, L. Ling, A. Gobourne, D. No, H. Liu, M. Kinnebrew, A. Viale, E. Littmann, M. R. M. van den Brink, R. R. Jenq, Y. Taur, C. Sander, J. Cross, N. C. Toussaint, J. B. Xavier, and E. G. Pamer, *Precision microbiome restoration of bile acid-mediated resistance to *Clostridium difficile**, *Nature* **517** (2015), no. 7533 205–208.
- [60] V. Bucci and J. B. Xavier, *Towards predictive models of the human gut microbiome*, *Journal of Molecular Biology* **426** (2014), no. 23 3907 – 3916.
- [61] A. Rineh, M. J. Kelso, F. Vatansever, G. P. Tegos, and M. R. Hamblin, **Clostridium difficile* infection: molecular pathogenesis and novel therapeutics*, *Expert Rev Anti Infect Ther* **12** (2014), no. 1 131–150.
- [62] S. Cohen, D. Gerding, S. Johnson, C. Kelly, V. Loo, L. McDonald, J. Pepin, and M. Wilcox, *Clinical practice guidelines for *Clostridium difficile* infection in adults: 2010 update by the Society for Healthcare Epidemiology of America (SHEA) and the Infectious Diseases Society of America (IDSA)*, *Infection Control and Hospital Epidemiology* **31** (2010), no. 5 431–455.
- [63] T. L. A. Nguyen, S. Vieira-Silva, A. Liston, and J. Raes, *How informative is the mouse for human gut microbiota research?*, *Disease Models and Mechanisms* **8** (2015), no. 1 1–16, [<http://dmm.biologists.org/content/8/1/1.full.pdf>].
- [64] B. D. Muegge, J. Kuczynski, D. Knights, J. C. Clemente, A. González, L. Fontana, B. Henrissat, R. Knight, and J. I. Gordon, *Diet drives convergence in gut microbiome functions across mammalian phylogeny and within humans*, *Science* **332** (2011), no. 6032 970–974.
- [65] Clarence L. Lehman and David Tilman, *Biodiversity, stability, and productivity in competitive communities.*, *The American Naturalist* **156** (2000), no. 5 534–552.
- [66] P. Hemarajata and J. Versalovic, *Effects of probiotics on gut microbiota: mechanisms of intestinal immunomodulation and neuromodulation*, *Therap Adv Gastroenterol* **6** (2013), no. 1 39–51.

- [67] P. J. Turnbaugh, R. E. Ley, M. A. Mahowald, V. Magrini, E. R. Mardis, and J. I. Gordon, *An obesity-associated gut microbiome with increased capacity for energy harvest*, *Nature* **444** (2006), no. 7122 1027–131.
- [68] D. A. Burns and N. P. Minton, *Sporulation studies in Clostridium difficile*, *Journal of Microbiological Methods* **87** (2011), no. 2 133 – 138.
- [69] M. C. Abt, P. T. McKenney, and E. G. Pamer, *Clostridium difficile colitis: pathogenesis and host defence*, *Nat Rev Micro* **14** (2016), no. 10 609–620.
- [70] A. S. Perelson and P. W. Nelson, *Mathematical analysis of HIV-I: Dynamics in vivo*, *SIAM Review* **41** (1999), no. 1 3–44.
- [71] J. B. Deris, M. Kim, Z. Zhang, H. Okano, R. Hermsen, A. Groisman, and T. Hwa, *The innate growth bistability and fitness landscapes of antibiotic resistant bacteria*, *Science* **342** (2013), no. 6162 1237435–1237435.
- [72] P. Spigaglia, *Recent advances in the understanding of antibiotic resistance in Clostridium difficile infection*, *Ther Adv Infect Dis* **3** (2016), no. 1 23–42.
- [73] E. Ibargüen-Mondragón, S. Mosquera, M. Cerón, E. M. Burbano-Rosero, S. P. Hidalgo-Bonilla, L. Esteva, and J. P. Romero-Leitón, *Mathematical modeling on bacterial resistance to multiple antibiotics caused by spontaneous mutations*, *Biosystems* **117** (2014) 60 – 67.
- [74] I. H. Spicknall, B. Foxman, C. F. Marrs, and J. N. S. Eisenberg, *A modeling framework for the evolution and spread of antibiotic resistance: Literature review and model categorization*, *Am J Epidemiol* **178** (2013), no. 4 508–520.
- [75] L. Opatowski, D. Guillemot, P.-Y. Boëlle, and L. Temime, *Contribution of mathematical modeling to the fight against bacterial antibiotic resistance*, *Current Opinion in Infectious Diseases* **24** (2011), no. 3.
- [76] P. Spigaglia, F. Barbanti, P. Mastrantonio, , G. Ackermann, C. Balmelli, F. Barbut, E. Bouza, J. Brazier, M. Delmée, D. Drudy, E. Kuijper, H. Ladas, P. Mastrantonio, E. Nagy, H. Pituch, I. Poxton, M. Rupnik, M. Wullt, and M. Yücesoy, *Multidrug resistance in european Clostridium difficile clinical isolates*, *Journal of Antimicrobial Chemotherapy* **66** (2011), no. 10 2227.
- [77] E. Jones, P. Roemer, M. Raghupathi, and S. Pankavich, *Analysis and Simulation of the Three-Component Model of HIV Dynamics*, *SIAM Undergraduate Research Online* (2013) [[arXiv:1312.3671](https://arxiv.org/abs/1312.3671)].
- [78] L. Edelstein-Keshet, *Mathematical Models in Biology*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2005.

- [79] L. V. McFarland, G. W. Elmer, and C. M. Surawicz, *Breaking the cycle: treatment strategies for 163 cases of recurrent Clostridium difficile disease*, *Am J Gastroenterol* **97** (2002), no. 7 1769–1775.
- [80] M. J. Koenigsnecht, C. M. Theriot, I. L. Bergin, C. A. Schumacher, P. D. Schloss, and V. B. Young, *Dynamics and establishment of Clostridium difficile infection in the murine gastrointestinal tract*, *Infect Immun* **83** (2015), no. 3 934–941.
- [81] F. Gil, S. Lagos-Moraga, P. Calderón-Romero, M. Pizarro-Guajardo, and D. Paredes-Sabja, *Updates on Clostridium difficile spore biology*, *Anaerobe* **45** (2017), no. Supplement C 3–9.
- [82] D. I. Andersson and B. R. Levin, *The biological cost of antibiotic resistance*, *Current Opinion in Microbiology* **2** (1999), no. 5 489–493.
- [83] A. H. Melnyk, A. Wong, and R. Kassen, *The fitness costs of antibiotic resistance mutations*, *Evolutionary Applications* **8** (2015), no. 3 273–283.
- [84] D. I. Andersson and D. Hughes, *Antibiotic resistance and its cost: is it possible to reverse resistance?*, *Nat Rev Micro* **8** (2010), no. 4 260–271.
- [85] J. L. Martinez and F. Baquero, *Mutation frequencies and antibiotic resistance*, *Antimicrob Agents Chemother* **44** (2000), no. 7 1771–1777.
- [86] Y. Wakamoto, N. Dhar, R. Chait, K. Schneider, F. Signorino-Gelo, S. Leibler, and J. D. McKinney, *Dynamic persistence of antibiotic-stressed mycobacteria*, *Science* **339** (2013), no. 6115 91–95.
- [87] G. I. Lang, D. P. Rice, M. J. Hickman, E. Sodergren, G. M. Weinstock, D. Botstein, and M. M. Desai, *Pervasive genetic hitchhiking and clonal interference in 40 evolving yeast populations*, *Nature* **500** (2013), no. 7464 571–574.
- [88] J. S. Bakken, T. Borody, L. J. Brandt, J. V. Brill, D. C. Demarco, M. A. Franzos, C. Kelly, A. Khoruts, T. Louie, L. P. Martinelli, T. A. Moore, G. Russell, and C. Surawicz, *Treating Clostridium difficile infection with fecal microbiota transplantation*, *Clin Gastroenterol Hepatol* **9** (2011), no. 12 1044–1049.
- [89] D. Barnes and K. T. Park, *Donor considerations in fecal microbiota transplantation*, *Current Gastroenterology Reports* **19** (2017), no. 3 10.
- [90] T. J. Borody, S. Paramsothy, and G. Agrawal, *Fecal microbiota transplantation: Indications, methods, evidence, and future directions*, *Curr Gastroenterol Rep* **15** (2013), no. 8 337.

- [91] G. C. Batzias, G. A. Delis, and L. V. Athanasiou, *Clindamycin bioavailability and pharmacokinetics following oral administration of clindamycin hydrochloride capsules in dogs*, *The Veterinary Journal* **170** (2005), no. 3 339 – 345.
- [92] M. E. Klepser, M. A. Banevicius, R. Quintiliani, and C. H. Nightingale, *Characterization of bactericidal activity of clindamycin against bacteroides fragilis via kill curve methods.*, *Antimicrobial Agents and Chemotherapy* **40** (1996), no. 8 1941–1944.
- [93] R. R. Regoes, C. Wiuff, R. M. Zappala, K. N. Garner, F. Baquero, and B. R. Levin, *Pharmacodynamic functions: a multiparameter approach to the design of antibiotic treatment regimens*, *Antimicrobial Agents and Chemotherapy* **48** (2004), no. 10 3670–3676.
- [94] D. Tilman, *The importance of the mechanisms of interspecific competition*, *The American Naturalist* **129** (1987), no. 5 769–774.
- [95] P. J. Wangersky, *Lotka-Volterra population models*, *Annual Review of Ecology and Systematics* **9** (1978), no. 1 189–218.
- [96] R. Tibshirani, *Regression shrinkage and selection via the lasso*, *Journal of the Royal Statistical Society* **58** (1996), no. 1 267–288.
- [97] G. Lillacci and M. Khammash, *Parameter estimation and model selection in computational biology*, *PLoS Computational Biology* **6** (2010), no. 3 1–17.
- [98] C. K. Fisher and P. Mehta, *Identifying keystone species in the human gut microbiome from metagenomic timeseries using sparse linear regression*, *PLoS ONE* **9** (2014), no. 7 1–10.
- [99] C. J. Briggs, R. A. Knapp, and V. T. Vredenburg, *Enzootic and epizootic dynamics of the chytrid fungal pathogen of amphibians*, *Proceedings of the National Academy of Sciences* (2010).
- [100] J. Lloyd-Price, G. Abu-Ali, and C. Huttenhower, *The healthy human microbiome*, *Genome Medicine* **8** (2016), no. 1 51.
- [101] S. Widder, R. J. Allen, T. Pfeiffer, T. P. Curtis, C. Wiuf, W. T. Sloan, O. X. Cordero, S. P. Brown, B. Momeni, W. Shou, H. Kettle, H. J. Flint, A. F. Haas, B. Laroche, J.-U. Kreft, P. B. Rainey, S. Freilich, S. Schuster, K. Milferstedt, J. R. van der Meer, T. Großkopf, J. Huisman, A. Free, C. Picioreanu, C. Quince, I. Klapper, S. Labarthe, B. F. Smets, H. Wang, I. N. I. Fellows, and O. S. Soyer, *Challenges in microbial ecology: building predictive understanding of community function and dynamics*, *The ISME Journal* **10** (2016) 2557.

- [102] J. A. Papin, J. L. Reed, and B. O. Palsson, *Hierarchical thinking in network biology: the unbiased modularization of biochemical networks*, *Trends in Biochemical Sciences* **29** (2004), no. 12 641 – 647.
- [103] D. Merenstein, N. El-Nachef, and S. V. Lynch, *Fecal microbial therapy: Promises and pitfalls*, *Journal of Pediatric Gastroenterology and Nutrition* **59** (2014), no. 2.
- [104] V. Bucci, B. Tzen, N. Li, M. Simmons, T. Tanoue, E. Bogart, L. Deng, V. Yeliseyev, M. L. Delaney, Q. Liu, B. Olle, R. R. Stein, K. Honda, L. Bry, and G. K. Gerber, *Mdsine: Microbial dynamical systems inference engine for microbiome time-series analyses*, *Genome Biology* **17** (2016), no. 1 121.
- [105] Z. Hou, B. Lisena, M. Pireddu, F. Zanolin, S. Ahmad, and I. Stamova, *Lotka-Volterra and Related Systems: Recent Developments in Population Dynamics*. De Gruyter Series in Mathematics and Life Sciences. De Gruyter, 2013.
- [106] S. Wiggins, *Introduction to applied nonlinear dynamical systems and chaos*, vol. 2. Springer Science & Business Media, 2003.
- [107] S. S. Hota and S. M. Poutanen, *Is a single fecal microbiota transplant a promising treatment for recurrent Clostridium difficile infection?*, *Open Forum Infect Dis* **5** (2018), no. 3.
- [108] C. Gu, *Model order reduction of nonlinear dynamical systems*. PhD thesis, UC Berkeley, 2011.
- [109] A. Goeke, S. Walcher, and E. Zerz, *Classical quasi-steady state reduction—a mathematical characterization*, *Physica D: Nonlinear Phenomena* **345** (2017) 11 – 26.
- [110] P. I. Costea, F. Hildebrand, M. Arumugam, F. Bäckhed, M. J. Blaser, F. D. Bushman, W. M. de Vos, S. Ehrlich, C. M. Fraser, M. Hattori, C. Huttenhower, I. B. Jeffery, D. Knights, J. D. Lewis, R. E. Ley, H. Ochman, P. W. O’Toole, C. Quince, D. A. Relman, F. Shanahan, S. Sunagawa, J. Wang, G. M. Weinstock, G. D. Wu, G. Zeller, L. Zhao, J. Raes, R. Knight, and P. Bork, *Enterotypes in the landscape of gut microbial community composition*, *Nature Microbiology* **3** (2018), no. 1 8–16.
- [111] V. B. Young, *The role of the microbiome in human health and disease: an introduction for clinicians*, *BMJ* **356** (2017).
- [112] J. E. Belizário and M. Napolitano, *Human microbiomes and their roles in dysbiosis, common diseases, and novel therapeutic approaches*, *Frontiers in Microbiology* **6** (2015) 1050.

- [113] T. J. Borody and A. Khoruts, *Fecal microbiota transplantation and emerging applications*, *Nature Reviews Gastroenterology & Hepatology* **9** (2011) 88 EP –.
- [114] R. D. Heath, C. Cockerell, R. Mankoo, J. A. Ibdah, and V. Tahan, *Fecal microbiota transplantation and its potential therapeutic uses in gastrointestinal disorders*, *North Clin Istanb* **5** (2018), no. 1 79–88.
- [115] L. E. Hudson, S. E. Anderson, A. H. Corbett, and T. J. Lamb, *Gleaning insights from fecal microbiota transplantation and probiotic studies for the rational design of combination microbial therapies*, *Clinical Microbiology Reviews* **30** (2017), no. 1 191–231.
- [116] J. Bilinski, P. Grzesiowski, N. Sorensen, K. Madry, J. Muszynski, K. Robak, M. Wroblewska, T. Dzieciatkowski, G. Dulny, J. Dwilewicz-Trojaczek, W. Wiktor-Jedrzejcak, and G. W. Basak, *Fecal microbiota transplantation in patients with blood disorders inhibits gut colonization with antibiotic-resistant bacteria: Results of a prospective, single-center study*, *Clinical Infectious Diseases* **65** (2017), no. 3 364–370.
- [117] Y. Taur, K. Coyte, J. Schluter, E. Robilotti, C. Figueroa, M. Gjonbalaj, E. R. Littmann, L. Ling, L. Miller, Y. Gyaltshen, E. Fontana, S. Morjaria, B. Gyurkocza, M.-A. Perales, H. Castro-Malaspina, R. Tamari, D. Ponce, G. Koehne, J. Barker, A. Jakubowski, E. Papadopoulos, P. Dahi, C. Sauter, B. Shaffer, J. W. Young, J. Peled, R. C. Meagher, R. R. Jenq, M. R. M. van den Brink, S. A. Giralt, E. G. Pamer, and J. B. Xavier, *Reconstitution of the gut microbiota of antibiotic-treated patients by autologous fecal microbiota transplant*, *Science Translational Medicine* **10** (2018), no. 460.
- [118] D. J. Messenheimer, S. M. Jensen, M. E. Afentoulis, K. W. Wegmann, Z. Feng, D. J. Friedman, M. J. Gough, W. J. Urba, and B. A. Fox, *Timing of pd-1 blockade is critical to effective combination immunotherapy with anti-ox40*, *Clinical Cancer Research* **23** (2017), no. 20 6165–6177.
- [119] S. Wang, *Optimal sequential immunization can focus antibody responses against diversity loss and distraction*, *PLoS Computational Biology* **13** (2017), no. 1 1–27.
- [120] D.-W. Kang, J. B. Adams, A. C. Gregory, T. Borody, L. Chittick, A. Fasano, A. Khoruts, E. Geis, J. Maldonado, S. McDonough-Means, E. L. Pollard, S. Roux, M. J. Sadowsky, K. S. Lipson, M. B. Sullivan, J. G. Caporaso, and R. Krajmalnik-Brown, *Microbiota transfer therapy alters gut ecosystem and improves gastrointestinal and autism symptoms: an open-label study*, *Microbiome* **5** (2017), no. 1.
- [121] S. Huang, I. Ernberg, and S. Kauffman, *Cancer attractors: a systems view of tumors from a gene network dynamics and developmental perspective*, in *Seminars in Cell & Developmental Biology*, vol. 20, pp. 869–876, Elsevier, 2009.

- [122] L.-Z. Wang, R.-Q. Su, Z.-G. Huang, X. Wang, W.-X. Wang, C. Grebogi, and Y.-C. Lai, *A geometrical approach to control and controllability of nonlinear dynamical networks*, *Nature Communications* **7** (2016), no. 1 11323.
- [123] A. B. Shreiner, J. Y. Kao, and V. B. Young, *The gut microbiome in health and in disease*, *Current Opinion in Gastroenterology* **31** (2015), no. 1 69.
- [124] Y. K. Lee and S. K. Mazmanian, *Has the microbiota played a critical role in the evolution of the adaptive immune system?*, *Science* **330** (2010), no. 6012 1768–1773.
- [125] D. E. Voth and J. D. Ballard, *Clostridium difficile toxins: mechanism of action and role in disease*, *Clinical Microbiology Reviews* **18** (2005), no. 2 247–263.
- [126] L. V. McFarland, G. W. Elmer, and C. M. Surawicz, *Breaking the cycle: treatment strategies for 163 cases of recurrent Clostridium difficile disease*, *The American Journal of Gastroenterology* **97** (2002), no. 7 1769–1775.
- [127] L. J. Brandt, O. C. Aroniadis, M. Mellow, A. Kanatzar, C. Kelly, T. Park, N. Stollman, F. Rohlke, and C. Surawicz, *Long-term follow-up of colonoscopic fecal microbiota transplant for recurrent Clostridium difficile infection*, *American Journal of Gastroenterology* **107** (2012), no. 7 1079–1087.
- [128] P. Spigaglia, F. Barbanti, P. Mastrantonio, E. S. G. on Clostridium difficile (ESGCD), G. Ackermann, C. Balmelli, F. Barbut, E. Bouza, J. Brazier, M. Delmée, *et. al.*, *Multidrug resistance in European Clostridium difficile clinical isolates*, *Journal of Antimicrobial Chemotherapy* **66** (2011), no. 10 2227–2234.
- [129] A. Khoruts and M. J. Sadowsky, *Understanding the mechanisms of faecal microbiota transplantation*, *Nature Reviews Gastroenterology & Hepatology* **13** (2016), no. 9 508.
- [130] H. E. Vuong and E. Y. Hsiao, *Emerging roles for the gut microbiome in autism spectrum disorder*, *Biological Psychiatry* **81** (2017), no. 5 411–423.
- [131] R. J. Perry, L. Peng, N. A. Barry, G. W. Cline, D. Zhang, R. L. Cardone, K. F. Petersen, R. G. Kibbey, A. L. Goodman, and G. I. Shulman, *Acetate mediates a microbiome–brain– β -cell axis to promote metabolic syndrome*, *Nature* **534** (2016), no. 7606 213–217.
- [132] D.-W. Kang, J. B. Adams, D. M. Coleman, E. L. Pollard, J. Maldonado, S. McDonough-Means, J. G. Caporaso, and R. Krajmalnik-Brown, *Long-term benefit of microbiota transfer therapy on autism symptoms and gut microbiota*, *Scientific Reports* **9** (2019), no. 1 5821.

- [133] D. Xu, V. L. Chen, C. A. Steiner, J. A. Berinstein, S. Eswaran, A. K. Waljee, P. D. Higgins, and C. Owyang, *Efficacy of fecal microbiota transplantation in irritable bowel syndrome: a systematic review and meta-analysis*, *American Journal of Gastroenterology* **114** (2019), no. 7 1043–1050.
- [134] M. B. Biggs, G. L. Medlock, G. L. Kolling, and J. A. Papin, *Metabolic network modeling of microbial communities*, *Wiley Interdisciplinary Reviews: Systems Biology and Medicine* **7** (2015), no. 5 317–334.
- [135] D. S. Wishart, Y. D. Feunang, A. C. Guo, E. J. Lo, A. Marcu, J. R. Grant, T. Sajed, D. Johnson, C. Li, Z. Sayeeda, *et. al.*, *Drugbank 5.0: a major update to the drugbank database for 2018*, *Nucleic Acids Research* **46** (2018), no. D1 D1074–D1082.
- [136] J. Mounier, C. Monnet, T. Vallaey, R. Arditi, A.-S. Sarthou, A. Hélias, and F. Irlinger, *Microbial interactions within a cheese microbial community*, *Appl. Environ. Microbiol.* **74** (2008), no. 1 172–181.
- [137] M. G. de Vos, M. Zagorski, A. McNally, and T. Bollenbach, *Interaction networks, ecological stability, and collective antibiotic tolerance in polymicrobial infections*, *Proceedings of the National Academy of Sciences* **114** (2017), no. 40 10666–10671.
- [138] J. Friedman, L. M. Higgins, and J. Gore, *Community structure follows simple assembly rules in microbial microcosms*, *Nature Ecology & Evolution* **1** (2017), no. 5 1–7.
- [139] L. Li and Z. S. Ma, *Testing the neutral theory of biodiversity with human microbiome datasets*, *Scientific Reports* **6** (2016) 31448.
- [140] K. Faust, F. Bauchinger, B. Laroche, S. De Buyl, L. Lahti, A. D. Washburne, D. Gonze, and S. Widder, *Signatures of ecological processes in microbial community time series*, *Microbiome* **6** (2018), no. 1 1–13.
- [141] W. T. Sloan, M. Lunn, S. Woodcock, I. M. Head, S. Nee, and T. P. Curtis, *Quantifying the roles of immigration and chance in shaping prokaryote community structure*, *Environmental Microbiology* **8** (2006), no. 4 732–740.
- [142] D. M. Anderson, *Approaches to monitoring, control and management of harmful algal blooms (HABs)*, *Ocean & Coastal Management* **52** (2009), no. 7 342 – 347.
- [143] R. E. Gozlan, S. St-Hilaire, S. W. Feist, P. Martin, and M. L. Kent, *Disease threat to European fish*, *Nature* **435** (2005), no. 7045 1046–1046.
- [144] R. E. Thresher, K. Hayes, N. J. Bax, J. Teem, T. J. Benfey, and F. Gould, *Genetic control of invasive fish: technological options and its role in integrated pest management*, *Biological Invasions* **16** (2014), no. 6 1201–1216.

- [145] H. E. Jakobsson, C. Jernberg, A. F. Andersson, M. Sjölund-Karlsson, J. K. Jansson, and L. Engstrand, *Short-term antibiotic treatment has differing long-term impacts on the human throat and gut microbiome*, *PLoS ONE* **5** (2010), no. 3 1–12.
- [146] R. E. Thresher, *Autocidal technology for the control of invasive fish*, *Fisheries* **33** (2008), no. 3 114–121.
- [147] P. Van den Abbeele, W. Verstraete, S. El Aidy, A. Geirnaert, and T. Van de Wiele, *Prebiotics, faecal transplants and microbial network units to stimulate biodiversity of the human gut microbiome*, *Microbial Biotechnology* **6** (2013), no. 4 335–340.
- [148] P. J. Taylor, *Consistent scaling and parameter choice for linear and generalized Lotka-Volterra models used in community ecology*, *Journal of Theoretical Biology* **135** (1988), no. 4 543 – 568.
- [149] O. Malcai, O. Biham, P. Richmond, and S. Solomon, *Theoretical analysis and simulations of the generalized Lotka-Volterra model*, *Phys. Rev. E* **66** (2002) 031102.
- [150] L. L. Eberhardt, *Applying difference equations to wolf predation*, *Canadian Journal of Zoology* **76** (1998), no. 2 380–386.
- [151] H. Fort, *On predicting species yields in multispecies communities: Quantifying the accuracy of the linear Lotka-Volterra generalized model*, *Ecological Modelling* **387** (2018) 154–162.
- [152] O. S. Venturelli, A. V. Carr, G. Fisher, R. H. Hsu, R. Lau, B. P. Bowen, S. Hromada, T. Northen, and A. P. Arkin, *Deciphering microbial interactions in synthetic human gut microbiome communities*, *Molecular Systems Biology* **14** (2018), no. 6.
- [153] M. Fan, K. Wang, and D. Jiang, *Existence and global attractivity of positive periodic solutions of periodic n -species Lotka-Volterra competition systems with several deviating arguments*, *Mathematical Biosciences* **160** (1999), no. 1 47 – 61.
- [154] J. A. Vano, J. C. Wildenberg, M. B. Anderson, J. K. Noel, and J. C. Sprott, *Chaos in low-dimensional Lotka–Volterra models of competition*, *Nonlinearity* **19** (2006), no. 10 2391–2404.
- [155] R. M. May, *Will a large complex system be stable?*, *Nature* **238** (1972), no. 5364 413–414.
- [156] S. Allesina and S. Tang, *Stability criteria for complex ecosystems*, *Nature* **483** (2012), no. 7388 205–208.

- [157] T. Gibbs, J. Grilli, T. Rogers, and S. Allesina, *Effect of population abundances on the stability of large random ecosystems*, *Phys. Rev. E* **98** (2018) 022410.
- [158] T. S. Rayner and R. G. Creese, *A review of rotenone use for the control of non-indigenous fish in Australian fresh waters, and an attempted eradication of the noxious fish, *Phalloceros caudimaculatus**, *New Zealand Journal of Marine and Freshwater Research* **40** (2006), no. 3 477–486.
- [159] T. Scotti, M. Mimura, and J. Y. Wakano, *Avoiding toxic prey may promote harmful algal blooms*, *Ecological Complexity* **21** (2015) 157 – 165.
- [160] S. J. Whipple, J. S. Link, L. P. Garrison, and M. J. Fogarty, *Models of predation and fishing mortality in aquatic ecosystems*, *Fish and Fisheries* **1** (2000), no. 1 22–40.
- [161] A. L. Hartman, D. M. Lough, D. K. Barupal, O. Fiehn, T. Fishbein, M. Zasloff, and J. A. Eisen, *Human gut microbiome adopts an alternative state following small bowel transplantation*, *Proceedings of the National Academy of Sciences* **106** (2009), no. 40 17187–17192.
- [162] C. De Filippo, D. Cavalieri, M. Di Paola, M. Ramazzotti, J. B. Poullet, S. Massart, S. Collini, G. Pieraccini, and P. Lionetti, *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** (2010), no. 33 14691.
- [163] S. Jalili-Firoozinezhad, F. S. Gazzaniga, E. L. Calamari, D. M. Camacho, C. W. Fadel, A. Bein, B. Swenor, B. Nestor, M. J. Cronce, A. Tovaglieri, O. Levy, K. E. Gregory, D. T. Breault, J. M. S. Cabral, D. L. Kasper, R. Novak, and D. E. Ingber, *A complex human gut microbiome cultured in an anaerobic intestine-on-a-chip*, *Nature Biomedical Engineering* **3** (2019), no. 7 520–531.
- [164] K.-H. Lin, B.-Y. Liao, H.-W. Chang, S.-W. Huang, T.-Y. Chang, C.-Y. Yang, Y.-B. Wang, Y.-T. K. Lin, Y.-W. Wu, S.-L. Tang, and H.-T. Yu, *Metabolic characteristics of dominant microbes and key rare species from an acidic hot spring in Taiwan revealed by metagenomics*, *BMC Genomics* **16** (2015), no. 1 1029.
- [165] *Structure, function and diversity of the healthy human microbiome*, *Nature* **486** (2012), no. 7402 207–214.
- [166] L. Fan, D. Reynolds, M. Liu, M. Stark, S. Kjelleberg, N. S. Webster, and T. Thomas, *Functional equivalence and evolutionary convergence in complex communities of microbial sponge symbionts*, *Proceedings of the National Academy of Sciences* **109** (2012), no. 27 E1878–E1887.

- [167] S. Azaele, S. Suweis, J. Grilli, I. Volkov, J. R. Banavar, and A. Maritan, *Statistical mechanics of ecological systems: Neutral theory and beyond*, *Rev. Mod. Phys.* **88** (2016) 035003.
- [168] M. Advani, G. Bunin, and P. Mehta, *Statistical physics of community ecology: a cavity solution to MacArthur's consumer resource model*, *Journal of Statistical Mechanics: Theory and Experiment* **2018** (2018), no. 3 033406.
- [169] J. Bertram and R. C. Dewar, *Combining mechanism and drift in community ecology: a novel statistical mechanics approach*, *Theoretical Ecology* **8** (2015) 419–435.
- [170] J. N. E. Lucero, A. Mehdizadeh, and D. A. Sivak, *Optimal control of rotary motors*, *Phys. Rev. E* **99** (2019) 012119.