

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Eco-evolutionary dynamics in high dimensions

Permalink

<https://escholarship.org/uc/item/8t30776w>

Author

Martis, Stephen

Publication Date

2021

Peer reviewed|Thesis/dissertation

Eco-evolutionary dynamics in high dimensions

by

Stephen Martis

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Oskar Hallatschek, Chair

Professor Daniel Rokhsar

Professor Yun S. Song

Summer 2021

Eco-evolutionary dynamics in high dimensions

Copyright 2021
by
Stephen Martis

Abstract

Eco-evolutionary dynamics in high dimensions

by

Stephen Martis

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Oskar Hallatschek, Chair

With the advent of many new and revolutionary technologies (e.g. cheap high-throughput sequencing, precise gene editing, and more), biology has become a much more quantitative science over the past few decades. With these technologies in hand, we find ourselves able to probe some of the “fundamental” questions and assumptions that undergird the fields of evolution and ecology, both in controlled laboratory environments and in natural settings. However, with these new, quantitative observations, it is clear that some of the frameworks we use to think about biological populations are insufficient to describe relatively simple scenarios.

One of the key assumptions shared by population genetics and theoretical ecology is that these two fields are distinct. In other words, it has been taken for granted that ecological and evolutionary processes act separately and on disparate timescales. However, this may not necessarily be the case, where even in controlled laboratory evolution experiments, ecological structure frequently *evolves* on human observable times. This leads to an interesting (and broad) set of questions – what are the new timescales to consider in a joint eco-evolutionary process? How does eco-evolutionary feedback affect known observables? What new observables might be relevant? And, importantly, what should we find surprising in such a setting?

The contents of this dissertation hope to *start* to answer some of these questions by proposing and analyzing relatively simple models of eco-evolutionary dynamics. Since the task involves combining models that fall under the distinct classes of population genetics and ecology, the resulting joint models are necessarily more complex. However, by taking cues from statistical physics, I study these models in an explicitly high-dimensional setting, finding some forms of simplification.

First, inspired by experimental observations of diversification resulting from the evolution of novel resource preferences, I (in joint work with Benjamin Good and Oskar Hallatschek)

propose a minimal model of evolution in the setting of resource consumption with trade-offs. This model combines aspects of niche construction theory from the realm of ecology, with directional selection from the realm of population genetics. We study the low and high dimensional behavior of the model and describe its relatively simple steady state behavior which is dominated by resource generalists.

Second, I extend this model to include epistasis, or ‘rugged’ trade-offs. I show that the simple behavior of the non-epistatic model yields to a richer phase diagram when there is even weak epistasis. I show that in the many resource limit, the resource generalist state becomes ‘fragile’ to small epistatic fitness differences. This results in a transition in which the steady state gives way to a state of ‘punctuated equilibrium’ in which the ecosystem spends long times waiting for fitness mutations which bring about rapid rearrangement of the resource strategies of resident strains. This can be understood in light of the form of the Lyapunov function, which naturally separates into fitness specific and ecology specific components.

Finally, I propose a simple model of predator-prey co-evolution in a high dimensional setting. Using a combination of stochastic and deterministic simulations and theory inspired by the physics of disordered spin systems, I show that co-evolution stabilizes such populations for sufficiently variable interactions and for sufficiently high mutation rates, which stands in sharp contrast to expectations from ecological models alone. I also derive the phase boundary between a stable eco-evolutionary phase and an extinct phase, showing the dependence on relevant parameter combinations.

To my family and my friends.

Contents

Contents	ii
1 Introduction	1
1.1 Set and setting	1
1.2 Observations and experiments	3
1.3 Ecology and mean-field population dynamics	8
1.4 Evolutionary dynamics and population genetics	16
1.5 Eco-evolutionary dynamics	21
1.6 The structure of this dissertation	23
1.7 Bibliography	25
2 Adaptation limits ecological diversification and promotes ecological tinkering during the competition for substitutable resources	31
2.1 Abstract	32
2.2 Introduction	32
2.3 Evolutionary model of resource competition	34
2.4 Analysis	37
2.5 Discussion	49
2.6 Supporting Information (SI)	52
2.7 Bibliography	89
3 Quenched disorder and directional selection lead to punctuated equilibrium in the competition for substitutable resources	94
3.1 Abstract	95
3.2 Introduction	95
3.3 Consumer resource model with quenched disorder	97
3.4 Evolution in the binary resource mutational architecture	98
3.5 Results	100
3.6 Discussion	108
3.7 Supporting Information (SI)	111
3.8 Bibliography	128

4	Eco-evolutionary feedback can stabilize multi-strain predator-prey communities	131
4.1	Abstract	132
4.2	Main	132
4.3	Model description	133
4.4	Population stability with strong selection, strong mutation	136
4.5	Total population sizes are stable in the many strain limit	139
4.6	Discussion	140
4.7	Materials and Methods	141
4.8	Supporting Information (SI)	143
4.9	Bibliography	169
5	Conclusions	172
5.1	Bibliography	174

Acknowledgments

‘It takes a village...’ as the saying goes.

I feel fortunate to have joined Oskar Hallatschek’s group, and am especially grateful for Oskar’s wisdom in allowing me seemingly unlimited freedom to explore and to find my niche. It was a long and winding road but I think I got there! I wish Oskar the best as he moves to his new position in Leipzig and hope that our paths may cross again in the future.

One of the great joys of being a member of Team Hallatschek has been the fellowship and camaraderie with the other team members, past and present. I’d especially like to thank Benjamin Good who has been my closest collaborator and whose attitude and ingenuity has been a constant source of inspiration. I’d also like to thank Aditya Prasad, Carl Schreck, Daniel Weissmann, Diana Fusco, Jayson Paulose, Joao Ascensao, Jona Kayser, Jonas Denk, Marie-Cecilia Duvernoy, Matti Gralka, Morgan Delarue, Qinqin Yu, Pawel Gniewek, Takashi Okada, Wolfram Möbius, and Yuya Karita. Thank you for all the long lunches, elaborate coffee drinks, beers, trips to Tahoe, and countless memories.

I’m also grateful to my committee members Yun Song and Dan Rokhsar. Yun’s theoretical immunology seminar was a landmark in my personal development and I’m excited that my next step is a postdoc in that very field! In addition, Yun has been a great moral support over the years. And I will never forget when I first met Dan, I rolled up to his office without warning and we ended up chatting for much longer than I’ve come to expect from an unannounced visit to a tenured professor. Since then, I’ve found our conversations invaluable, be they discussions pertaining to research or incidental chats about the history of physics over Facebook Messenger.

I would be remiss to fail to mention some of the professors I’ve had the pleasure of GSling for, especially Ian Holmes and Mike Zaletel. They are both great examples of what the teaching part of being a professor ought to be. In the process of teaching with them, I learned as much about computational biology and statistical physics as from any course I’ve taken as a student. In addition to the professors, I’d like to thank the students I’ve taught over the years, across all the courses I’ve taught – I’ve learned so much from you all!

And then there are all the friends that I’ve encountered along the way, those I met in grad school and those I met at house shows, or at poetry readings, or at the myriad other strange places I seem to have found myself over the years. There are too many to name individually but you know who you are and you know how much you mean to me.

Finally, I’d like to thank my parents, who have been with me every step of the way. After these prodigal years in California, I’m coming home!

Chapter 1

Introduction

Dongh was worried by these multiple-choice futures, but Lyubov enjoyed them. In diversity is life and where there's life there's hope, was the general sum of his creed, a modest one to be sure.

— Ursula K. LeGuin, *The Word for World is Forest* [43]

1.1 Set and setting

The amount of biological diversity that exists and persists at many scales of observation is remarkable. Even in the well-manicured neighborhoods of Berkeley, California, many different types of macroscopic and microscopic diversity are readily observed. For instance, one might encounter many different species of tree standing side-by-side in a yard. Looking up at a particular tree, one might observe several species of bird sitting in its branches. A particularly entrepreneurial reader might go on to examine a soil sample from the base of the tree under a microscope and see hundreds or even thousands of types of microbes that coexist in that environment. It is a wonder that all of these living things, from the trees to the birds to the microbes, have descended from a common ancestor billions of years ago and yet persist together. A natural (and everlasting) question, then, is “how might this all have come to be?”

Dobzhansky’s famous dictum that “nothing makes sense in biology except in the light of evolution” [19] has served as a loose organizing principle (really a sort of repeated, if imprecise, mantra) for modern thinking about this question. In recent years, by coupling mathematical models of evolutionary dynamics with advances in sequencing technology, we have started to take both this question and this principle more “seriously.” Specifically, many researchers have begun to engage with them *quantitatively*, attacking more precise questions about the “tempo and mode” [66] of evolution – both in controlled laboratory settings and through observations of natural populations.

Many attempts to model the process of evolution have relied on certain simplifying assumptions, specifically that we can roll all of an evolving system’s biological complexity into a small set of effective parameters, including, perhaps most importantly, a scalar quantity that has been referred to as ‘fitness’ [33, 11]. Fitness can mean many things and can be measured differently depending on the specific context or even who you ask [44, 1, 62]. At

its most essential, fitness is a phenotype associated with a genotype. This phenotype is often defined as the genotype's growth rate when compared to the growth rate of the genotypes that comprise its surrounding population. Fitness, so defined, can take on positive, negative or null values, referring to genotypes that are growing, declining or indistinguishable from average within a population. In general, a genotype's fitness can vary with time, but in the simplest case, we can assume that it is constant for all time.¹

In this constant scalar fitness framework, mutations can result in fitness changes, which in turn accrue within a population over time. We can then picture evolution as an incremental process by which mutation seeds types with increasingly higher fitness, marching towards higher and higher values along an abstract fitness axis. Indeed, there has been much work to construct explicit models of this process, which has borne out some clear predictions [17] and even some correspondence with carefully controlled laboratory experiments [6]. However, with an influx of quantitative data, there is increasing evidence that these simple mathematical models of evolution, while incredibly rich and fruitful, are insufficient to describe even relatively simple laboratory experiments [30, 34, 12], let alone what we observe in complex natural settings [29, 55].

Returning to Dobzhansky's metaphor, we can understand that the "light of evolution" illuminates myriad other processes that occur within and between populations: competition, altruism, predation and many more complex scenarios. In addition, it is intuitively clear that individuals adapt to the abiotic characteristics of their environments, consuming resources and developing mechanisms to avoid toxins. These additional processes and externalities, which can loosely be described as *ecology*, in turn strongly influence what can and will evolve. So even in the simplest scenarios, the evolutionary process and a genotype's 'fitness' can be extremely contingent on the environmental or ecological setting in which it proceeds.

So in some sense we ought to refine the metaphor. Evolution is still the light that guides us, but ecology is the battery that powers the light, and as such ought to directly inform any attempt to understand the evolutionary process. But even this isn't the complete picture. Oftentimes, the timescales of environmental and evolutionary change are comparable (and even directly coupled). When this is the case, the more apt description is that evolution and ecology form a closed loop (Fig. 1.1), a sort-of perpetual motion machine in which light and battery power each other – or put more simply, the metaphor breaks down. How then should we assess our simple models of adaptation along a single dimension?

We can go all the way back to Darwin, who acknowledged such a potentially complicated relationship between ecology and evolution in his original treatise "On the Origin of Species," where he alluded to an "entangled bank" of species evolving together in a shared and co-created environment [13]. So instead of starting from a reductionist view of evolution as an incremental process along a one-dimensional fitness axis, we might instead embrace the messy complexity of a joint eco-evolutionary process. Taking cues from statistical physics, we

¹Needless to say, this sort of vast simplification has accrued years of criticism and debate, which I shall not discuss here. Instead we will hold up scalar fitness models as a sort of 'straw man' which, while undoubtedly useful, might be worth reconsidering. Parties interested in critical discussions of scalar fitness might see [4] and many more.




Figure 1.1: **An eco-evolutionary metaphor.** Evolution illuminates biological observation and experiment, but ecology ‘powers’ the evolutionary process. In this sense ecology and evolution form a closed feedback loop, that cannot necessarily be disentangled.

can look at *large*, complex eco-evolutionary systems with the hope that there is simplification when there are sufficiently many constituent parts (e.g. species, resources, genotypes, etc). Such *high-dimensional* models are the core contents of this dissertation.

The rest of this introduction will lay the groundwork for the coming chapters by providing an aerial view of some of the established (and most illustrative) low- and high-dimensional evolutionary and ecological frameworks that have received attention over the years. I will also identify some of the difficulties in combining these processes in the most general case and will discuss some of the extant models of joint eco-evolutionary dynamics. But first, I will discuss the relevant experimental and observational work that informs the different modeling choices and motivates why we even need models in the first place.

1.2 Observations and experiments

Relevant scales

Before we get into any specifics, it’s instructive to get a feel for the numbers and scales that we typically deal with when we think about evolutionary and ecological processes. Depending on the particular system and question of interest, the relevant numbers can vary greatly, but I will do my best to call attention to when this may be relevant to our estimates.

There is a vast range of timescales that we must consider if we go back to the very beginning of life on planet Earth. Current fossil evidence suggests that microbial life was present on earth as early as 4.28 billion years ago in the purported remnants of hydrothermal vents [20]. This is not so long after the Earth itself formed approximately 4.5 billion years ago.

In natural communities like the gut microbiome, the time it takes for a microbe to reproduce is approximately one day or $\sim 10^5$ s [40]. Taking this to be the mean reproduction time for all microbes, over all time, we can estimate (conservatively) the number of microbial generations since the earliest known:

$$\frac{10^{17} \text{ s}}{10^5 \text{ s/generation}} \sim 10^{12} \text{ generations}$$

Another important scale to get a handle on is the typical size of the communities we care about. The population size can serve as a rough indicator of the relevance of noise in a system, with large populations typically being less prone to stochastic effects. Sticking with the human-relevant example of the human microbiome as a “typical” microbial community, there are $10^2 - 10^3$ or more unique species in a given human microbiome, inhabiting different compartments at a broad range of abundances. Among all these species, there are some estimates that there are roughly 10^{14} microbial cells in and on a human body (with the majority residing in the gut) [65]. This is roughly as many human cells as there are in an average human, which is quite remarkable (and perhaps even shocking). If this population size is roughly constant, then over the course of a single day, almost all of these microbial cells will have been replaced by their descendants. As such, this will have allowed for many mutations to have entered this metapopulation.

We can get a handle on the mutation rate by looking to experimental and observational data. Mutation rates can vary greatly across species, but we will stick specifically to microbes. A reasonable estimate of the per base mutation rate of a bacterium is 10^{-10} mutations/nucleotide/reproduction event [7]. Of course, this neglects other types of mutational events, including rearrangements, horizontal gene transfer, deletions, and insertions, but it serves as a good base case to begin with. We expect that a typical microbe has roughly 10^6 nucleotides in its genome [53], this gives us 10^{-4} single site mutations per reproduction event. Plugging in our population size estimate for a single human microbiome, this amounts to 10^{10} single site mutations *per day per microbiome* (consistent with the lower end of the estimate in [75]). If we assume that this single microbiome has existed since the dawn of life, we have that there have been 10^{22} single site mutations over the course of its entire ~ 4 billion year history. This may seem like an extremely large number, but we need to compare it to the size of ‘genome space’ in order to get a true sense of scale.

Given that a microbe has about 10^6 nucleotides in its genome, that means that there are roughly 10^{10^6} possible genomes of that length. For our single microbiome example, if every mutation resulted in a new, never before seen genome, there would have only ever been 10^{22} unique genomes over its history. Given estimates that the global bacterial population size is roughly 10^{31} , that makes 10^{43} bacteria that have ever lived (assuming a fixed population size). This would have resulted in the (vast overestimate) of 10^{39} unique genomes. This is orders of magnitude greater than the number of stars in the observed universe. However, this is *tiny* when compared to the size of total genome space. In fact, it corresponds to completely exploring the space of a single stretch of DNA of less than 100 nucleotides in length. This is much shorter than the length of even a typical gene.

So now we can try to synthesize some of this information. Sticking specifically to microbes, we can see that despite rapid adaptation, the evolutionary search of genome space is in no sense even close to exhaustive. Furthermore, we can also see that since many microbes can coexist in shared environments (e.g. the human gut), there ought to be some sort of coupling between the biotic and abiotic surroundings of a species and its evolutionary future. Additionally, at least empirically, despite all of these interactions and rapid changes, natural ecosystems seem to be remarkably stable on human relevant timescales and beyond. The very large numbers at play (and their relative scales), the evolutionary coupling of chance to contingency, and the observed qualitative stability are all important to keep in mind when considering modeling choices and some of the results presented in future chapters.

Relevant observations

Equipped with a sense of scale, we can discuss some relevant observational and experimental work which lends credence to an underlying eco-evolutionary process. The observational examples have a large impact on human health and global environmental health. The evolution experiment that I discuss is relevant here because it was specifically designed to avoid ecological structure, which repeatedly turned up by evolutionary means.

Human gut microbiome

The human gut microbiome is a fascinating system that plays a crucial role in human health. As mentioned in the previous section, an extremely large number of species at fairly large abundances live within our large intestine. The gut microbiome has been implicated in many pathologies in mouse and human models, from ones localized to the gut like inflammatory bowel disease [16], to obesity [46], to pathologies that affect brain health and function [56, 37]. As such, it is clearly an important system to understand.

As of now, there is very little in the way of a unifying framework to understand gut microbiome composition and function. However, with the relative ease of modern sequencing technology, we are accruing a collection of observations that can begin to inform and constrain theoretical models. Among the important observations is that the gut microbiome seems to be resilient (or stable, at the species and even strain level), both over long times [21] and in the face of relatively large perturbations, including antibiotic treatment [61]. In addition, there is mounting evidence that adaptation occurs on human relevant timescales within the gut [75, 24], which is consistent with our discussion in the previous section.

This presents an interesting theoretical challenge then – if the gut is typically so stable (at coarse taxonomic scales), but is also adapting at the genetic level, how can we reconcile these two phenomena? One might expect that adaptation would greatly change the composition of microbial types present, at least if we were to rely on a fitness-based description of the evolutionary process. This leaves us with a few possibilities: mutations occurring in gut microbes are effectively neutral, some type of spatial or niche structure prevents sub populations from directly competing with each other for shared resources, or something more

complicated might be happening, like monopolization effects [14]. In reality, all of these and more might be contributing factors to the observed data, but both more data needs to be analyzed and more careful models need to be proposed to constrain the scenarios which might be relevant to the human gut.

Host-pathogen co-evolution

A prime and timely example of eco-evolutionary feedback is that of the co-evolutionary dynamics of pathogens and their hosts. I would be remiss not to mention the COVID-19 pandemic, which is ongoing at the time of writing of this dissertation. The qualitative dynamics of this global pathogen can be described as ‘co-evolutionary’, as the global population develops immunity and new variants begin to evade this immunity. However, since this is an ongoing crisis, I cannot precisely comment on how this fits within the context of this dissertation. It will remain to be seen whether eco-evolutionary thinking can bring much to bear on the understanding and outcome of the pandemic.

However, one extremely well-studied and relevant example of host-pathogen co-evolution is the adaptation of influenza to human immunity. Influenza is a viral pathogen that infects about 1 billion people every year, globally [57]. Since it persists over sufficiently long timescales (many years), it has measurably evolved and rapidly at that. It evolves so rapidly in fact, that a new vaccine is tailored to the specific strain that is dominant in a given year [25]. There is a whole cottage industry of scientists and organizations which contribute to this process, sequencing strains globally and developing methods to better predict which strain will be the next to rise to high frequency [51].

In addition to constant evasion of human immune recognition in the short term, influenza has branched into different co-circulating lineages on longer (but still observable) time scales [63]. This sort of quasi-speciation can be attributed to many things including allopatry via multiple jumps into human hosts [38], the dynamical process of immune escape [74] and canalization in high-dimensional ‘shape spaces’ [8]. Regardless, this sort of coexistence phenomenon has an inherently ‘ecological’ character which is driven by evolutionary change. Understanding the exact mechanism by which multiple strains evolve and persist is crucial for mitigating the endemic spread of this disease.

Prochlorococcus diversity

Recent progress in sequencing has allowed for the detailed characterization of genetic diversity in non-human related natural populations. One such population is that of the ocean-dwelling cyanobacterium *Prochlorococcus*, a so-called ‘global species’ [36]. *Prochlorococcus* is the most abundant photosynthetic organism on the planet, with an estimated population size of 10^{27} cells spread out across the world’s oceans [9]. Despite its characterization as a single ‘species,’ there is substantial variation within this taxonomic class, with many different clades or *ecotypes* having been observed in different samples.

With the help of single-cell whole-genome sequencing of many individual cells, it was found that even within ecotypes there are still finer-scaled clusters of diversity [36]. These ‘genomic backbones’ are estimated to have diverged potentially millions of years ago (corresponding to many millions of generations). One possible ecological explanation for this genealogical structure is ‘ancient, stable niche partitioning’ [36]. However this presents an interesting theoretical problem since there has clearly been sufficient time for substantial adaptation of these diverged lineages, which would naively reduce any diversity.

Lenski’s longterm evolution experiment (LTEE)

An extremely simple and invaluable laboratory experiment that serves as an important point of reference for evolutionary and ecological theory is Richard Lenski’s long-term evolution experiment (LTEE) [45]. The essential goal of the LTEE was to simply watch evolution proceed in a sufficiently simple and well-controlled laboratory setting. This was carried out by taking a well-characterized laboratory strain of *Escherichia coli* and propagating it for ~ 30 years (or $\sim 60,000$ generations) in a minimal, fixed environment [48]. The null expectation was that over time, the *E. coli* strain would acquire mutations so that evolved strains would have improved fitness in the environment in comparison to their ancestors. And this was in fact observed [71], though the complete picture is much more complex.

After several years of ‘simple’ directional adaptation, it was observed that in one population, two different phenotypic types emerged and were maintained at intermediate frequencies [64]. The types were observed to have a colony-level phenotype, and are referred to as ‘morphs.’ The first morph, the so-called L morph, exhibited relatively large colonies when plated on rich media. The second morph, the S morph, exhibited small colonies on plates. Importantly, these phenotypic characteristics are heritable, so one could infer that the phenotypic changes were caused by changes at the genetic level. See Fig. 1.2 for an image of these morphs.

By counting colonies, it was initially determined that these types had emerged and fluctuated at intermediate frequencies for many thousands of generations [64]. Later, using molecular and sequencing techniques, it was discovered that the S morph had acquired the ability to process acetate, a byproduct of the L (as well as the ancestor) type’s metabolism, which primarily used glucose as a carbon source. The molecular changes that necessitated this phenotypic change were later teased apart [59]. It was shown that only a small set of mutations were sufficient for the S morph to evolve, allowing for its relatively rapid emergence.

However, this is only a single population out of twelve initial populations in the LTEE. Perhaps this observed diversification was the result of chance? Luckily, part of the experimental design of the LTEE was such that samples of populations were frozen and stored away every 500 generations, thus preserving a ‘fossil record’ of the evolutionary process. A recent study went back and examined this fossil record, sequencing many of the time points to observe the dynamics of adaptation at the genomic level [30]. Remarkably, it was discovered that many populations consisted of multiple haplotypes co-existing at intermediate frequen-




Figure 1.2: **S and L morphs from the LTEE.** The S and L morphs plated on maltose media. Only the S morph can utilize maltose which allows for the differential coloration of the two types. On more standard plates, the size differential between colonies would be more readily observed, although it is still noticeable here. Photo c/o Joao Ascensao.

cies. So apparently in *E. coli* grown in minimal media, diversification is not necessarily a rare occurrence. Indeed diversification has also been observed in an independent experiment in *E. coli* [22]. Moreover, it was observed that over time there was continued evolution on the diversified lineages, leading to perturbations in lineage frequencies, while overall the rate of adaptation slowed. Despite the simple nature of this model system, it would seem that the timescales of ecology and evolution are comparable even in this context.

1.3 Ecology and mean-field population dynamics

Now that we have some sense of what is out there in nature and in laboratories, we can start to think about the right ingredients in modeling the processes that underlie these systems. First, we discuss models of ecology. A core facet of mathematical ecology is population dynamics. Population dynamics is the study of how populations grow, fluctuate and interact, primarily by means of difference equations and differential equations. This is obviously an extremely broad field that would take several textbooks [49, 52], lecture notes [5, 58] and Youtube videos [32] to provide a satisfying overview. I will not attempt to do this here. Instead, I will focus on some of the relevant background for the remaining

chapters of the dissertation. The main thrust of this section will be to describe ‘mean-field’ population dynamics models, which are commonly employed deterministic models of population growth. Any biological population is subject to stochastic forces. However, if the population size is large enough, it can often be assumed (with impunity) that stochastic effects are subleading and that discrete population variables can be treated as continuous quantities. To be pedagogically complete, I will illustrate how these deterministic models can connect to microscopic stochastic dynamics, since, as I will discuss in the next section, stochasticity cannot be ignored in evolutionary dynamics.

The simplest models: exponential and logistic growth of single species

In the simplest case, a fixed environment with limitless resources and no interactions with other species or abiotic elements, each identical individual in a population will grow at a fixed rate, s . This can be captured by the following mean-field differential equation model for the number of individuals n :

$$\frac{d}{dt}n = sn \Rightarrow n(t) = n(0)e^{st}$$


This is an incredibly simple, but also incredibly rich model that is relevant to population genetics (as we will discuss in the next section).

However, resources are never limitless and so exponential growth cannot proceed indefinitely (i.e. it is reasonable to assume that there is some sort of population size regulation). One common way to account for this is via the phenomenological logistic growth model:

$$\frac{d}{dt}n = sn \left(1 - \frac{n}{K}\right)$$

in which K is the carrying capacity or total population size. In this model a population at low abundance ($n \ll K$) will grow exponentially at rate s until nonlinear effects kick in and the population saturates at K individuals. Carrying capacity or population size control is the simplest form of ecological interaction in a population dynamic model. To see what I mean precisely, we need to look at the underlying *stochastic reaction process*.

If we use the variable X to denote an exchangeable individual in the population, we can define two types of interactions with accompanying rates:



If we go through the process of writing down a master equation for the dynamics of the probability distribution of the population size, $P(n, t)$ (the probability that there are n individuals of type X present at time t), we recover the logistic equation when we take the mean-field limit [23]. A large class of theoretical models from population genetics have the

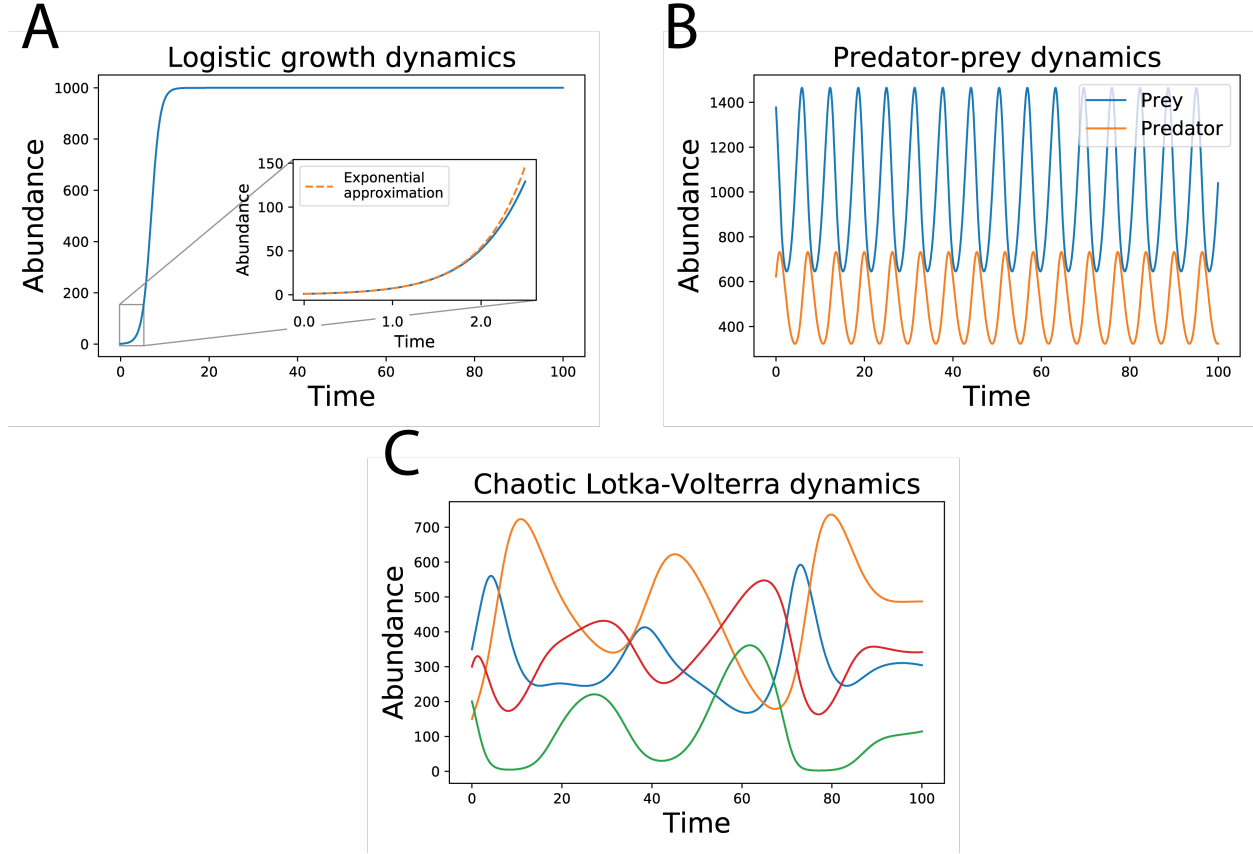



Figure 1.3: **Different examples of mean-field population dynamics.** Panel **A** shows simple logistic growth. For short times, the logistic function can be approximated by exponential growth, but eventually nonlinear effects kick in and cause the population abundance to saturate at some final equilibrium value. Panel **B** shows predator-prey dynamics for a single predator and a single prey type. These dynamics follow a limit cycle in the mean-field limit. However, this limit cycle is marginally stable so that stochastic fluctuations will drive rapid extinctions. Panel **C** shows chaotic (competitive) Lotka-Volterra dynamics of 4 species. Notice that individual population abundances can get close to zero so that such chaotic systems can be especially susceptible to stochastic extinctions. The parameters used were found and described in [70].

logistic model as a mean-field limit albeit with a different underlying reaction process. We will discuss this in a later section of the introduction.


In many cases in ecology and in population genetics, we are specifically interested in models with multiple types, i.e. when there are sets of individuals which are not exchangeable. In considering such cases, we need to have a sufficiently reasonable description of what constitutes a type. We start with a version of the logistic reaction scheme above, except

we have N types that interact. In the most general case, a single type i is described by up to $2N + 1$ parameters, s and K_{ij} (as well as K_{ji}). This means that we can have up to $N(2N + 1) \sim N^2$ rate parameters describing one population. In proceeding, we will start from the level of stochastic interactions to get a handle on the microscopic details of the model and then assert the relevant mean-field equation and how it depends on these details. However, before we discuss the general case, we will start from a simplified case with maximally symmetric interactions where $K_{ij} = K_{ji} = K$.

The first process we can consider is that each type possesses its own intrinsic growth rate, s_i , which gives us the set of reactions:



Now since the population size control is implemented by an interaction at the microscopic level we can make several choices for its generalization to the multi-type case. On the one hand, we can assume that all types have the same interaction parameter, K , so that we have a set of interactions:




where this ‘death reaction’ is proportional to the growth rate of the dying type. This results in the mean field equations for the population size of the types:

$$\frac{d}{dt}n_i = s_i n_i \left(1 - \frac{\sum_j n_j}{K}\right)$$

In this model, it can be demonstrated that the type with the highest growth rate will comprise the entire population at long times, which is known as the *competitive exclusion principle*.

But why should interactions take this specific form? Even if we limit ourselves to interactions that are a simple product of the abundances of the reactants, we are left with an embarrassment of options. We might have the case where interactions are proportional to the growth rate of the surviving type:




which results in a different set of mean field equations with *qualitatively* different behaviors:

$$\frac{d}{dt}n_i = n_i \left(s_i - \frac{\sum_j s_j n_j}{K}\right)$$

These equations have been extensively studied under the umbrellas of population genetics, ecology, physics and game theory. They are often referred to as the *replicator equations* and serve as a minimal model of evolutionary games.

In the most extreme case, we can have a unique rate for every interaction:



so that there are N growth rate parameters and N^2 interaction parameters. This yields the mean-field equations:

$$\frac{d}{dt}n_i = s_i n_i - \sum_j a_{ij} n_i n_j$$

These equations are referred to as *competitive Lotka-Volterra equations*, and these have also been studied in many contexts. Interestingly, it was proved that for systems with 5 or more species, depending on the choice of parameters, any dynamics could be expected: stability, limit cycles, or chaos [67]. We will discuss this model and its extensions to non-competitive cases shortly.

From this simple example we can start to recognize a couple of things. First of all, there are many particular choices that we can make in describing the population dynamics of a system. What the ‘relevant’ choices are and whether or not there are useful forms of universality among the space of models is an important and as yet unsolved problem. The other key observation is that even in this relatively simple model, there is an explosion in the number of parameters that rapidly grows with the order of the interactions. This is the first suggestion that understanding these sorts of population models *statistically* might be a fruitful strategy in understanding their typical behavior. This is a line of inquiry many others have pursued and which is continued in this dissertation.

Abiotic interactions: MacArthur’s consumer-resource model

Now that we have discussed the case where we have multiple ‘types’ growing with and without population size control, we can start to include more details, i.e. more types and specific mechanisms of population size control. One of the classic models of size control via external factors is MacArthur’s consumer-resource model. I go on to analyze this model in following chapters so it bears utility to discuss it in some detail here.

The consumer-resource model partitions a system into two classes, unsurprisingly referred to as consumers and resources. The consumers require resources in order to grow, otherwise they are removed (via death or some other mechanism like the outflow of a chemostat). The resources can typically be biotic or abiotic, but we will only consider the latter case here. The biotic case bears some relation to predator-prey models which we will discuss in the next section and in chapter 4 of the dissertation.

We will denote abundance of consumers by the variables C_μ and the abundance of resources by the variables R_i . In this case, instead of starting from individual-based interactions, we’ll jump directly to the mean-field and justify this post-hoc. The dynamics of the consumers take the form:

$$\frac{d}{dt}C_\mu = \left[g_\mu(\vec{R}) - \delta_\mu \right] C_\mu$$

where $g_\mu(\vec{R})$ defines the growth rate of consumer μ as a function of the set of resource abundances and δ is a death or outflow rate. We have used vectorial short hand in defining

the growth rate, but this should not be taken to indicate any special properties of the function g_μ .

In the most general case, abiotic resources will obey the equations:

$$\frac{d}{dt}R_i = s_i - h_i(\vec{C}, R_i)$$

Such abiotic resources are supplied at a fixed rate s_i and are depleted at a rate h_i , which is a function of the consumer abundances and the abundance of resource i . More terms might be added, for instance an outflow rate (perhaps modeling a chemostat environment), but typically these can be incorporated into the generic depletion rate h_i . Other potential complications include scenarios in which the supply rates are functions of the consumer abundances (for instance in the case of cross-feeding type behavior), or interactions between the resources. However, we do not consider cases like this in this dissertation although they present an interesting avenue for study and are almost certainly relevant in natural and experimental settings.

In the generic form that we have cast them, not much can be said about the properties of solutions to these equations other than that there is a fixed point when:

$$g_\mu(\vec{R}^*) = \delta_\mu \quad \text{and} \quad s_i = h_i(\vec{C}^*, R_i^*)$$

We need to more precisely specify the model in order to determine whether or not this fixed point is feasible and stable. There are several assumptions that have commonly been used to make progress along these lines. The original consumer-resource model implementation [47] (which has been reanalyzed in the high-dimensional limit in [2]) made the following relatively simple assumptions for the growth rate of consumers and consumption rate of the resources:

$$g_\mu(\vec{R}) = \sum_i a_{\mu i} w_i R_i \quad \text{and} \quad h_i(\vec{C}, R_i) = \sum_\mu a_{\mu i} C_\mu R_i$$

We can eliminate the resource abundances to get an equilibrium condition for the type abundances:

$$\sum_i \frac{a_{\mu i} w_i s_i}{\sum_\mu a_{\mu i} C_\mu^*} = \delta_\mu$$

Since there is an equilibrium (which can proven to be unique, see chapters 2 and 3), any fluctuations will be around this equilibrium. We can in turn show that fluctuations are subleading in the limit of large population sizes, so that our mean field assumption is justified [29].


In the model where the death rate is constant across types, $\delta_\mu = \delta$, as might be the case in a chemostat, and in which the resource strategies \vec{a}_μ are constrained by linear trade-offs (e.g. $\sum_i a_{\mu i} = \text{const.} \quad \forall \mu$), the model has the remarkable feature that an assembled ecosystem can easily break the bounds set by competitive exclusion [60]. What this means is that at equilibrium, the number of species at finite abundance can be greater than (in fact arbitrarily greater than) the number of resources. However, upon introducing variance

in the death rate δ_μ , this symmetry is broken unless very particular conditions are met and the number of species is bounded by the number of resources [68]. In chapters 2 and 3, I show that evolutionary dynamics can break this sort of fragile excess coexistence. However, in chapter 3, I show that a type of non-equilibrium high diversity state can be achieved with a more complex trade-off constraint surface.


Biotic interactions: the Lotka-Volterra framework

Interactions need not necessarily be mediated by externally supplied resources. Types might interact with each other through competitive mechanisms (which we have already discussed), as well as through predation, cooperation, and potentially even higher order interactions. This is not surprising as we typically think of natural ecosystems comprised of a food web of different consumers and different prey. In our discussion of logistic-like population dynamic models, we came up with a set of equations describing competition for a single resource dimension.

These *competitive Lotka-Volterra equations* can be generalized to incorporate other types of interactions. For instance predation:



and cooperation:



The so called *generalized Lotka-Volterra equations* take on a form that is at least visually similar to the competitive case:

$$\frac{d}{dt}n_i = \left(s_i + \sum_j a_{ij}n_j \right) n_i$$

However, now the matrix of interactions $A = a_{ij}$ can take on entries with arbitrary sign and has no strict a priori structure. The matrix A is referred to as the community matrix and its properties will be analyzed shortly.

A special case: predator-prey interactions

Predator-prey interactions are a classic subclass of this model that have been studied in great detail. In the simplest case, we might have a single prey and a single predator strain, both of which are assumed to be far from their carrying capacities. We can assume that without sufficient prey, the predator would die off and that without predators, the prey would grow exponentially. This leads to the coupled set of equations for the prey population, X , and

the predator population, Y :

$$\begin{aligned}\frac{d}{dt}X &= (b - aY)X \\ \frac{d}{dt}Y &= (-d + cX)Y\end{aligned}$$

We can readily show that there is the trivial fixed point, $(X^*, Y^*) = (0, 0)$, as well as a non-trivial fixed point:

$$Y^* = b/a, \quad X^* = d/c$$

We can study the stability properties of these fixed points. The trivial fixed point is unstable, which can be shown rather simply. For the nontrivial fixed point we obtain the following linearized equations for small perturbations δX and δY :

$$\begin{aligned}\frac{d}{dt}\delta X &= -aX^*\delta Y \\ \frac{d}{dt}\delta Y &= cY^*\delta X\end{aligned}$$

We can compute the eigenvalues:

$$\lambda_{\pm} = \pm i\sqrt{bd}$$

Since the eigenvalues are pure imaginary, we can infer that there is a closed orbit around the fixed point. These orbits are determined by the initial conditions of the system and are given by isoclines of the conserved quantity:

$$V \equiv cX - d \ln X + aY - c \ln Y$$

Many species: May's stability result

An extremely important result in the history of theoretical ecology is Robert May's sweeping yet simple comment on the stability of generic high-dimensional Lotka-Volterra models [50]. May's argument starts with the community matrix of a generalized Lotka-Volterra model (or really the Jacobian matrix of any arbitrary high-dimensional differential equation model), which determines the stability of the assumed fixed point of the equations. In other words, for the system under consideration, perturbations around a fixed point \vec{n}^* will be described by:

$$\frac{d}{dt}\delta\vec{n} = A\delta\vec{n}$$

One can then go on to assume that the diagonal of the matrix has values of -1 , which indicates carrying capacity constraints on the individual types. This gives us:

$$A = B - I_n$$

where I_n indicates the n -dimensional identity and B is a matrix with random entries. Therefore, the eigenvalues of A , which determine stability, are given by $\lambda_B - 1$ where λ_B is an eigenvalue of B .

The entries of B are drawn from a distribution with mean 0 and variance σ^2 . Importantly, these sorts of matrices had been studied in the context of large interacting quantum systems and so the statistical properties of their eigenvalues had been worked out by Wigner, Ginibre, Mehta and others. The distribution of eigenvalues for a random, not necessarily symmetric matrix of the sort with which we are concerned is given by a ‘circular law’ (in contrast to the semi-circular law of Wigner, which applies to random symmetric matrices). The eigenvalues of the matrix are uniformly distributed within a circle of radius $\sigma\sqrt{n}$ centered at $(0, 0)$ on the complex plane. The condition for stability is that the largest real part of the eigenvalues is less than zero. This translates to:

$$\sigma\sqrt{n} < 1$$

which is the result reported in May and others [50, 3].

We can learn a lot from this exceedingly brief and simple ‘proof’ (which is only slightly shorter than May’s original paper). First and foremost it is readily apparent that statistical thinking can be extremely useful when considering large ecological models and can lead to sweeping, generic results. Second, that we have our work cut out for us – it would seem all but the most particular ecological models will be unstable. Moreover it seems that we can expect any small heterogeneity to destabilize a sufficiently large ecosystem. We will return to this topic in chapter 4, where I show that for *evolving* systems, the situation can be quite different: high variance interactions can actually be a prerequisite for stability.

1.4 Evolutionary dynamics and population genetics

The evolutionary process acts on top of the population dynamics of a system. However, in the spirit of simplification, much of population genetics has treated situations where the population size is approximately fixed. In this section, I will present some of the important aspects of evolutionary dynamics and population genetics. Where possible, I will make explicit connections to population dynamic models from the previous section.

The basics: selection, mutation, and drift

There are three ‘fundamental’ forces at play in any evolutionary process: selection, mutation, and drift. The scare-quotes should be taken to indicate that fundamental is meant in its most qualitative sense. Of course, there are other important biological processes which complement these three basic forces, such as recombination and horizontal gene transfer. However, despite their clear relevance in many natural systems, we will not discuss them here since they are not necessarily relevant to the particular systems that I study in the following chapters.

Selection, broadly defined, is the process by which differences in genotypes and phenotypes in a population lead to changes in the frequency of those types. An important phenotypic quantity that tracks these changes is the so-called ‘fitness’ of an individual type. Fitness is an inherently difficult term to define precisely, since an individual’s fitness can be both population and environment dependent [39]. But in essence, fitness refers to the growth rate of an individual with respect to its surrounding population, e.g. with respect to the population mean. Types can be partitioned into three classes then: 1) positive fitness, or having higher fitness than the population mean 2) negative fitness, or declining with respect to the mean of the population 3) neutral, having fitness that is identical to the mean.

Mutation is the process by which genotypic and phenotypic diversity are generated within a population. There are many methods by which this can occur, but one of the most commonly discussed methods is the single-site mutational process. In this process, a single nucleotide position in the genome undergoes a change, which can have downstream effects on the protein that it encodes. Some mutations are nonsynonymous and change the actual amino acid sequence that is encoded. Since the genetic code has some redundancies, some mutations are synonymous, and do not change the encoded amino acid sequence. It is often assumed that synonymous mutations are ‘neutral’ and have no fitness effect, while nonsynonymous mutations can have substantial fitness effects. However, increasing evidence shows that this might not be the case for synonymous mutations [42], perhaps due to the physical process of translation that occurs in the ribosome [35].

Drift is the final ‘fundamental’ force of evolutionary dynamics, and in essence it encodes the statement that any population biology process will be noisy. Drift is the process by which stochasticity perturbs the frequencies of types within a population so that it is separate from selection. It can take many forms and can be due to many particular subprocesses, although in practice it is assumed to come as a result of the distribution of offspring numbers or sampling noise [27]. One of the key historical debates of the population genetics community is whether or not evolutionary dynamics are drift-dominated or selection-dominated [41]. In the first case, all genotypes and mutations are effectively neutral with respect to each other so that fluctuations drive the dynamics of types over relatively long timescales (which are typically proportional to the population size). In the second case, mutation seeds appreciable fitness differences, which leads to rapid changes in the composition of types.

The Wright-Fisher model and single locus dynamics

The canonical model that incorporates these processes is known as the Wright-Fisher model. The Wright-Fisher model assumes a population of constant size N , on top of which mutations and selection occur. Drift is encoded in the fact that the population is finite (though perhaps large). In the model, time is discrete and measured in ‘generations.’ In the current generation, let us assume that there are N individuals distributed among M distinct neutral genotypes, the ‘parents.’ The next generation of N ‘children’ is generated by taking a multinomial sample of the parents, with weights defined by each parent type’s frequency. In the case where there are fitness differences, these enter into the weights as well, with positive




Figure 1.4: **Wright-Fisher model and its scaling limit.** Illustration of the Wright-Fisher model and the limiting process when the population size becomes large. In this limit the time between generations is small enough to be approximately continuous, $\Delta t = 1/N$. In this scaling limit N formally goes to infinity, but $s \sim \mathcal{O}(1/N)$ so that the product $Ns \sim \mathcal{O}(1)$.

fitness types getting a boost according to their fitness coefficient, s . In addition, mutations occur in the process of creating the next generation with a probability μ per individual.

In order to build intuition, we can turn our attention to a simple case within the Wright-Fisher framework, known as the *single locus model*. Let us consider the dynamics of a single binary genetic site with two allelic states. WLOG we can denote one of these alleles as the ‘wild type’ allele and measure fitness with respect to this type. We assume that the population is clonal in the wild type until $t = 0$, when a single mutant individual is born, so that $f_{mut}(0) = 1/N$. Let us say that this type has fitness $1 + s$, compared to the wildtype fitness 1. In the next generation after its initial appearance, this type is sampled according to the binomial distribution:

$$P(\# \text{ mutants} = n) = \binom{N}{n} w^n (1 - w)^{N-n}$$

where the normalized weight is given by:

$$w = \frac{f_{mut} e^{s\tau}}{f_{mut} e^{s\tau} + (1 - f_{mut})}$$

where τ is the generation time. We notice that in the limit $s \rightarrow 0$, this reproduces what we would expect for a neutral mutation, namely $w = f_{mut}$.

Coarse-graining over generations assuming the product $s\tau$ is small, so that frequency changes over a generation are small, we obtain the following Langevin equation as the ap-

appropriate continuum limit for the selection-drift dynamics of the mutant frequency:

$$\frac{d}{dt}f = \underbrace{sf(1-f)}_{\text{selection}} + \underbrace{\sqrt{\frac{f(1-f)}{N}}\eta(t)}_{\text{drift}}$$

where we have dropped the subscript of the frequency (e.g. due to the population size constraint, for the single-locus case the frequency of the mutant uniquely defines the state of the system). The first term on the right hand side of the equation can be understood as the deterministic contribution of selection and is formally identical to the logistic model we defined in our discussion of ecological models. The second term is the contribution of genetic drift (i.e. binomial sampling noise).

This model and its variants have been studied in detail in many other contexts [28], so I will not reproduce that here, except for a key point: drift is a *singular perturbation* and can never completely be neglected. The fate of a mutation is intimately linked to its dynamics when rare, e.g. in a drift dominated boundary layer [27]. We can see this by appealing to a heuristic analysis of the relevant timescales on which drift and selection substantially perturb frequencies.

Let us say that we have a focal genotype in a population at frequency f , with selection coefficient s . From the deterministic component of our Langevin equation, we see that the change in frequency due to selection is given by:

$$\delta f \sim |s|f(1-f)\delta t$$

so that selection will change the frequency by a factor of $\delta f/f \sim \mathcal{O}(1)$ on the timescale

$$\tau_{\text{sel}} \sim 1/|s|$$

Similarly we have that the change in frequency due to drift is given by:

$$\delta f \sim \mathcal{N}\left(0, \sqrt{\frac{f(1-f)\delta t}{N}}\right)$$

The timescale on which drift will perturb the frequency by a factor of $\mathcal{O}(1)$ is:

$$\tau_{\text{drift}} \sim Nf(1-f)$$

We can then compare these timescales:

$$\frac{\tau_{\text{drift}}}{\tau_{\text{sel}}} = N|s|f(1-f)$$

This ratio allows us to separate the Langevin equation into a ‘bulk’ and two boundary layer regimes where selection and drift, respectively, dominate the dynamics:

$$\frac{d}{dt}f \approx \begin{cases} \sqrt{\frac{f}{N}}\eta(t), & \text{for } f \lesssim 1/N|s| \\ sf(1-f), & \text{for } f(1-f) \gg 1/N|s| \\ \sqrt{\frac{1-f}{N}}\eta(t), & \text{for } 1-f \lesssim 1/N|s| \end{cases}$$

Importantly for spontaneous (small effect) mutations, which inevitably start at frequency $1/N$, they will start in the first (stochastic) regime so that stochastic effects cannot necessarily be ignored when considering evolution.

Linked selection

In the previous section we described how independent mutations arise and fix. However, mutations do not typically fix independently. Instead, a new mutation will typically come up on a background that includes other older mutations. The fitness effects of these mutations will then combine in some way to give the fitness of the mutant. However, the map between a genotype and fitness is not necessarily ‘additive,’ as evidenced by experimental observations. In such a situation, in which we have *epistasis*, we can appeal to Sewall Wright’s classic conceptual model of a ‘fitness landscape’ [72]. When the genotype-phenotype map is epistatic, a population might follow a complex trajectory along a ‘rugged’ high-dimensional landscape, eventually reaching a local fitness peak. Such epistatic evolution can impact notions of predictability in evolution and can interact with genetic drift effects in non-intuitive ways [15].

Even when the situation is at its simplest, and the fitness landscape is additive, linked selection can be quite complex. This is no more apparent than when the mutation rate is high compared with the time it takes a mutation to fix in the population. In such a case, multiple mutant lineages will compete within a population at once. This is the so-called *clonal interference* regime of population genetics. Competition between lineages can slow any individual lineage from fixing as we can see from a simple heuristic analysis.

We can follow the analysis of Desai and Fisher for illustrative purposes [17]. Let’s say we have a population with total number of individuals N . We assume that there is a typical beneficial fitness effect, s , and a beneficial mutation rate U_b . From our analysis in the previous section, a single mutation will escape the drift dominated regime on timescale $\tau_{\text{est}} \sim 1/NU_b s$. When τ_{est} is sufficiently small in comparison to the typical fixation time of a mutant ($NU_b \gg \frac{1}{\log N s}$), multiple mutants compete in the population at one time. In this case it can be shown that the population fitness distribution takes on an approximately Gaussian shape in fitness space at long times.

The rightmost edge of this ‘solitary wave,’ referred to as the *nose*, consists of a small population of very fit types that will make up the bulk of the population at some later time. Since it consists of relatively few individuals, the nose will be highly subject to stochastic drift effects until it increases in size and is pushed sufficiently far into the bulk of the wave packet. This process of the nose being pushed into the bulk must exactly balance the process of seeding new highly fit types for the wave’s shape to be stable, so that we can derive a typical speed of adaptation. Importantly, observables from such models drastically differ from expectations from the simple Wright-Fisher picture described above [54, 28].

1.5 Eco-evolutionary dynamics

Now that we have illustrated some of the basic characteristics of both ecological and evolutionary processes, and have hopefully built some intuition for the relevant behaviors when they are treated separately, we can turn to the joint eco-evolutionary process. We have learned several key factors in the previous sections. First, we saw that ecological interactions can lead to complex behavior. Even when ecology results in a stable population size, near the fixed point interactions will *by definition* balance clonal growth rate terms, and so that the interactions cannot necessarily be ignored. In the evolutionary context, we have seen that stochasticity can play a crucial part in adaptation due to the existence of a ‘drift boundary,’ i.e. mutations are subject to strong fluctuations when they first arise. In addition, evolution dictates that linkage and lineage structure and the phenotype ‘landscape’ should play some part in a joint eco-evolutionary framework.

We can consider a ‘generic’ eco-evolutionary model, which accounts for all of these forces, deterministic and stochastic. The ‘ecological’ component will consist of two terms: 1) the growth of different genotypes 2) the interactions between different genotypes. The ‘evolutionary’ component will be comprised of two terms as well: 1) mutation between genotypes 2) drift or noise. To make things slightly more precise, we can write down a representative Langevin equation of these dynamics:

$$\frac{d}{dt}n(\vec{g}) = \underbrace{s(\vec{g})n(\vec{g}) + \sum_{\vec{g}'} A(\vec{g}, \vec{g}')n(\vec{g}')n(\vec{g})}_{\text{‘ecology’}} + \underbrace{\sum_{|\vec{g}-\vec{g}'|=1} [\mu_{\vec{g}' \rightarrow \vec{g}}(\vec{n}) - \mu_{\vec{g} \rightarrow \vec{g}'}(\vec{n})]}_{\text{‘evolution’}} + \underbrace{\sqrt{D(\vec{n})}\eta(t)}_{\text{drift}}$$

There are two key subtleties that I will call attention to. First, the mutation term can depend on the composition of the population \vec{n} because it depends on the specific birth process that underlies each genotype. This is discussed in a more concrete setting in chapter 4. Second, the drift term can be fairly complex, but can be derived by considering the microscopic interactions via the chemical Langevin equation, only given an understanding of the underlying interaction structure [23].

However, while we can write down a plausible generic model, it is exceedingly complex and a general solution is not readily found. In order to proceed, assumptions must be made about the form of different terms or at least about their corresponding timescales. There have been many attempts to formally and verbally articulate the expected behaviors of such restricted eco-evolutionary models (which assume a separation of timescales or which might apply to specific systems). I will go on to discuss some special cases that fall within the generic eco-evolutionary class presented above. Aside from adaptive dynamics in particular, the mathematical analysis of these cases are relatively new and underexplored.

Adaptive dynamics

The adaptive dynamics framework (also known as evolutionary invasion analysis) is a sweeping one that describes *generic* eco-evolutionary process in the ‘strong selection, weak mutation’ regime [26]. This means that mutations occur slowly, while selection acts rapidly so that there is no interference between the dynamics of mutations. In the language of the generic eco-evolutionary model from above, the ‘ecology’ terms act on timescales that are much faster than the ‘mutation’ terms. Furthermore, adaptive dynamics operates on abstract (often continuous and low-dimensional) *phenotype space*, with mutations moving along a phenotype axis in small increments.

In this regime, the main quantity that is relevant for a mutant is its *invasion fitness*, or the growth rate of a type when it is rare in a population (which is a function of its abstract phenotype). The invasion fitness is a function of the mutant phenotype, y , and the environmental conditions, E , which are set by the resident population phenotypes, \vec{x} :

$$s_{\text{inv}}(y, E_{\vec{x}})$$

Consider an initially clonal population and consider a rare mutant with invasion fitness s_{inv} . If the invasion fitness of the mutant type is positive, it will invade. If its invasion fitness is negative, it will not. However, invasion itself does not determine the final state of the population (i.e. just because a type invades does not mean it will fix in the population).

The final fate of the population can be determined by a ‘reciprocal invasion analysis.’ If the mutant can invade the initially clonal wild type population and the wild type can invade an initial clonal mutant population, then the two types will coexist at some intermediate frequency. This can be formalized by examining the form of the interaction kernel, specifically its second derivative. One can then formalize different points around which such ‘evolutionary branching’ will occur in the population. These are known as evolutionary singular states, where the second derivative of the invasion fitness vanishes. The great result of adaptive dynamics was to show that there are attracting evolutionary singular states around which branching is inevitable.

There is an extensive literature studying the different outcomes of different adaptive dynamics scenarios [18]. There have even been some attempts to connect the theory with experimental observations [22]. However, because of limitations of the assumptions and the extreme generality of the model, direct quantitative applicability seems to be an ongoing challenge.

The Red Queen hypothesis

The red queen hypothesis originated as a primarily ‘verbal’ model of co-evolution proposed by Van Valen [69]. The Red Queen hypothesis gets its name from Lewis Carroll’s classic book *Through the Looking-glass*. In the book [10], the titular Red Queen tells Alice:

Now, *here*, you see, it takes all the running you can do, to keep in the same place.

Van Valen’s proposition was that this might serve as a reasonable metaphor for the evolutionary process of species that stably exist in competition with each other. It can be assumed that coexisting types will adapt at relatively similar rates. As such the mean *relative* fitness of any particular competitive ecosystem ought not to change too much. The population is left effectively ‘running in place.’

As it is formulated verbally this is inconsistent, since fitness gains result in higher frequency, which results in the increased chance for further adaptation along a particular lineage rather than as an entire population. However, it was recently shown that a similar type of behavior arises from a mathematical model of host-pathogen co-evolution [74]. In this scenario a pathogen population, in order to evade host immunity must constantly evolve. In doing so it forms a stable distribution of ‘effective’ fitness, in analogy to traveling wave models of directional selection [17, 54], although the types that comprise the population constantly change.

The co-evolving Kill-the-Winner model

The kill-the-winner model is another primarily ‘verbal’ model that seeks to explain the maintenance of diversity in phage-bacteria systems. The gist of the model is that phage will be more likely to attack a bacterial population that is at high abundance rather than one at low abundance (since interactions scale with the product of abundances). Therefore, whatever bacterial strain is highly abundant at a given time will be driven to lower abundances by phage until another strain grows to sufficiently high abundance. At this point the phage will consume the newly abundant strain and the cycle will continue, so that on average all bacterial populations are maintained at intermediate abundances.

It was recently demonstrated that stochasticity can destroy this sort of coexistence since the phage are able to drive bacterial abundances low enough to be close to extinction. However, by adding a simple mutational kernel to both the bacterial and phage types, coexistence can be rescued [73]. Therefore, eco-evolutionary coupling seems to be crucial in this case, and a detailed exploration of the parameter combinations describing this coupling is discussed in Chapter 4.

1.6 The structure of this dissertation

The subsequent chapters will proceed as follows.

Chapter 2 is adapted from [31]. This work discusses a model of resource competition and how evolutionary dynamics impacts the types of communities that evolve. This model exploits properties of mean-field consumer-resource models, while respecting the ‘proper’ evolutionary dynamics (e.g. by respecting genetic drift). In this model, mutation is partitioned into ‘directional’ and ‘ecological’ components which arrive stochastically on top of the equilibrium population dynamics, mimicking observations from microbial evolution experi-

ments. Depending on their relative rates, we find that these two components balance each other, dictating the number of species that might coexist at long times.

Chapter 3 extends the evolving resource competition model by imposing that ecological mutations are accompanied by fitness costs and benefits, modeling epistasis. In this work we show that weak epistasis can destabilize the particular form of the diversification-selection balance found in Chapter 2. This results in a type of ‘phase’ behavior and a transition between a generalist dominated state and a phase in which the resident resource strategies continually change over long times. This transition can be understood in light of the form of the Lyapunov function of the consumer-resource model, which acts analogously to a free energy from statistical physics.

Chapter 4 discusses a fully stochastic model of rapid predator-prey co-evolution, that goes beyond the strong selection, weak mutation regime into the territory of clonal interference. In this work I show that despite the apparent complexity of the problem – chaotic population dynamics, stochastic evolutionary dynamics, clonal interference, and more – there is vast simplification when the number of strains is large. In this simplification there is an emergent feedback mechanism, which stabilizes the population dynamics. I propose the mechanism by which this stabilization occurs and show that highly variable interactions are counter-intuitively *more stable* in contrast to the classical intuition. I also characterize the phase diagram of the model and provide relevant parameter combinations that describe the transition to stability.

Finally, in Chapter 5, I conclude by speculating about the place of this work in the bigger picture. Here and throughout these chapters, I attempt to connect my results to concepts from statistical physics and established literature to lend to their interpretation.

1.7 Bibliography

- [1] M. Abrams. Measured, modeled, and causal conceptions of fitness. *Frontiers in Genetics*, 3:196, 2012.
- [2] 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(3):033406, 2018.
- [3] S. Allesina and S. Tang. The stability–complexity relationship at age 40: a random matrix perspective. *Population Ecology*, 57(1):63–75, 2015.
- [4] A. Ariew and R. C. Lewontin. The confusions of fitness. *British Journal for the Philosophy of Science*, 55(2), 2004.
- [5] M. Assaf and B. Meerson. Wkb theory of large deviations in stochastic populations. *Journal of Physics A: Mathematical and Theoretical*, 50(26):263001, 2017.
- [6] A. N. N. Ba, I. Cvijović, J. I. R. Echenique, K. R. Lawrence, A. Rego-Costa, X. Liu, S. F. Levy, and M. M. Desai. High-resolution lineage tracking reveals travelling wave of adaptation in laboratory yeast. *Nature*, 575(7783):494–499, 2019.
- [7] J. E. Barrick and R. E. Lenski. Genome dynamics during experimental evolution. *Nature Reviews Genetics*, 14(12):827–839, 2013.
- [8] T. Bedford, A. Rambaut, and M. Pascual. Canalization of the evolutionary trajectory of the human influenza virus. *BMC biology*, 10(1):1–12, 2012.
- [9] S. J. Biller, P. M. Berube, D. Lindell, and S. W. Chisholm. Prochlorococcus: the structure and function of collective diversity. *Nature Reviews Microbiology*, 13(1):13–27, 2015.
- [10] L. Carroll. *Through the looking glass and what Alice found there*. Penguin UK, 2010.
- [11] J. F. Crow, M. Kimura, et al. An introduction to population genetics theory. *An introduction to population genetics theory.*, 1970.
- [12] I. Cvijović, A. N. N. Ba, and M. M. Desai. Experimental studies of evolutionary dynamics in microbes. *Trends in Genetics*, 34(9):693–703, 2018.
- [13] C. Darwin. *On the origin of species: A facsimile of the first edition*. Harvard University Press, 1964.
- [14] L. De Meester, J. Vanoverbeke, L. J. Kilsdonk, and M. C. Urban. Evolving perspectives on monopolization and priority effects. *Trends in Ecology & Evolution*, 31(2):136–146, 2016.

- [15] J. A. G. De Visser and J. Krug. Empirical fitness landscapes and the predictability of evolution. *Nature Reviews Genetics*, 15(7):480–490, 2014.
- [16] A. K. DeGruttola, D. Low, A. Mizoguchi, and E. Mizoguchi. Current understanding of dysbiosis in disease in human and animal models. *Inflammatory bowel diseases*, 22(5):1137–1150, 2016.
- [17] M. M. Desai and D. S. Fisher. Beneficial mutation–selection balance and the effect of linkage on positive selection. *Genetics*, 176(3):1759–1798, 2007.
- [18] U. Dieckmann and M. Doebeli. On the origin of species by sympatric speciation. *Nature*, 400(6742):354–357, 1999.
- [19] T. Dobzhansky. Nothing in biology makes sense except in the light of evolution. *The american biology teacher*, 75(2):87–91, 2072.
- [20] M. S. Dodd, D. Papineau, T. Grenne, J. F. Slack, M. Rittner, F. Pirajno, J. O’Neil, and C. T. Little. Evidence for early life in earth’s oldest hydrothermal vent precipitates. *Nature*, 543(7643):60–64, 2017.
- [21] J. J. Faith, J. L. Guruge, M. Charbonneau, S. Subramanian, H. Seedorf, A. L. Goodman, J. C. Clemente, R. Knight, A. C. Heath, R. L. Leibel, et al. The long-term stability of the human gut microbiota. *Science*, 341(6141), 2013.
- [22] M. L. Friesen, G. Saxer, M. Travisano, and M. Doebeli. Experimental evidence for sympatric ecological diversification due to frequency-dependent competition in escherichia coli. *Evolution*, 58(2):245–260, 2004.
- [23] C. Gardiner. *Stochastic methods*, volume 4. Springer Berlin, 2009.
- [24] N. R. Garud, B. H. Good, O. Hallatschek, and K. S. Pollard. Evolutionary dynamics of bacteria in the gut microbiome within and across hosts. *PLoS biology*, 17(1):e3000102, 2019.
- [25] C. Gerdil. The annual production cycle for influenza vaccine. *Vaccine*, 21(16):1776–1779, 2003.
- [26] S. A. Geritz, G. Mesze, J. A. Metz, et al. Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary ecology*, 12(1):35–57, 1998.
- [27] J. H. Gillespie. *Population genetics: a concise guide*. JHU Press, 2004.
- [28] B. H. Good. *Molecular Evolution in Rapidly Evolving Populations*. PhD thesis, Harvard University, 2016.

- [29] B. H. Good and O. Hallatschek. Effective models and the search for quantitative principles in microbial evolution. *Current opinion in microbiology*, 45:203–212, 2018.
- [30] B. H. Good, M. J. McDonald, J. E. Barrick, R. E. Lenski, and M. M. Desai. The dynamics of molecular evolution over 60,000 generations. *Nature*, 551(7678):45–50, 2017.
- [31] B. H. Good, S. Martis, and O. Hallatschek. Adaptation limits ecological diversification and promotes ecological tinkering during the competition for substitutable resources. *Proceedings of the National Academy of Sciences*, 115(44):E10407–E10416, 2018.
- [32] J. Grilli. ‘ecology & evolution (qls-ee)’. [YouTube playlist]. URL <https://www.youtube.com/watch?v=cS22wJKr0kM&list=PLp0hSY2uBeP8-UMVA0X2R1hcd3770chSs&index=1>.
- [33] J. B. S. Haldane. A mathematical theory of natural and artificial selection. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 23, pages 363–372. Cambridge University Press, 1926.
- [34] E. R. Jerison and M. M. Desai. Genomic investigations of evolutionary dynamics and epistasis in microbial evolution experiments. *Current opinion in genetics & development*, 35:33–39, 2015.
- [35] M. Josuweit and J. Krug. The fitness landscapes of translation. *Physica A: Statistical Mechanics and its Applications*, 569:125768, 2021.
- [36] N. Kashtan, S. E. Roggensack, S. Rodrigue, J. W. Thompson, S. J. Biller, A. Coe, H. Ding, P. Marttinen, R. R. Malmstrom, R. Stocker, et al. Single-cell genomics reveals hundreds of coexisting subpopulations in wild *prochlorococcus*. *Science*, 344(6182):416–420, 2014.
- [37] S. Kim, H. Kim, Y. S. Yim, S. Ha, K. Atarashi, T. G. Tan, R. S. Longman, K. Honda, D. R. Littman, G. B. Choi, et al. Maternal gut bacteria promote neurodevelopmental abnormalities in mouse offspring. *Nature*, 549(7673):528–532, 2017.
- [38] A. Kitchen, L. A. Shackelton, and E. C. Holmes. Family level phylogenies reveal modes of macroevolution in rna viruses. *Proceedings of the National Academy of Sciences*, 108(1):238–243, 2011.
- [39] K.-I. Kojima. Is there a constant fitness value for a given genotype? no! *Evolution*, 25(2):281–285, 1971.
- [40] T. Korem, D. Zeevi, J. Suez, A. Weinberger, T. Avnit-Sagi, M. Pompan-Lotan, E. Matot, G. Jona, A. Harmelin, N. Cohen, et al. Growth dynamics of gut microbiota in health and disease inferred from single metagenomic samples. *Science*, 349(6252):1101–1106, 2015.

- [41] M. Kreitman. The neutral theory is dead. long live the neutral theory. *Bioessays*, 18(8):678–683, 1996.
- [42] E. Lebeuf-Taylor, N. McCloskey, S. F. Bailey, A. Hinz, and R. Kassen. The distribution of fitness effects among synonymous mutations in a gene under directional selection. *Elife*, 8:e45952, 2019.
- [43] U. K. LeGuin. *The Word for World is Forest*. Berkley Books, 1976.
- [44] R. Lenski. Comment on the use of selection rate versus relative fitness.
- [45] R. E. Lenski, M. R. Rose, S. C. Simpson, and S. C. Tadler. Long-term experimental evolution in escherichia coli. i. adaptation and divergence during 2,000 generations. *The American Naturalist*, 138(6):1315–1341, 1991.
- [46] R. E. Ley, P. J. Turnbaugh, S. Klein, and J. I. Gordon. Human gut microbes associated with obesity. *nature*, 444(7122):1022–1023, 2006.
- [47] R. MacArthur. Species packing and competitive equilibrium for many species. *Theoretical population biology*, 1(1):1–11, 1970.
- [48] R. Maddamsetti, R. E. Lenski, and J. E. Barrick. Adaptation, clonal interference, and frequency-dependent interactions in a long-term evolution experiment with escherichia coli. *Genetics*, 200(2):619–631, 2015.
- [49] R. May. Theoretical ecology. principles and applications. 1976.
- [50] R. M. May. Will a large complex system be stable? *Nature*, 238(5364):413–414, 1972.
- [51] D. H. Morris, K. M. Gostic, S. Pompei, T. Bedford, M. Łuksza, R. A. Neher, B. T. Grenfell, M. Lässig, and J. W. McCauley. Predictive modeling of influenza shows the promise of applied evolutionary biology. *Trends in microbiology*, 26(2):102–118, 2018.
- [52] J. D. Murray. *Mathematical Biology I. An Introduction*. Springer, 2002.
- [53] S. Nayfach and K. S. Pollard. Average genome size estimation improves comparative metagenomics and sheds light on the functional ecology of the human microbiome. *Genome biology*, 16(1):1–18, 2015.
- [54] R. A. Neher and O. Hallatschek. Genealogies of rapidly adapting populations. *Proceedings of the National Academy of Sciences*, 110(2):437–442, 2013.
- [55] R. A. Neher and A. M. Walczak. Progress and open problems in evolutionary dynamics. *arXiv preprint arXiv:1804.07720*, 2018.
- [56] K. Neufeld, N. Kang, J. Bienenstock, and J. A. Foster. Reduced anxiety-like behavior and central neurochemical change in germ-free mice. *Neurogastroenterology & Motility*, 23(3):255–e119, 2011.

- [57] W. H. Organization et al. Global influenza strategy 2019-2030. 2019.
- [58] G. Oster. Lectures in population dynamics. *Modern Modelling of Continuum Phenomena*, 16:149–190, 1977.
- [59] J. Plucain, T. Hindré, M. Le Gac, O. Tenaillon, S. Cruveiller, C. Médigue, N. Leiby, W. R. Harcombe, C. J. Marx, R. E. Lenski, et al. Epistasis and allele specificity in the emergence of a stable polymorphism in escherichia coli. *Science*, 343(6177):1366–1369, 2014.
- [60] A. Posfai, T. Taillefumier, and N. S. Wingreen. Metabolic trade-offs promote diversity in a model ecosystem. *Physical review letters*, 118(2):028103, 2017.
- [61] M. Roodgar, B. H. Good, N. R. Garud, S. Martis, M. Avula, W. Zhou, S. Lancaster, H. Lee, A. Babveyh, S. Nesamoney, et al. Longitudinal linked read sequencing reveals ecological and evolutionary responses of a human gut microbiome during antibiotic treatment. *bioRxiv*, 2019.
- [62] A. Rosenberg and F. Bouchard. Fitness. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2020 edition, 2020.
- [63] P. A. Rota, T. R. Wallis, M. W. Harmon, J. S. Rota, A. P. Kendal, and K. Nerome. Cocirculation of two distinct evolutionary lineages of influenza type b virus since 1983. *Virology*, 175(1):59–68, 1990.
- [64] D. E. Rozen and R. E. Lenski. Long-term experimental evolution in escherichia coli. viii. dynamics of a balanced polymorphism. *The American Naturalist*, 155(1):24–35, 2000.
- [65] R. Sender, S. Fuchs, and R. Milo. Revised estimates for the number of human and bacteria cells in the body. *PLoS biology*, 14(8):e1002533, 2016.
- [66] G. G. Simpson. *Tempo and mode in evolution*. Columbia University Press, 1944.
- [67] S. Smale. On the differential equations of species in competition. *Journal of Mathematical Biology*, 3(1):5–7, 1976.
- [68] M. Tikhonov and R. Monasson. Collective phase in resource competition in a highly diverse ecosystem. *Physical review letters*, 118(4):048103, 2017.
- [69] L. Van Valen. Molecular evolution as predicted by natural selection. *Journal of molecular evolution*, 3(2):89–101, 1974.
- [70] J. Vano, J. Wildenberg, M. Anderson, J. Noel, and J. Sprott. Chaos in low-dimensional lotka–volterra models of competition. *Nonlinearity*, 19(10):2391, 2006.

- [71] M. J. Wiser, N. Ribeck, and R. E. Lenski. Long-term dynamics of adaptation in asexual populations. *Science*, 342(6164):1364–1367, 2013.
- [72] S. Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proceedings of the VI International Congress of Genetics*, pages 356–366, 1932.
- [73] C. Xue and N. Goldenfeld. Coevolution maintains diversity in the stochastic “kill the winner” model. *Physical review letters*, 119(26):268101, 2017.
- [74] L. Yan, R. A. Neher, and B. I. Shraiman. Phylodynamic theory of persistence, extinction and speciation of rapidly adapting pathogens. *Elife*, 8:e44205, 2019.
- [75] S. Zhao, T. D. Lieberman, M. Poyet, K. M. Kauffman, S. M. Gibbons, M. Groussin, R. J. Xavier, and E. J. Alm. Adaptive evolution within gut microbiomes of healthy people. *Cell host & microbe*, 25(5):656–667, 2019.

Chapter 2

Adaptation limits ecological diversification and promotes ecological tinkering during the competition for substitutable resources

Preface

I start my discussion of eco-evolutionary models with work that was done in collaboration with Benjamin Good and Oskar Hallatschek. This work was partially inspired by a recent flurry of attention from the physics community directed at MacArthur’s consumer-resource model, specifically focused on ecological processes such as community assembly. However, very few works have addressed the topic of evolution in these types of community settings. Since the ecological dynamics are well-studied, the consumer-resource model seemed like a reasonable place from which to attack evolutionary questions. In addition, the observed diversification in Lenski’s long term evolution experiment, due in part to evolving resource preferences, provided an important empirical touchstone. At the point of writing this dissertation, the consumer-resource model is over 50 years old. Yet incredibly, it continues to bear insight.

There are two primary threads in this work. The first is low-dimensional eco-evolutionary dynamics for 2 resources. The second is a high-dimensional eco-evolutionary dynamics with relatively simple properties. What we find in this case is that the observed dynamics are completely analogous to the low-dimensional case. However, this analogy is fragile to modeling choices, which is explored in greater detail in the next chapter.

The remainder of this chapter was published in its entirety as:
 “Adaptation limits ecological diversification and promotes ecological tinkering during the competition for substitutable resources.” Benjamin H Good, Stephen Martis, Oskar Hallatschek. *PNAS* (2018).

2.1 Abstract

Microbial communities can evade competitive exclusion by diversifying into distinct ecological niches. This spontaneous diversification often occurs amid a backdrop of directional selection on other microbial traits, where competitive exclusion would normally apply. Yet despite their empirical relevance, little is known about how diversification and directional selection combine to determine the ecological and evolutionary dynamics within a community. To address this gap, we introduce a simple, empirically motivated model of eco-evolutionary feedback based on the competition for substitutable resources. Individuals acquire heritable mutations that alter resource uptake rates, either by shifting metabolic effort between resources or by increasing the overall growth rate. While these constitutively beneficial mutations are trivially favored to invade, we show that the accumulated fitness differences can dramatically influence the ecological structure and evolutionary dynamics that emerge within the community. Competition between ecological diversification and ongoing fitness evolution leads to a state of diversification-selection balance, in which the number of extant ecotypes can be pinned below the maximum capacity of the ecosystem, while the ecotype frequencies and genealogies are constantly in flux. Interestingly, we find that fitness differences generate emergent selection pressures to shift metabolic effort toward resources with lower effective competition, even in saturated ecosystems. We argue that similar dynamical features should emerge in a wide range of models with a mixture of directional and diversifying selection.

2.2 Introduction

Ecological diversification and competitive exclusion are opposing evolutionary forces. Conventional wisdom suggests that most new mutations are subject to competitive exclusion, while ecological diversification occurs only under highly specialized conditions [56]. Recent empirical evidence from microbial, plant, and animal populations has started to challenge this assumption, suggesting that the breakdown of competitive exclusion is a more common and malleable process than is often assumed [43, 48, 54]. Particularly striking examples have been observed in laboratory evolution experiments, in which primitive forms of ecology evolve from a single ancestor over years [53], months [30], and even days [52].

In the simplest cases, the population splits into a pair of lineages, or *ecotypes*, that stably coexist with each other due to frequency-dependent selection, leading to a breakdown of competitive exclusion [30, 53, 18, 17, 57]. The mechanism of coexistence can often be traced to differences in resource utilization, or to the accessibility of privileged spatial or temporal niches. Interestingly, these microbes rarely cease their evolution once ecological diversification has been achieved. Sequencing studies have shown that adaptive mutations continue to accumulate within each ecotype, even when population-wide fixations are rare [31, 63, 29]. This additional evolution can cause the ecological equilibrium to wander over longer timescales, as observed in the shifting population frequencies of the two ecotypes [53, 29]. In certain cases, these evolutionary perturbations can even drive one of the original

lineages to extinction, either through the outright elimination of the niche [17], or by the invasion of individuals that mutate from the opposing ecotype [63].

Pairwise coexistence is the simplest form of community structure, but similar dynamics have been observed in more complex communities as well. Some laboratory experiments diversify into three or more ecotypes [52, 13, 50], and it is likely that previously undetected ecotypes may be present in existing experiments [29]. Moreover, many natural microbial populations evolve in communities with tens or hundreds of ecotypes engaged in various degrees of competition and coexistence [62, 34, 16]. Although the evolutionary dynamics within these communities are less well-characterized, recent work suggests that similar short-term evolutionary processes can occur in these natural populations as well [5, 66, 20].

While the interactions between microbial adaptation and ecology are known to be important empirically, our theoretical understanding of this process remains limited in comparison. Early work in the field of adaptive dynamics [21] showed how ecological diversification emerges under very general models of frequency-dependent trait evolution, which are thought to describe the limiting behavior of a wide class of ecological interactions near the point of diversification. Numerous studies have also investigated the effects of evolution on ecological diversification and stability using computer simulations, in which the parameters of a particular ecological model are allowed to evolve over time [8, 9, 36, 1, 2, 55, 65]. Yet while both approaches can reproduce some of the qualitative behaviors observed in experiments, it has been difficult to forge a more quantitative connection between these models and the large amount of molecular data that is now available.

One reason that quantitative comparisons have been difficult is that evolution also selects for other traits that are not directly involved in diversification. For example, natural selection always works to maintain essential cellular functions, while there may be a benefit to removing costly functions that are not needed in the current environment. As a result, mutations that influence an ecologically-relevant phenotype like acetate metabolism might have to compete with constitutively beneficial mutations that are only tangentially related to metabolism [e.g. the loss of the yeast mating pathway [38]]. In some experiments, these constitutively beneficial mutations may even comprise the bulk of the mutations that reach observable frequencies [63, 39, 29]. Although many models exist for describing constitutively beneficial or deleterious mutations in the absence of ecology [47], we lack even a basic theoretical understanding of how they behave when they are linked to ecological phenotypes, and vice versa. The absence of quantitative theoretical predictions makes it difficult to draw any inferences from the vast molecular data that is now available.

To start to bridge this gap, we introduce a simple, empirically motivated model that describes the interplay between ecological diversification and directional selection at a large number of linked loci. The ecological interactions derive from a well-studied class of consumer resource models [61, 51, 59, 3], in which individuals compete for multiple substitutable resources (e.g. different carbon sources) in a well-mixed environment. We extend this ecological model to allow for heritable mutations in resource uptake rates, which can either divert metabolic effort between resources, or increase the growth rate on all resources. The latter class of mutations provides a natural way to model adaptation at linked genomic loci.

Constitutively beneficial mutations might seem like an ecologically trivial addition to the model, since they are always favored to invade on short timescales. On longer timescales, however, we will show that these accumulated fitness differences can dramatically influence both the ecological structure and the evolutionary dynamics that take place within the community. By focusing on the weak mutation limit, we derive analytical expressions for these dynamics in the two resource case, and we show how our results extend to larger communities as well. These analytical results provide a general framework for integrating ecological and population-genetic processes in evolving microbial communities, and suggest new ways in which these processes might be inferred from time-resolved molecular data.

2.3 Evolutionary model of resource competition

To investigate the interactions between ecological diversification and directional selection, we focus on a simple ecological model in which individuals compete for an assortment of externally supplied resources in a well-mixed, chemostat-like environment (Fig. 2.1). This resource-based model aims to capture some of the key ecological features observed in certain microbial evolution experiments [53, 18], as well as more complex ecosystems such as the gut microbiome [16], while remaining as analytically tractable as possible.

In our idealized setting, individuals compete for \mathcal{R} substitutable resources, which are supplied by the environment at a fixed rates (Fig. 2.1). Individuals are characterized by a resource utilization vector $\vec{r} = (r_1, \dots, r_{\mathcal{R}})$, which describes how well they can grow on each of the resources. We assume that the resource utilization phenotypes are constitutively expressed, so that we may neglect complicating factors like regulation. We will find it useful to decompose the phenotype \vec{r} into a normalized portion $\alpha_i = r_i / \sum_j r_j$, and an overall magnitude $X = \log(\sum_i r_i)$. The components of $\alpha_{\mu,i}$ represent the fractional effort devoted to growth on resource i , so we will refer to this quantity as the *resource strategy vector*. In contrast, the parameter X resembles environment-independent measure of *overall fitness*, an analogy that we will make more precise below.

We assume that individuals reproduce asexually, so that the state of the ecosystem can be described by the number of individuals n_{μ} with a given resource strategy vector $\vec{\alpha}_{\mu}$ and fitness X_{μ} . Under suitable assumptions, the ecosystem can be described by the stochastic differential equation,

$$\frac{\partial f_{\mu}}{\partial t} = \sum_{i=1}^{\mathcal{R}} \alpha_{\mu,i} \left[e^{X_{\mu} - \bar{X}_i(t)} - 1 \right] f_{\mu} + \frac{\xi_{\mu}(t)}{\sqrt{N}}, \quad (2.1)$$

where N is a fixed carrying capacity, $f_{\mu} = n_{\mu}/N$ is the relative frequency of strain μ , and $\xi_{\mu}(t)$ is a stochastic contribution arising from genetic drift (SI Appendix 2.6). The state of the environment is encoded by the set of resource-specific mean fitnesses,

$$\bar{X}_i(t) = \log \left(\frac{\sum_{\mu} \alpha_{\mu,i} e^{X_{\mu}} f_{\mu}}{\beta_i} \right), \quad (2.2)$$




Figure 2.1: **Ecological and evolutionary dynamics in a simplified consumer-resource model.** (a) Schematic depiction of ecological dynamics. Substitutable resources are supplied to the chemostat at constant rates β_i ($i = 1, \dots, \mathcal{R}$), measured in units of biomass ($\sum_i \beta_i = 1$). Cells import resources at genetically encoded rates, r_i , which define a normalized *resource strategy* $\alpha_i = r_i / \sum_j r_j$ and *overall fitness* $X = \log \sum_i r_i$. (b) Schematic depiction of evolutionary dynamics. Mutations that alter resource strategies ($\vec{\alpha}$) occur at rate U_{α} , while mutations that alter overall fitness (X) occur at rate U_X . (c-f) Simulated ecological and evolutionary dynamics, starting from a clonal ancestor, in an environment with $\mathcal{R} = 2$ resources. The four panels represent independent populations evolved under different sets of parameters, which differ only in the mutation rates and fitness benefits of pure fitness mutations (SI Appendix 2.6). Lines denote the population frequency trajectories of all mutations that reached frequency $\geq 10\%$ in at least one timepoint. Resource strategy mutations are shown in red, while pure fitness mutations are shown in blue.

where β_i denotes the fractional flux supplied by resource i . Eq. (2.1) is an example of a more general and well-studied class of consumer resource models introduced by Refs. [41, 61], whose ecological properties have been explored in several recent works [59, 51, 3].

However, the essential features of Eq. (2.1) are not limited to this consumer-resource framing. In SI Appendix 2.6, we argue that Eq. (2.1) captures the limiting behavior of a much larger class of models in the limit that X_{μ} and \bar{X}_i are both small compared to one. In

this case, we can expand the exponential in Eq. (2.1) to obtain the lowest order contribution,

$$\frac{\partial f_\mu}{\partial t} \approx \sum_{i=1}^{\mathcal{R}} \alpha_{\mu,i} [X_\mu - \bar{X}_i(t)] f_\mu + \frac{\xi_\mu(t)}{\sqrt{N}}. \quad (2.3)$$

When $\mathcal{R} = 1$, we recover the standard Wright-Fisher model of population genetics [11], with its logistic growth function, $\partial_t f_\mu = (X_\mu - \bar{X})f_\mu$. The parameter X_μ coincides with the standard measure of (log) fitness. In the presence of multiple resources, Eq. (2.3) can be regarded as the simplest generalization the Wright-Fisher model that incorporates multiple fitness dimensions. From an ecological perspective, Eq. (2.3) can also be viewed as a special version of the Lotka-Volterra model that arises when the interactions between strains are mediated by \mathcal{R} intensive variables (the resource-specific mean fitnesses, \bar{X}_i). This is an important simplification: although there is no closed-form solution for $f_\mu(t)$ when $\mathcal{R} > 1$, Eq. (2.1) still possesses a convex Lyapunov function (SI Appendix 2.6), which implies that $f_\mu(t)$ must approach a unique and stable equilibrium at long times.

The ecological model in Eq. (2.1) only describes the competition between a fixed set of strains. To incorporate evolution, we also allow for new strains to be created through the process of mutation. We will show that it is useful to distinguish between two broad classes of mutations. The first class, which we will refer to as *strategy mutations*, alter an individual's resource uptake strategy. For simplicity, we assume that these mutations constitute a perfect tradeoff, so that the overall fitness X remains unchanged (though we will eventually relax this assumption below). We assume that strategy mutations occur at a per genome rate U_α and result in a new resource strategy $\vec{\alpha}'$ drawn from some distribution $\rho_\alpha(\vec{\alpha}'|\vec{\alpha})$. In addition to strategy mutations, we consider a second class of *pure fitness mutations*, which alter the overall fitness X but leave the resource strategy $\vec{\alpha}$ unchanged. These mutations capture the effects of directional selection at a large number of other loci throughout the genome, which may only be tangentially related to the resource utilization strategy. We assume that these fitness mutations arise at a per-genome rate U_X , and that they increment X by an amount s drawn from the distribution of fitness effects, $\rho_X(s)$. For simplicity, we assume that there is no macroscopic epistasis for fitness [27], so that $\rho_X(s)$ remains the same for all genetic backgrounds.

We note that this division into fitness and strategy mutations is neither exhaustive nor unambiguous. Some changes in resource strategy may also incur a fitness cost (see below), and one can simulate a pure fitness mutation by shifting metabolic effort away from resources that are not present in the current environment (i.e., those with $\beta_i = 0$). Nevertheless, considering these idealized cases as our fundamental axes will prove to be a useful conceptual tool, which provides additional insight into the behavior of our model.

For example, pure fitness mutations might seem like an ecologically trivial addition to the model, because they are always favored to invade. However, computer simulations show that these accumulated fitness differences can still have a dramatic influence on both the ecological structure and the evolutionary dynamics that arise in a given population. Figures 2.1C-F depict individual-based simulations of four populations, which are subject to

the same environmental conditions and the same supply of strategy mutations, but have different values of U_X and $\rho_X(s)$ (SI Appendix 2.6). Depending on the supply of fitness mutations, the behaviors can include rapid diversification and stasis (Fig. 2.1C), unstable but continually renewed coexistence (Fig. 2.1D), stable coexistence and rapid within-clade evolution (Fig. 2.1E), or the permanent disruption of coexistence (Fig. 2.1F). In this way, the seemingly simple model in Eq. (2.1) can produce a diverse range of behaviors, which at least superficially resemble the complex dynamics observed in some microbial evolution experiments.

To understand these different behaviors and how they depend on the underlying parameters, we will start by analyzing the simplest non-trivial scenario, in which the strains evolve in an environment with just two resources. In this case, the environmental supply rates and resource uptake strategies can be described by scalar parameters $\vec{\beta} = (\beta, 1 - \beta)$ and $\vec{\alpha} = (\alpha, 1 - \alpha)$, respectively. This case will already be sufficient to elucidate many of the key qualitative behaviors and fundamental timescales involved, while maximizing analytical tractability. In the second section, we will extend this analysis to larger numbers of resources, and comment on the additional features that are unique to this more complicated scenario.

2.4 Analysis

Selection for ecosystem to match environment, stable coexistence

We will begin by considering the dynamics in the absence of fitness differences ($U_X = 0$, $X_\mu = 0$). The ecological dynamics in this “neutral” scenario have recently been described by Ref. [51], and it will be useful to build on these results when we introduce fitness differences below.

We begin by considering a single strategy mutation that occurs in a clonal population of type α_1 , creating a new strain of type α_2 . The initial dynamics of this mutation can be described by a branching process with growth rate $S_{\text{inv}} = \langle \partial_t f \rangle / f$, also known as the *invasion fitness* (SI Appendix 2.6). In this case, the invasion fitness is given by

$$S_{\text{inv}} = \frac{\Delta\alpha(\beta - \alpha_1)}{\alpha_1(1 - \alpha_1)}, \quad (2.4)$$

where $\Delta\alpha = \alpha_2 - \alpha_1$ is the difference between the mutant and wildtype uptake rates. The invasion fitness is positive whenever $\Delta\alpha$ and $\beta - \alpha_1$ have the same sign: if $\alpha_1 < \beta$, then selection will favor mutations that increase α , while if $\alpha_1 > \beta$, selection will favor mutations that decrease α . In this way, selection tries to tune the population uptake rate to match the environmental supply rate. If $\alpha_1 = \beta$, then the invasion fitness vanishes for all further strategy mutants. This constitutes a marginal evolutionarily stable state (ESS). Using the definition of $\bar{X}_i(t)$ in Eq. (2.2), we can see that the universal dynamics in Eq. (2.3) correspond to a *near-ESS limit*, where the resource uptake rates remain close to β . For the sake of generality, we will focus on this limit for the remainder of the main text. The full expressions for the microscopic model in Eq. (2.1) are listed in the SI Appendix.

When α_1 and α_2 are both close to β , Eq. (2.4) reduces to the quadratic form,

$$S_{\text{inv}} \approx \frac{\Delta\alpha(\beta - \alpha_1)}{\beta(1 - \beta)}. \quad (2.5)$$

Since all mutations first arise in a single individual, many will be lost to genetic drift, even when their invasion fitness is positive. With probability $\sim S_{\text{inv}} \ll 1$, the mutant lineage will survive drift long enough to reach frequency $f \sim 1/NS_{\text{inv}}$, and will then start to increase deterministically at rate S_{inv} . In sufficiently large populations, the transition to deterministic growth will occur long before the mutant starts to influence its own growth rate, so that the constant invasion fitness assumption is justified (SI Appendix 2.6).

At long times, the ecological dynamics will lead to one of two final states: the mutant will either replace the wildtype (competitive exclusion) or the two will coexist at some intermediate frequency (Fig. 2.2A). The latter scenario will occur if and only if the wildtype can re-invade a population of mutants, which requires that the reciprocal invasion fitness, $S_{\text{inv}}^R \approx \Delta\alpha(\alpha_2 - \beta)/\beta(1 - \beta)$, is also positive. By examining this expression, we see that the mutant will outcompete the wildtype if its strategy lies between β and α_1 , while stable coexistence occurs when α_1 and α_2 span β (i.e., $\alpha_1 < \beta < \alpha_2$ or vice versa). When this condition for coexistence is met, the steady-state frequencies are determined by the linear equation,

$$\bar{\alpha} \equiv \sum_{\mu} \alpha_{\mu} f_{\mu}^* = \beta, \quad (2.6)$$

whose solution is given by $f^*/(1 - f^*) = (\beta - \alpha_2)/(\alpha_1 - \beta)$ [51]. In other words, the relative frequencies of the strains are inversely proportional to their distance from the environmental supply rate. According to Eq. (2.6), these frequencies are chosen such that the population-averaged uptake rate $\bar{\alpha} = \sum_{\mu} \alpha_{\mu} f_{\mu}^*$ exactly balances the resource supply rate β . This provides an intuitive explanation for the cause of coexistence: by maintaining the strains at intermediate frequencies, the population is able to match the environmental supply rate more closely than it could with either strain on its own.

Once this ecological equilibrium is attained, number fluctuations will continuously perturb the true frequency away from f^* (Fig. 2.2A), subject to a linearized restoring fitness $\sim \Delta\alpha^2/\beta(1 - \beta)$ (SI Appendix 2.6). The restoring force is strong compared to genetic drift when $N\Delta\alpha^2/\beta(1 - \beta) \gg 1$, which leads to linearized fluctuations of order $\delta f \sim \sqrt{\beta(1 - \beta)/N\Delta\alpha^2}$, and a lifetime for the stable state that is exponentially long in $\sqrt{N\Delta\alpha^2}$. At this point, additional strategy mutants are subject to very weak selection pressures: fluctuations will induce momentary invasion fitnesses of order $\delta S_{\text{inv}} \sim |\alpha_3 - \alpha_2|/\sqrt{N\beta(1 - \beta)}$ (which can be large compared to $1/N$), but these fitness effects are quickly averaged to zero during the $\sim 1/\delta S_{\text{inv}}$ generations required for such a mutation to establish (SI Appendix 2.6). Thus, once the population diversifies to fill the two niches, the rate of evolution dramatically slows down (as in Fig. 2.1A), since the relevant timescales are controlled by genetic drift. In this way, a large effect mutation can allow the ecosystem as a whole to reach an effective ESS, long before any of the constituent strains reach the ESS on their own.

Diversification load

We are now in a position to analyze how fitness alters the basic picture above. We begin by revisiting the invasion of a mutant strain in an initially clonal population, this time allowing for a fitness difference ΔX between the mutant and wildtype. In this case, the new invasion fitness is given by a simple linear combination,

$$S_{\text{inv}}(\Delta\alpha, \Delta X) \approx \Delta X + S_{\text{inv}}(\Delta\alpha), \quad (2.7)$$

where $S_{\text{inv}}(\Delta\alpha)$ is the invasion fitness for a pure strategy mutation from Eq. (2.4). This result describes, in quantitative terms, how selection balances its ecological preferences ($\bar{\alpha} \rightarrow \beta$) with its desire to maximize fitness ($\bar{X} \rightarrow \infty$). When the uptake rate of the resident population is far from the environment supply rate [$\beta - \alpha_1 \sim \mathcal{O}(1)$], the ecological selection pressures can be quite strong, with invasion fitnesses as high as 10% – 100%. This implies that strongly deleterious mutations of order

$$\Delta X_{\text{min}} \approx -\frac{\Delta\alpha(\beta - \alpha_1)}{\beta(1 - \beta)} \quad (2.8)$$

can hitchhike to fixation when the population colonizes a new ecological niche (a form of *diversification load*).

Fitness differences perturb ecological equilibria

In addition to shifting the invasion fitness of a new mutation, fitness differences can also alter the long-term ecological equilibrium between mutant and wildtype in Eq. (2.6). In the extreme limit, this can disrupt the stable coexistence altogether. If the mutant is less fit than the wildtype ($\Delta X < 0$), this will occur whenever ΔX is less than the maximum diversification load ΔX_{min} in Eq. (2.8). On the other hand, if $\Delta X > 0$, extinction will occur when the wildtype no longer has positive invasion fitness, or when ΔX exceeds a threshold

$$\Delta X_{\text{max}} \approx \frac{\Delta\alpha(\alpha_2 - \beta)}{\beta(1 - \beta)}. \quad (2.9)$$

We note that the fitness differences in Eqs. (2.8) and (2.9) are lower than the values required for the mutant or wildtype to dominate in *all* environmental conditions (SI Appendix 2.6). Instead, the fitness thresholds strongly depend on how the resource strategies differ from each other, and how they differ from the environmental supply rate. When $\Delta\alpha \sim \epsilon$, even small fitness differences ($\Delta X \sim \epsilon^2$) can disrupt the stable ecology, while for $\Delta\alpha \sim \mathcal{O}(1)$, much larger fitness differences ($\Delta X \gtrsim 100\%$) can be tolerated.

When $\Delta X_{\text{min}} < \Delta X < \Delta X_{\text{max}}$, the two strains continue to coexist, but their equilibrium frequency is no longer given by Eq. (2.6). In this case, the competing drive to maximize fitness means that selection will no longer favor an ecology that matches the environmental

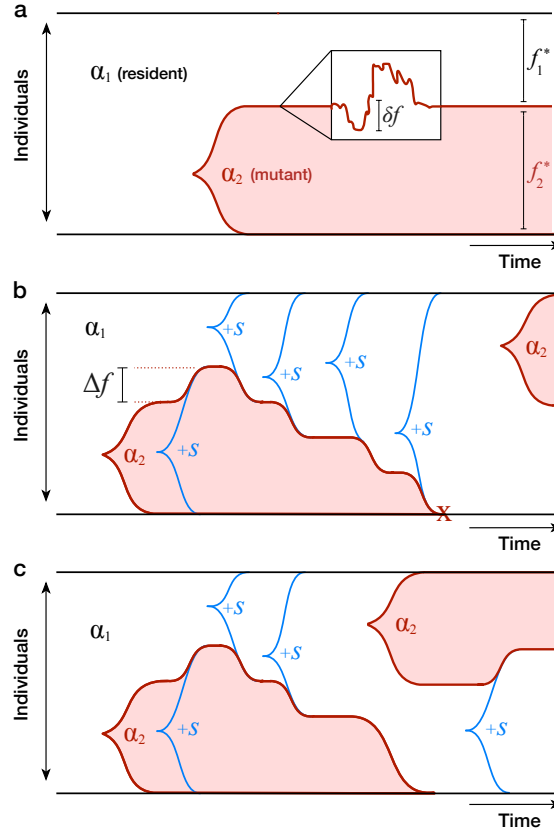


Figure 2.2: **Schematic illustration of key eco-evolutionary processes in a two-resource ecosystem.** (a) Ecological diversification from a clonal ancestor. In the absence of fitness mutations, strains coexist at a stable equilibrium (f^*) with fluctuations (δf) controlled by genetic drift. Further strategy mutations are not favored to invade. (b) Pure fitness mutations that sweep within an ecotype shift the stable equilibrium by Δf ; accumulated fitness differences can ultimately drive ecotypes to extinction. Further strategy mutations allow the winning clade to re-diversify at a later time. (c) Occupied niches can also be invaded by strategy mutations that arise in fitter genetic backgrounds. In this case, the original ecotype lineage is driven to extinction while the ecological structure of the community is preserved.

resource distribution, at least not perfectly. In SI Appendix 2.6, we show that the new equilibrium frequency is given by

$$f^*(\Delta X) \approx f_0^* + \frac{\beta(1-\beta)}{\Delta\alpha^2} \cdot \Delta X, \quad (2.10)$$

where f_0^* is the neutral ecological equilibrium from Eq. (2.6). From this expression, we can read off the typical fitness differences required to perturb f^* from its present value. This

fitness sensitivity is again determined by the distance between the two resource strategies. If $\Delta\alpha \sim \mathcal{O}(1)$, large fitness differences ($\Delta X \gtrsim 100\%$) are required to change the equilibrium frequency, while for $\Delta\alpha \sim \epsilon$, even very small fitness differences ($\Delta X \sim \epsilon^2$) can generate large changes in the equilibrium frequency.

Further fitness evolution and diversification-selection balance

Once the population achieves the stable ecology in Eq. (2.10), additional fitness mutations will occur in each strain with probability proportional to the equilibrium frequency f^* . In our model, the invasion fitness of such a mutation is simply its fitness effect s , independent of the ecological state of the population. With probability $\sim s$, this mutation will sweep through its parent clade, changing the fitness difference between the clades by $\pm s$ and the equilibrium frequency by $\Delta f = f^*(\Delta X \pm s) - f^*(\Delta X)$ (Fig. 2.2B). In the linear regime of Eq. (2.10), the frequency and fitness changes are directly related,

$$\Delta f \approx \pm \left(\frac{s}{s_c} \right), \quad (2.11)$$

where $s_c = \Delta\alpha^2/\beta(1-\beta)$ is the fitness scale that determines changes in equilibrium frequency. If $s \gg s_c f^*(1-f^*)$, then the stable coexistence will be disrupted, and the mutant clade will take over the population. We will refer to such a scenario as *ecosystem collapse*, since one of the niches is no longer occupied.

Similar behavior can occur when $s \ll s_c$ as well, except that now the ecosystem collapse occurs due to cumulative effect of many pure fitness mutations. When the fitness mutations accumulate independently, this process can be described by an effective diffusion model,

$$s_c \frac{\partial f^*}{\partial t} \approx 2NU_X s^2 (2f^* - 1) + \sqrt{2NU_X s^3} \cdot \eta(t) \quad (2.12)$$

with a bias that reflects the higher probability of producing a mutation in a larger clade (SI Appendix 2.6). Eq. (2.12) superficially resembles the drift-induced perturbations at ecological equilibrium, except that the bias is now unstable rather than restoring. When $2f^* - 1 \ll \sqrt{s/s_c}$, the mutation bias is weak, and the clade frequencies undergo a random walk ($\delta f^* \sim \sqrt{\frac{NU_X s^3}{s_c^2}} \cdot \delta t$). But after a time of order $\tau_{\text{drift}} \sim \frac{s_c}{NU_X s^2}$, the frequency differential grows large enough that the more prevalent clade will deterministically produce more beneficial mutations, so that it is destined for fixation. After a time of order $\tau_{\text{collapse}} \sim \frac{s_c}{NU_X s^2} \log\left(\frac{s_c}{s}\right)$, the fitness difference between the clades grows so large that the ecosystem finally collapses (Fig. 2.2B). This timescale sets an upper limit on the lifetime of the stable state when many fitness mutations are available.

Once the ecosystem collapses, there will be a strong selection pressure for the winning clade to re-diversify through additional strategy mutations, and restart this process from the beginning (Fig. 2.2B). To gain insight into these dynamics, we first consider the case where the resource strategies are controlled by a single genetic locus, with fixed phenotypes

α_1 and α_2 , and mutations that alternate between the two states at rate U_α . After an ecosystem collapse, Eq. (2.4) shows that the invasion fitness for the opposite strategy is given by $S_{\text{inv}} \sim s_c$, so the collapsed state will persist for a time of order $\tau_{\text{diversify}} \sim 1/NU_\alpha s_c$, until the stable ecology is re-established. If the two strategies are symmetric about β , so that $f^*(0) = 1/2$, the new stable state will persist for $\sim \tau_{\text{collapse}}$ generations in the absence of additional strategy mutations, and the process will then repeat itself. The relative probability of observing the population in the collapsed ($\mathcal{S} = 1$) or saturated ($\mathcal{S} = 2$) states is therefore given by

$$\frac{\Pr[\mathcal{S} = 2]}{\Pr[\mathcal{S} = 1]} \approx \frac{\tau_{\text{collapse}}}{\tau_{\text{diversify}}} \sim \begin{cases} \frac{U_\alpha}{U_X} \left(\frac{s_c}{s}\right) & \text{if } s \gg s_c, \\ \frac{U_\alpha}{U_X} \left(\frac{s_c}{s}\right)^2 \log\left(\frac{s_c}{s}\right) & \text{if } s \ll s_c. \end{cases} \quad (2.13)$$

This expression shows the minimum amount of strategy mutations, or the maximum amount of pure fitness mutations, that allow the population to maintain a saturated ecosystem. We will refer to this dynamic steady state as *diversification-selection balance*, in analogy to mutation-selection balance in population genetics [24]. Note that this balance crucially depends on the state of the ecosystem through $s_c \sim \Delta\alpha^2/\beta(1-\beta)$. All else being equal, ecosystems with more similar resource uptake strategies will be disrupted more easily than those with a higher degree of specialization.

Invading ecotypes can delay ecosystem collapse

Strictly speaking, the derivation of Eq. (2.13) is only valid in the limit that $\tau_{\text{collapse}} \ll \tau_{\text{diversify}}$, since we neglected mutations between α_1 and α_2 when both niches were filled. When $\tau_{\text{collapse}} \gtrsim \tau_{\text{diversify}}$ (i.e., when the ecosystem spends an appreciable amount of time in the saturated state), we must also account for mutations between the two strategies that arise before the ecosystem collapses. Those mutations that arise in the less-fit clade will have little chance of invading. However, a mutation from the more-fit to the less-fit strategy will establish with probability $\sim |\Delta X|$, where ΔX is the current fitness difference between the two clades. If this mutation is successful, it will outcompete the resident lineage with the corresponding value of α , and reset the fitness difference to $\Delta X = 0$ (Fig. 2.2C). In this way, invasion from one ecotype to another can significantly delay the process of ecosystem collapse, since it relieves the tension between fitness maximization and $(\bar{X} \rightarrow \infty)$ and selection to match the environment $(\bar{\alpha} \rightarrow \beta)$.

To analyze this process, we note that successful invasion events will occur as an inhomogeneous poisson process with rate $\lambda(t) \sim NU_\alpha f_{\text{argmax}(X_i)}^* |\Delta X|$, where $f^*(t)$ and $\Delta X(t)$ are again determined by the diffusion model in Eq. (2.12). This leads to a characteristic invasion

timescale

$$\tau_{\text{invade}} \sim \begin{cases} \frac{1}{NU_X s} & \text{if } U_\alpha \gg U_X, \\ \frac{1}{NU_X s} \left(\frac{U_X}{U_\alpha} \right)^{2/3} & \text{if } U_\alpha \gg U_X \left(\frac{s}{s_c} \right)^{3/2}, \\ \frac{s_c}{NU_X s^2} \log \left(\frac{U_X^2 s^3}{U_\alpha^2 s_c^3} \right) & \text{if } U_\alpha \gg U_X \left(\frac{s}{s_c} \right)^2, \\ \infty & \text{else.} \end{cases} \quad (2.14)$$

which is derived in SI Appendix 2.6. Each of these regimes corresponds to a different intuitive picture of the dynamics. In the first case, strategy mutations are frequent compared to pure fitness mutations, and invasion occurs almost immediately after the first fitness mutation arises. In the second case, invasion occurs after multiple fitness mutations have accumulated, but when the frequencies of the clades still wander diffusively relative to each other [$f^* \approx 1/2 \pm \mathcal{O}(\sqrt{s/s_c})$]. In the third regime, invasion occurs after one of the clades has grown to a sufficiently large frequency that it would have deterministically led to ecosystem collapse. When the invading mutant establishes, it will therefore cause a rapid shift in the frequencies of the ecotypes as $f^*(\Delta X)$ returns to $f^*(0)$.

Finally, when $U_\alpha \ll U_X \left(\frac{s}{s_c} \right)^2$, strategy mutants are sufficiently rare that the ecosystem will typically collapse and re-diversify before invasion can occur. This sets the region of validity of the diversification-selection balance in Eq. (2.13). Interestingly, Eq. (2.13) shows that collapse and re-diversification can still dominate over invasion even when both niches are typically filled ($\Pr[\mathcal{S} = 2] \gg \Pr[\mathcal{S} = 1]$). In this case, both the genealogical structure and the typical state of the ecosystem will resemble the invasion regime, but the historical record would contain a series of punctuated extinction and diversification events, interspersed with long periods of gradual fitness evolution.

Fitness differences create opportunities for ecological tinkering

Our derivation of Eqs. (2.12) and (2.14) assumed that the two ecotypes were fixed by the genetic architecture of the organism. Individuals could mutate between α_1 and α_2 , but mutations to other points in strategy space were forbidden. In the absence of fitness differences, we saw that selection for these additional strategy mutants is weak once both niches have been filled ($S_{\text{inv}} \lesssim 1/N$), potentially justifying the single-locus assumption in terms of a priority effect. However, the previous analysis shows that there can be strong selection to switch strategies once fitness mutations accumulate, so it is also plausible that fitness differences could lead to selection for strategy mutations more generally.

To investigate these selection pressures, we consider a population that is currently described by the steady state in Eq. (2.10). We then consider strategy mutations that occur on the background of α_2 , altering its strategy to α_3 while leaving its fitness intact. The invasion

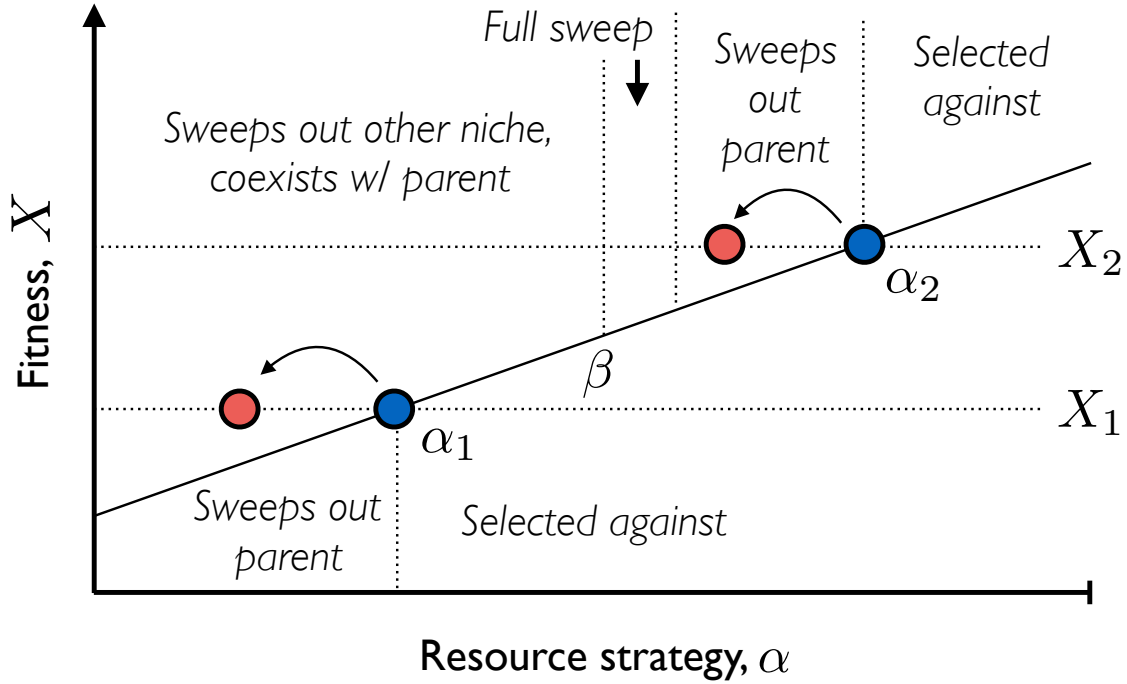


Figure 2.3: **Invasion fitness landscape for additional strategy mutations in a two-resource ecosystem.** The two resident ecotypes are illustrated by blue circles, while red circles denote mutant strains created by strategy mutations on one of the ecotype backgrounds. The solid black line denotes the effective mean fitness, $\sum_i \alpha_i \bar{X}_i$, experienced by a given resource strategy. Strains with overall fitness (X) above this line are favored to invade, while others are selected against. If a mutant successfully invades, its effect on the ecosystem is indicated by the text.

fitness for such a mutation is given by

$$S_{\text{inv}} \approx \frac{(\alpha_2 - \alpha_3)\Delta X}{\Delta \alpha}, \quad (2.15)$$

As anticipated, fitness differences create additional selection pressure for strategy mutations beyond the simple switching behavior considered above.

The direction of selection is determined by the sign of ΔX . In the background of the fitter clade ($\Delta X > 0$), selection favors mutations that increase the strategy in the direction of β (a form of generalism), while simultaneously disfavoring mutations that lead to increased specialization (Fig. 2.3). The opposite behavior occurs in the less-fit background, with selection favoring mutations that increase the distance from β , leading to increased specialization. Both behaviors have an intuitive explanation in terms of individuals preferring

to allocate their metabolic energy toward the resource with the least-fit consumers, thereby minimizing the effective competition that they experience.

Once a successful strategy mutation arises, it will sweep through part of the population and alter the ecological equilibrium (Fig. 2.3). Mutations in the less-fit clade are straightforward to analyze. Since these are always directed away from both β and α_1 , these mutants will sweep through their parent clade and increase the equilibrium frequency according to Eq. (2.10). Successful mutations in the fitter clade have a wider range of outcomes, since these are always directed toward β and α_1 . If $\alpha_3 < \beta$, the mutant lineage will outcompete the less-fit strain α_1 , and will stably coexist with its parent clade α_2 at an equilibrium frequency $f^* = (\beta - \alpha_2)/(\alpha_3 - \alpha_2)$. On the other hand, if $\beta < \alpha_3 < \alpha_2$, the mutant lineage will always sweep through its parent clade α_2 . If α_3 is sufficiently close to α_2 , this will simply lead to an increase in frequency according to Eq. (2.10). However, if α_3 is close enough to β that $\Delta X_{\max}(\alpha_3)$ becomes less than the actual fitness difference, ΔX , then the mutant will sweep out both clades and lead to an ecosystem collapse and subsequent re-diversification. Thus, in addition to creating a larger target for invasion events, these additional strategy mutations can also enhance the probability of ecosystem collapse. The balance between these competing tendencies will depend on the genetic architecture of the resource strategies, $\rho_\alpha(\alpha'|\alpha)$, which is poorly parameterized by existing data. A detailed analysis of the potential regimes will be left for future work.

Beyond pairwise coexistence

Our previous analysis focused on environments with only two substitutable resources, where at most two strains can coexist at equilibrium. In this case, the structure of the stable ecosystem was simple enough to admit a full analytical solution, which we could use to derive explicit predictions for many evolutionary quantities of interest. However, many microbial communities are found in environments with large numbers of potential resources, and flexible gene pools that allow them to alter their resource uptake rates through horizontal gene transfer [46]. It is therefore natural to ask how our results generalize to these more complicated environments as well. A full analysis of this case is beyond the scope of the present work, as there are even fewer constraints on the space of ecological and evolutionary parameters compared to the two resource case. Nevertheless, it is still useful to know whether our qualitative results extend beyond $\mathcal{R} = 2$, and whether there are fundamentally new behaviors that only arise in higher dimensions.

For a general ecological equilibrium, a mutation that alters the phenotype of a resident strain from $(X_\mu, \vec{\alpha}_\mu)$ to $(X_\mu + s, \vec{\alpha}_\mu + \Delta\vec{\alpha})$ will have an invasion fitness

$$S_{\text{inv}} \approx s - \sum_i \Delta\alpha_i \bar{X}_i, \quad (2.16)$$

where the resource-specific mean fitnesses in Eq. (2.2) must be evaluated at the equilibrium frequencies f_μ^* (SI Appendix 2.6). Increases in α_i are favored when \bar{X}_i is lower than the




Figure 2.4: **Diversification-selection balance when $\mathcal{R} \gg 1$.** Circles depict the long-term steady state from SSWM simulations of a binary resource usage model in a nearly uniform environment (SI Appendix 2.6). Each point denotes an average over multiple timepoints from 3 independent replicates; solid lines indicate the minimum and maximum replicate. (a) The standard deviation in \bar{X}_i across the \mathcal{R} resources. (b) The number of coexisting ecotypes. The colored dashed lines denote the maximum ecosystem capacity $\mathcal{S} = \mathcal{R}$. The black dashed line depicts the scaling relation in Eq. (2.19) which applies for $\mathcal{S} \ll \mathcal{R}$, with an $\mathcal{O}(1)$ prefactor of $1/\sqrt{2\pi}$ included for visualization.

“effort-averaged” \bar{X}_i for the other resources, and vice-versa. Thus, similar to the two-resource case, there is still a sense in which selection favors mutations that flow from high values of \bar{X}_i to lower values of \bar{X}_i , though there are now $\binom{\mathcal{R}}{2}$ beneficial directions, $\bar{X}_i \rightarrow \bar{X}_j$, rather than just one.

The invasion fitness in Eq. (2.16) depends on the current community composition only through the intensive variables \bar{X}_i . In a *saturated* ecosystem, where the number of coexisting strains is equal to the number of resources, these can be directly obtained by a matrix

inversion of Eq. (2.1),

$$\bar{X}_i \approx \sum_{\mu} \alpha_{i,\mu}^{-1} X_{\mu}, \quad (2.17)$$

where $\alpha_{i,\mu}^{-1}$ is the left inverse of $\alpha_{\mu,i}$. Thus, we see that in a saturated ecosystem, the \bar{X}_i are given by linear combinations of the strain fitnesses X_{μ} , justifying their interpretation as resource-specific mean fitnesses. Moreover, perturbation expansions of $\alpha_{i,\mu}^{-1}$ suggest that the prefactor is still inversely proportional to an effective distance between the strategies (SI Appendix 2.6), similar to the two-resource case in Eq. (2.15). We note that the equilibrium values of \bar{X}_i are conditionally independent of both the resource supply vector β_i and the strain frequencies f_{μ}^* ; these quantities influence \bar{X}_i only through shaping the set of resource strategies that coexist at equilibrium. Thus, these saturated ecosystems dynamically adjust their composition to screen the internal selection pressures \bar{X}_i from the external environmental conditions. Similar findings were recently obtained for the neutral case [where $\bar{X}_i = 0$ [51]], as well as in certain community assembly processes in the $\mathcal{R} \rightarrow \infty$ limit [59, 3]. Eq. (2.17) shows that this is a generic property that occurs whenever the number of surviving species is equal to the number of resources.

In this limit, the steady-state frequencies f_{μ}^* can be obtained from a similar matrix inversion,

$$f_{\mu}^* \approx \sum_i \beta_i \alpha_{i,\mu}^{-1} - \sum_{i,\nu} \beta_i \alpha_{i,\mu}^{-1} \alpha_{i,\nu}^{-1} (X_{\mu} - X_{\nu}), \quad (2.18)$$

which serves as the generalization of Eq. (2.10) for multiple resources (SI Appendix 2.6). As in the two-resource case, small fitness differences perturb the neutral ecological equilibrium via linear combinations of the strain fitnesses, with a prefactor that is inversely proportional to the square of the effective distance between the resource strategies.

While the saturated case is particularly simple, we saw above that fitness mutations can drive the number of surviving species below this saturated value. In contrast to the two-resource case, these *unsaturated* ecosystems can now harbor multiple coexisting strains when $\mathcal{R} > 2$, leading to a continuous generalization of the diversification-selection balance in Eq. (2.13). To investigate this effect, we performed computer simulations of a binary strategy space model, in which individuals can either utilize or not utilize a given resource, with mutations that toggle individual uptake rates on and off (SI Appendix 2.6). The results recapitulate the qualitative behavior observed for $\mathcal{R} = 2$ resources, in that a sufficiently high rate of pure fitness mutations can constrain the number of distinct strategies that are able to coexist (Fig. 2.4B). To compensate for the strong ecological selection pressures that can arise when $\mathcal{S} \ll \mathcal{R}$, the populations are forced to evolve consortia of “generalist” strains such that $\sum_{\mu} \bar{\alpha}_{\mu} f_{\mu}^*$ is still close to $\vec{\beta}$, at least at lowest order (Fig. 2.4A).

In a nearly uniform environment [$\beta_i \propto 1 \pm \mathcal{O}(\epsilon)$], simulations show that the steady-state ecosystem tends to be dominated by a single “generalist” strain ($\alpha_{\mu,i} \propto 1$), and a collection of $\mathcal{S} - 1$ single loss-of-function variants ($\alpha_{\mu,i} \propto 1 - \delta_{\mu,i}$) that recently descended




Figure 2.5: **Schematic of (a) ecological and (b) genealogical structure at the evolutionary steady-state described in Eq. (2.19).** In (b), blue dots represent pure fitness mutations and red dots represent loss-of-function strategy mutations.

from mutations in the generalist background (Figs. 2.5 and 2.7-2.9). This is reminiscent of a mutation-load argument [23], in which the preference for the generalist strain is balanced by the greater entropy of loss-of-function mutants. However, a key difference in this case is that the generalist strain is not actually favored by selection. By definition, all of the transient states in Fig. 2.5 are ecologically stable, while the preference for the generalist strain arises dynamically from the race to acquire pure fitness mutations.

The dynamics of this process can be characterized analytically in the weak mutation limit, yielding a simple heuristic expression for the diversification-selection balance,

$$\mathcal{S} \sim \frac{1}{\mathcal{R}} \left(\frac{U_\alpha \epsilon^2}{U_X s^2} \right), \quad (2.19)$$

which is valid for $\mathcal{S} \ll \mathcal{R}$ (SI Appendix 2.6). The transition to the fully saturated state ($\mathcal{S} = \mathcal{R}$) requires an even more stringent condition, which implies that unsaturated ecosystems are obtained for a very broad parameter regime (Fig. 2.4B). In both cases, a larger number of substitutable resources will lead to a less diverse ecosystem at diversification-selection balance. This is ultimately due to the fact that the difference between generalists and single loss-of-function variants becomes increasingly small as $\mathcal{R} \rightarrow \infty$.

This suggests that the relative frailty of the diversification-selection balance in Eq. (2.19) may be a pathological feature of the simple genetic architecture that we have assumed, in which fit generalist phenotypes are easily accessible. If we instead impose an upper limit $\mathcal{R}_c \ll \mathcal{R}$ on the number of resources that a strain can metabolize, heuristic calculations suggest that diversification-selection balance will be achieved for substantially higher values

of \mathcal{S} , even for large \mathcal{R} (SI Appendix 2.6). In this case, the ecological and genealogical structures that are attained at this evolutionary steady state will be considerably more complex than the shallow star-shaped genealogies in Fig. 2.5. A more detailed analysis of this steady state will be considered in future work.

2.5 Discussion

In microbial populations, primitive ecological interactions can evolve spontaneously over years [53], months [30], and even days [52]. Yet this process rarely takes place in isolation. In rapidly evolving populations, diversifying selection must compete with directional selection acting on other loci throughout the genome. Here, we have introduced a simple mathematical framework to model the interactions between these two processes in asexually reproducing organisms.

The ecological interactions in our model emerge from the competition for substitutable resources (e.g. different carbon sources), according to a well-studied class of models from theoretical ecology [61, 51, 59, 3]. To incorporate evolution into this model, we assumed that individuals can acquire mutations that alter their resource uptake rates. We showed that it is useful to distinguish between two characteristic types of mutations: (i) strategy mutations, which divert metabolic effort from one resource to another and (ii) fitness mutations, which increase the overall growth rate but leave the relative uptake rates unchanged. Strategy mutations enable ecological diversification, while fitness mutations capture the effects of directional selection at other genomic loci.

This classification scheme is best viewed as a conceptual tool, rather than a statement about the underlying biology. We have mostly focused on mutations with either a perfect tradeoff or a perfect benefit, but Eqs. (2.7) and (2.16) apply equally well in more realistic cases where a shift in resource strategy is accompanied by a change in the overall growth rate. These expressions can be used to predict when the costs of an opportunistic mutation will outweigh its ecological benefits, or vice versa. Similarly, true fitness mutations (e.g. an increase in ATP efficiency) are assumed to be rare in nature, since they could have fixed in the population long ago. In practice, *effective* fitness mutations are more likely to correspond to strategy mutations whose tradeoffs are simply not exposed by the current environmental conditions. In this picture, the overall fitness X_μ can be viewed as an emergent trait that arises whenever we project high-dimensional cellular phenotypes onto a restricted set of axes (SI Appendix 2.6). The presence of effective fitness mutations may therefore be a more general phenomenon than their name would seem to imply.

The creation of new strains via mutation bears a superficial resemblance to immigration from a fixed species pool, which is the traditional scenario considered in theoretical ecology. However, this analogy is exact only in the absence of inheritance, when the phenotypes of nearby genotypes are uncorrelated from each other. In contrast, when the effects of mutations are heritable, we have seen that directional selection can produce dramatic departures from traditional ecological predictions.

Similar to immigration [60, 3], strategy mutations allow an initially clonal population to diversify into stably coexisting ecotypes, whose upper bound is set by the number of resources. Yet because fitness mutations are heritable, further evolution will lead to fitness differences between the clades, which can dynamically shift the ecological equilibrium over time, and eventually drive less fit clades to extinction. The mere observation that selection can disrupt coexistence is not surprising, since drug resistance or other harsh selection regimes provide striking examples of this effect. However, our quantitative analysis shows that this collapse can happen long before any clade is universally inferior to another, and that it can result from the compound effect of many small effect mutations that would not lead to extinction on their own. These results suggest that ongoing directional selection may have a larger impact on the structure of microbial communities than is often assumed. In particular, while previous ecological analyses suggest that the number of ecotypes should meet [59, 3] or even exceed [51] the number of resources, our results raise the possibility that they could also reside at a *diversification-selection balance* below the maximum capacity of the ecosystem.

In addition to their influence on coexistence, we also found that fitness differences accrued via directional selection will generate emergent selection pressures for continual evolution of the ecological phenotypes, even in a saturated ecosystem. While these internal selection pressures are reminiscent of the Red Queen effect [64], our quantitative analysis shows that they select for different phenotypes than in the standard predator-prey setting. In particular, less fit clades do not experience increased selection pressure to narrow their fitness deficit by accumulating fitness mutations. Instead, selection favors mutations that divert metabolic effort toward resources with lower effective competition, even at the cost of widening the fitness deficit. Moreover, the direction of selection toward any given resource can shift dynamically as the fitness differences and resource uptake strategies evolve over time.

Most of our analysis focused on the strong-selection weak-mutation regime, in which the current ecological equilibrium is attained before the next mutation occurs. In this limit, when the resource uptake strategies are sufficiently close to the supply rates, our model takes on a universal form that closely resembles traditional models of adaptive dynamics [21, 10]. The key difference is that directional selection behaves as an additional trait dimension, which is effectively constrained to remain far from its optimum at all times (Fig. 2.6, SI Appendix 2.6). Our results show that this simple broken symmetry can lead to dramatic deviations from the standard adaptive dynamics picture.

In contrast to adaptive dynamics, we also allow for mutations that have non-infinitesimal effects on resource uptake rates, which turn out to play a key role in controlling the dynamical behaviors that we observe. In practice, the genetic architectures of most ecological interactions remain poorly characterized empirically. In a few well-studied cases, ecological diversification can be traced to a single large-effect mutation [17, 45], while in others, a series of smaller mutations have been implicated [49]. Our present analysis suggests new ways in which we might constrain this key parameter experimentally, either by analyzing fluctuations in ecotype frequencies on long timescales [29], or by measuring the joint distribution invasion fitness (S_{inv}) and ecological perturbation (Δf) across a panel of engineered mutations [49].

Of course, the present work has focused on a highly simplified model, which omits many of the complicating factors expected in either natural or laboratory settings. A particularly important limitation is our focus on the weak mutation limit ($NU \ll 1$). While analytically convenient, this assumption is violated by many of the laboratory experiments that motivated this study. In SI Appendix 2.6, we describe a preliminary extension of our results to the case where $NU \gg 1$, which builds on the intuition gleaned from the weak mutation limit. However, a more thorough investigation of this regime is required to quantitatively fit the model to evolutionary measurements (e.g. in an approximate-likelihood framework [27]). Future work will also be required to fully explore the effects of crossfeeding, time-varying environments, recombination, and other additions to our basic model [e.g. K -selection [42]]. We believe that our results provide a promising analytical framework in which to investigate these effects, which have mostly been confined to simulations so far.

It is also interesting to ask whether our results can mapped onto more diverse modes of ecological interaction, or whether there are other universality classes yet to be discovered. Since our model can be viewed as the simplest generalization of population genetics with multiple fitness axes, we hypothesize that it may capture the limiting behavior of a broader class of ecological interactions that are mediated by a small number of intensive variables. If so, its analytical tractability may offer a promising avenue for investigating the interactions between ecology and evolution more generally.

2.6 Supporting Information (SI)

Derivation of the model

In this section, we show how the coarse-grained Langevin dynamics in Eqs. (2.1) and (2.2) emerge from two different microscopic models. The first is a simplified class of consumer-resource models described in the main text. We also describe a second implementation that is a more direct extension of the traditional Wright-Fisher model. Motivated by results from population genetics, we then describe a special limit of Eq. (2.1) that is expected to describe the behavior of a much broader class of models, which differ in some of their microscopic details. To illustrate this limit, we show how it applies to a class of consumer-resource models analyzed by Refs. [59, 3].

Consumer-resource model

Our consumer-resource derivation closely follows the one described by Ref. [51], except that we now allow strains to vary in their total energy budget. In particular, we assume that all strains μ and resources i are present in a well-mixed volume V , which is diluted at rate D . In the consumer-resource framework, the per capita growth rate of each strain is mediated by the resource concentrations,

$$\partial_t n_\mu = g_\mu(\vec{c}) n_\mu - D n_\mu + \sqrt{n_\mu D} \cdot \eta_\mu(t), \quad (2.20)$$

where n_μ is the absolute number of individuals of strain μ , $g_\mu(\vec{c})$ is a strain-specific growth function, and $\eta_\mu(t)$ is a Brownian noise term [19] whose mean and covariance are given by

$$\langle \eta_\mu(t) \rangle = 0, \quad (2.21)$$

$$\langle \eta_\mu(t) \eta_\nu(t') \rangle = \delta_{\mu,\nu} \delta(t - t'). \quad (2.22)$$

The resource concentrations (in units of V^{-1}) obey a second set of equations,

$$\partial_t c_i = S_i - D c_i - \sum_\mu \frac{d_{\mu,i}(\vec{c}) n_\mu}{V} + \sqrt{\frac{c_i D}{V}} \cdot \eta_i(t), \quad (2.23)$$

where S_i is the input flux of resource i , $d_{\mu,i}(\vec{c})$ is the per capita depletion rate of resource i by strain μ , and $\eta_i(t)$ is an analogous set of noise terms that describe fluctuations in the resource concentrations. This general class of models has been studied previously by Refs. [41, 61], and others. Following Ref. [51], we consider a restricted subset of models where the growth and depletion functions take on a particularly simple form:

$$g_\mu(\vec{c}) = \sum_i b_{\mu,i} d_{\mu,i}(\vec{c}), \quad (2.24)$$

$$d_{\mu,i}(\vec{c}) = r_{\mu,i} \lambda_i(\vec{c}). \quad (2.25)$$

The first assumption states that the resources are *effectively substitutable*, i.e. biomass can be produced equally well from suitably normalized versions of any imported resource. The constant normalization factor $b_{\mu,i}^{-1}$ can be interpreted as the amount of imported resource i necessary to create one cell of strain μ . The second assumption states that the resource uptake rates can be factored into a species- and resource-specific (but concentration independent) factor $r_{\mu,i}$, and a species-independent (but resource and concentration-specific) function $\lambda_i(\vec{c})$. For example, $\lambda_i(\vec{c})$ could denote the uptake rate of a pathway that imports resource i , while $r_{\mu,i}$ denotes the constitutive expression of that pathway in an individual of strain μ . In this picture, strains can differ in their overall expression of a given pathway, but have limited ability to tune its biochemical properties. This should be a good approximation for strains that have recently descended from a common ancestor, though it may be violated for more rapidly evolving enzymes in distantly related species.

We assume that the resource fluxes and concentrations are both large, such that the dilution and noise terms can be neglected in Eq. (2.23). Following Ref. [51], we also assume a separation of timescales between the dynamics of resource concentrations, such that the resource concentrations reach a quasi-equilibrium $S_i V \approx \sum_{\mu} d_{\mu,i}(\vec{c}) n_{\mu}$ before the strain abundances start to change significantly. Under these assumptions, we can eliminate the concentration variables entirely, and obtain a set of coarse-grained dynamics for the strain abundances:

$$\partial_t n_{\mu} = \left[-D + \sum_i \frac{S_i V b_{\mu,i} r_{\mu,i}}{\sum_{\nu} r_{\nu,i} n_{\nu}} \right] n_{\mu} + \sqrt{n_{\mu} D} \cdot \eta_{\mu}(t). \quad (2.26)$$

In this model, the dynamics of the total number of individuals, $\hat{N}(t) = \sum_{\mu} n_{\mu}(t)$, does not close, due to the μ dependence in the biomass conversion factor $b_{\mu,i}$:

$$\frac{\partial \hat{N}}{\partial t} = -D \hat{N} + \sum_i S_i V \left[\frac{\sum_{\mu} b_{\mu,i} r_{\mu,i} n_{\mu}}{\sum_{\nu} r_{\nu,i} n_{\nu}} \right] + \sqrt{\hat{N} D} \cdot \eta_{\mu}(t). \quad (2.27)$$

However, if we assume that the strains share similar biomass conversion factors $b_{\mu,i} \approx b_i$ (similar to our previous assumption that $\lambda(\vec{c})$ is independent of μ), then the equation for $\hat{N}(t)$ closes. We find that $\hat{N}(t)$ rapidly approaches a steady-state value $N \equiv \sum_i S_i b_i V / D$ on a timescale of order $1/D$, with fluctuations of order \sqrt{N} . Such fluctuations become irrelevant in the large N limit, which suggests that we rewrite the dynamics in terms of the strain frequencies, $f_{\mu} = n_{\mu} / \sum_{\nu} n_{\nu}$ and measure time in units of D^{-1} . Following the derivation in Ref. [26], the dynamics of the frequencies f_{μ} can be shown to satisfy

$$\begin{aligned} \frac{\partial f_{\mu}}{\partial t} = & \left[-1 + \sum_i \frac{\beta_i \alpha_{\mu,i} e^{X_{\mu}}}{\sum_{\nu} \alpha_{\nu,i} e^{X_{\nu}} f_{\nu}} \right] f_{\mu} \\ & + \sum_{\nu} [\delta_{\mu,\nu} - f_{\mu}] \sqrt{\frac{f_{\nu}}{N}} \eta_{\nu}(t), \end{aligned} \quad (2.28)$$

where we have defined the normalized parameters

$$\beta_i = \frac{S_i b_i}{\sum_j S_j b_j}, \quad (2.29)$$

$$\alpha_{\mu,i} = \frac{r_{\mu,i}}{\sum_j r_{\mu,j}}, \quad (2.30)$$

$$X_\mu = \log \sum_i r_{\mu,i}. \quad (2.31)$$

The stochastic noise term $\xi_\mu(t)$ in Eq. (2.1) in the main text can therefore be identified with the linear combination

$$\xi_\mu(t) = \sum_\nu [\delta_{\mu,\nu} - f_\mu] \sqrt{f_\nu} \eta_\nu(t), \quad (2.32)$$

whose correlation structure ensures that $\sum_\mu f_\mu(t) = 1$ at all times.

Deterministic Lyapunov function

The deterministic part of Eq. (2.1) possesses a Lyapunov function,

$$\Lambda(\vec{f}) = - \sum_\mu f_\mu + \sum_i \beta_i \log \left(\sum_\mu \alpha_{\mu,i} e^{X_\mu} f_\mu / \beta_i \right), \quad (2.33)$$

$$= - \sum_\mu f_\mu + \sum_i \beta_i \bar{X}_i, \quad (2.34)$$

which is convex and bounded from above, and for which

$$\frac{d\Lambda}{dt} = \sum_\mu \frac{1}{f_\mu} \left(\frac{df_\mu}{dt} \right)^2 \geq 0. \quad (2.35)$$

Among other things, this implies that the deterministic dynamics have a unique equilibrium that is approached at long times. We exploit this fact in the simulations in Appendix 2.6.

Subdivided environment model

The familiar form of Eq. (2.28) suggests that these dynamics can also be obtained from a generalization of the standard Wright-Fisher model [11], in which the population is periodically subdivided into separate environments. In this model, the strains in environment i produce a number of gametes proportional to their Wrightian fitness, $W_{\mu,i}$. After a period of growth, $N\beta_i$ gametes are chosen from each environment and mixed together to obtain the next generation. The expected fraction of individuals in the next generation is

$$\langle f_\mu(t + \Delta t) \rangle = \sum_i \beta_i \left[\frac{W_{\mu,i} f_\mu}{\sum_\nu W_{\nu,i} f_\nu} \right]. \quad (2.36)$$

When $\langle f_\mu(t+\Delta t) - f_\mu(t) \rangle$ is small, this update rule has the same continuum limit as Eq. (2.28), with

$$X_\mu = \log \left(\sum_i W_{\mu,i} \right), \quad (2.37)$$

$$\alpha_{\mu,i} = \frac{W_{\mu,i}}{\sum_j W_{\mu,j}}. \quad (2.38)$$

The near-ESS limit

Building on well-known results from population genetics, we expect that the model in Eq. (2.1) will attain its greatest generality in the limit that the intrinsic fitness differences X_μ and the resource-specific mean fitnesses \bar{X}_i are both small compared to one (though still nonzero). This ensures a separation of timescales, in which the population- or ecosystem-level dynamics take place over times that are much longer than a single generation. Previous work has shown that the population-level dynamics in this case become insensitive to many assumptions about the underlying birth-death process [11].

The requirement that X_μ is small is familiar from the standard diffusion limit of the Wright-Fisher model [11]. From the definition of the resource-specific mean fitness in Eq. (2.2), we see that a sufficient condition for \bar{X}_i to be small is that the resource strategies $\alpha_{\mu,i}$ are close to β_i . Since we have previously identified $\alpha_{\mu,i} = \beta_i$ as a marginal evolutionarily stable state (ESS), we have termed this regime the *near-ESS limit*. Alternatively, it can be viewed as a generalization of the standard diffusion limit of population genetics.

We can access the near-ESS limit of Eq. (2.1) in several different ways. We can either work with the full model and take the near-ESS limit in the end, or else we can work directly with the near-ESS limit of Eq. (2.1). We have employed for the former strategy for most of this work. However, in certain cases, it can be more convenient take the near-ESS limit of Eq. (2.1) as our microscopic model (just as it is convenient to work directly with the Wright-Fisher diffusion process).

To obtain the near-ESS limit of Eq. (2.1), we first rewrite the resource uptake strategies in the form

$$\alpha_{\mu,i} = \beta_i (1 + \gamma_{\mu,i}), \quad (2.39)$$

for some rescaled vector $\gamma_{\mu,i}$. The normalization conditions for $\alpha_{\mu,i}$ and β_i yield a corresponding condition for $\gamma_{\mu,i}$,

$$\sum_i \beta_i \gamma_{\mu,i} = 0. \quad (2.40)$$

If we substitute Eq. (2.39) into Eqs. (2.1) and (2.2) and expand to lowest order in X_μ , $\gamma_{\mu,i}$, and $1/N$, we obtain

$$\begin{aligned} \frac{\partial f_\mu}{\partial t} = & \left[X_\mu - \sum_\nu X_\nu f_\nu \right] f_\mu + \frac{\xi_\mu}{\sqrt{N}} \\ & - \sum_{i,\nu} \beta_i \gamma_{\mu,i} \gamma_{\nu,i} f_\nu f_\mu \\ & + \left(\sum_{i,\nu,\sigma} \beta_i \gamma_{\nu,i} \gamma_{\sigma,i} f_\nu f_\sigma \right) f_\mu. \end{aligned} \quad (2.41)$$

The first two terms coincide with the diffusion limit of the Wright-Fisher model, as expected, while the ecological interactions enter at $\mathcal{O}(\gamma^2)$ in the third term. These interactions take the form of a symmetric Lotka-Volterra model with a special interaction matrix formed by the outer product of the resource strategy vectors,

$$A_{\mu,\nu} = \sum_i \beta_i \gamma_{\mu,i} \gamma_{\nu,i}. \quad (2.42)$$

The rank of this matrix is at most \mathcal{R} , regardless of the number of strains. The final term is analogous to the mean fitness term in the Wright-Fisher model, and ensures that $\sum_\mu f_\mu = 1$ at all times. Although this three-body term is formally outside of the Lotka-Volterra model, it is often small in practice, and can be neglected in many calculations.

Fitness differences arising from variable death rates

To illustrate the generality of the near-ESS limit, we show how it can apply to a separate class of consumer-resource models that have been studied in the literature [59, 3], in which the fitness differences arise from differences in the underlying death rate. In other words, we assume that the dilution rate D in Eq. (2.20) can now vary between strains:

$$D \rightarrow D + m_\mu. \quad (2.43)$$

Although this model is formally different than the one we consider in Eq. (2.1), it produces the same limiting behavior as Eq. (2.41) when the fitness differences are small on the timescale of a single generation, i.e., when $m_\mu \ll D$. At lowest order, variation in the death rate will generate effective fitness differences of the form

$$X_\mu \rightarrow X_\mu - \frac{m_\mu}{D}. \quad (2.44)$$

Even when m_μ/D is small, we have shown that these fitness differences can produce dramatic effects when integrated over many generations. Yet at lowest order, the effect of death rate

variation is indistinguishable from the variation in growth rate that we have considered above.

Of course, the limiting behavior in Eq. (2.41) will cease to apply if the fitness differences between strains are no longer small. In this regime, however, the dynamics will often depend on aspects of the birth-death process that are not captured by toy models like Eq. (2.20) [e.g., cell-to-cell variation in growth rate [40], phenotypic delays in mutation penetrance [58], how genetic drift is implemented, etc.]. These would need to be carefully chosen to match the biological system of interest. To ensure the greatest generality, we therefore focus on the aspects of the model that can be captured by Eq. (2.41). Enumeration of other universality classes is an interesting topic for future work.

Competition for two resources

In this section, we derive our main results for the two resource case. The major advantage of this limit is that the multidimensional resource space reduces to the scalar interval $(0, 1)$. Without loss of generality, we will write everything in terms of the first resource component, defining $\beta = \beta_1$ and $\alpha_{\mu,1} = \alpha_\mu$, with the remaining components $\beta_2 = 1 - \beta$ and $\alpha_{\mu,2} = 1 - \alpha_\mu$ fixed by the normalization condition. Following the description in the main text, we will begin by analyzing the competition between two strains, and then consider the effects of adding a third strain to a pair of previously coexisting strains.

Competition between two strains

To analyze the competition between two strains, we let α_1 and α_2 denote the strategy vectors of the two strains, and let $\Delta X = X_2 - X_1$ denote the fitness difference between them. We arbitrarily designate strain 1 as the “wildtype” and consider the frequency of the “mutant strain”, $f \equiv f_2$. With these definitions, the stochastic term in Eq. (2.32) can be written as

$$\begin{aligned}\xi_2(t) &= \sqrt{\frac{(1-f)^2 f}{N}} \eta_2(t) - \sqrt{\frac{f^2 (1-f)}{N}} \eta_1(t), \\ &= \sqrt{\frac{f(1-f)}{N}} \eta(t),\end{aligned}\tag{2.45}$$

where $\eta(t)$ is a third Brownian noise term with $\langle \eta(t) \rangle = 0$ and $\langle \eta(t) \eta(t') \rangle = \delta(t - t')$. Eq. (2.1) can then be rewritten in the familiar population genetic form [11],

$$\frac{\partial f}{\partial t} = s_e(f) f(1-f) + \sqrt{\frac{f(1-f)}{N}} \eta(t),\tag{2.46}$$

where the effective frequency-dependent selection coefficient, $s_e(f)$, is obtained from the deterministic portion of Eq. (2.1),

$$\begin{aligned}s_e(f) &\equiv \frac{1}{f(1-f)} \left(\frac{\partial f_2}{\partial t} \right)_{\text{deterministic}}, \\ &= \frac{\beta [\alpha_2 e^{\Delta x} - \alpha_1]}{\alpha_1 + [\alpha_2 e^{\Delta x} - \alpha_1] f} \\ &\quad + \frac{(1-\beta) [(1-\alpha_2) e^{\Delta X} - (1-\alpha_1)]}{1 - \alpha_1 + [(1-\alpha_2) e^{\Delta X} - (1-\alpha_1)] f}.\end{aligned}\tag{2.47}$$

Our main results can be derived from limiting versions of this basic model.

Invasion of a new strain

The invasion of a new strain corresponds to the $f \rightarrow 0$ limit, in which Eq. (2.46) reduces to the linearized form,

$$\frac{\partial f}{\partial t} = S_{\text{inv}} f + \sqrt{\frac{f}{N}} \eta(t), \quad (2.48)$$

with an invasion fitness S_{inv} defined by

$$\begin{aligned} S_{\text{inv}} &\equiv \lim_{f \rightarrow 0} s_e(f), \\ &= (e^{\Delta X} - 1) + e^{\Delta X} \left[\frac{(\beta - \alpha_1)(\alpha_2 - \alpha_1)}{\alpha_1(1 - \alpha_1)} \right]. \end{aligned} \quad (2.49)$$

Eq. (2.48) can be solved using standard methods [35]. We will simply quote the relevant results here, while a more pedagogical exposition can be found in Chapter 1 of Ref. [25].

For initial frequencies small compared to the $1/NS_{\text{inv}}$, the genetic drift term dominates, and there is a high probability that the mutant will drift to extinction. However, with probability $p_{\text{est}} = 2NS_{\text{inv}}f(0)$, the mutant will drift to frequency $\sim 1/NS_{\text{inv}}$, after which point the selection term dominates over genetic drift. This “established” lineage will then grow deterministically as $f(t) = \frac{1}{2NS_{\text{inv}}} e^{S_{\text{inv}}t}$, which can be matched onto the full nonlinear (but deterministic) solution as f increases further. The full solution is somewhat unwieldy, but the first-order nature of the ODE shows that $f(t)$ cannot decrease as $t \rightarrow \infty$. Thus, once the mutant establishes, the deterministic dynamics will never drive the mutant close enough to the drift barrier that extinction becomes likely again. This suggests that the branching process description will be valid as long as $f(t)$ remains sufficiently small during the duration of the establishment process that $f(t) \ll 1$ and $s_e(f) \approx S_{\text{inv}}$. This will be true provided that these conditions are satisfied at the drift barrier, $1/NS_{\text{inv}}$, which leads to the conditions

$$NS_{\text{inv}} \gg 1, \quad (2.50)$$

$$\frac{NS_{\text{inv}}\alpha_1}{\alpha_2 e^{\Delta X} - \alpha_1} \gg 1, \quad (2.51)$$

$$\frac{NS_{\text{inv}}(1 - \alpha_1)}{(1 - \alpha_2)e^{\Delta X} - (1 - \alpha_1)} \gg 1. \quad (2.52)$$

These conditions can be satisfied simultaneously for sufficiently large N .

Stable coexistence

If $S_{\text{inv}} > 0$ and the mutant is lucky enough to establish, then the frequency-dependent selection term will either drive the mutant to fixation ($f = 1$) or else stabilize at some intermediate frequency f^* . As described in the main text, stable coexistence requires that

the reciprocal invasion fitness,

$$\begin{aligned} S_{\text{inv}}^R &\equiv \lim_{f \rightarrow 1} -s_e(f), \\ &= (e^{-\Delta X} - 1) + e^{-\Delta X} \left[\frac{(\beta - \alpha_2)(\alpha_1 - \alpha_2)}{\alpha_2(1 - \alpha_2)} \right], \end{aligned} \quad (2.53)$$

is also positive. Solving this equation when $S_{\text{inv}}^R = 0$ yields the critical fitness threshold

$$\Delta X_{\text{max}} = \log \left(1 + \frac{(\alpha_1 - \alpha_2)(\beta - \alpha_2)}{\alpha_2(1 - \alpha_2)} \right), \quad (2.54)$$

which reduces to Eq. (2.9) in the main text in the near-ESS limit. We might naively assume that this threshold would be equivalent to the fitness that gives strain 2 a higher uptake rate on *both* individual resources, i.e. $\alpha_2 e^{\Delta X} \geq \alpha_1$ and $(1 - \alpha_2)e^{\Delta X} \geq 1 - \alpha_1$. Although this is indeed a sufficient condition for strain 2 to fix, the true thresholds in Eqs. (2.9) and (2.54) are much weaker conditions, which depend on the environmental supply vector β . This means that in practice, stable coexistence will be disrupted long before one of the strains is uniformly better than the other.

When the conditions for stable coexistence are met, the equilibrium frequency f^* is obtained from the condition that $s_e(f^*) = 0$. From Eq. (2.47), we see that this can only happen if $\alpha_2 e^{\Delta X} - \alpha_1$ and $(1 - \alpha_2)e^{\Delta X} - (1 - \alpha_1)$ have different signs, i.e. neither strain is uniformly better than the other. Solving for f^* , we find that

$$f^* = \frac{f_0^* + \left[f_0^* + \frac{\alpha_1(1 - \alpha_1)}{\Delta \alpha} \right] (e^{\Delta X} - 1)}{\left[1 + \frac{\alpha_2}{\Delta \alpha} (e^{\Delta X} - 1) \right] \left[1 - \frac{(1 - \alpha_2)}{\Delta \alpha} (e^{\Delta X} - 1) \right]}, \quad (2.55)$$

where $f_0^* = (\beta - \alpha_1)/\Delta \alpha$ is the equilibrium frequency in the absence of any fitness differences. When $f = f^*$, the resource-specific mean fitnesses \bar{X}_i take on the values

$$\begin{aligned} \bar{X}_1 &= -\log \left[1 - (1 - e^{-\Delta X}) \left(\frac{1 - \alpha_2}{\Delta \alpha} \right) \right], \\ \bar{X}_2 &= -\log \left[1 + (1 - e^{-\Delta X}) \left(\frac{\alpha_2}{\Delta \alpha} \right) \right], \end{aligned} \quad (2.56)$$

which are independent of the resource supply vector β . This extends the “environmental shielding” behavior derived in the neutral limit by Ref. [51]: when two strains coexist on two substitutable resources, the strain frequencies evolve so that the remaining selection pressures take on values that are independent of the environment, and depend only on the identities of the coexisting strains. We will revisit this behavior again in the multi-resource case below.

In the limit that fitness differences are small [specifically, when ΔX is small compared to 1, $\alpha_2/\Delta \alpha$, and $(1 - \alpha_2)/\Delta \alpha$], Eq. (2.56) reduces to the linearized version,

$$\bar{X}_1 = \frac{(1 - \alpha_2)}{\Delta \alpha} \Delta X, \quad \bar{X}_2 = -\frac{\alpha_2}{\Delta \alpha} \Delta X, \quad (2.57)$$

while Eq. (2.55) reduces to the linear relation quoted in Eq. (2.10) in the main text. This defines a second fitness scale,

$$\begin{aligned} X_f &\equiv f^*(1 - f^*) \left(\frac{\partial f^*}{\partial \Delta X} \right)^{-1}, \\ &= \frac{(\beta - \alpha_1)(\alpha_2 - \beta)}{\beta(1 - \beta) + (\beta - \alpha_1)(\alpha_2 - \beta)}, \end{aligned} \quad (2.58)$$

over which $f^*(\Delta X)$ changes significantly. Note that X_f has approximately the same scaling behavior for small and large $\beta - \alpha$ as the critical threshold ΔX_{\max} in Eq. (2.9).

For frequencies close to f^* , the selection term again grows small compared to the genetic drift term. Linearizing Eq. (2.46) around $f \approx f^*$, the fluctuations $\delta f = f - f^*$ are described by

$$\frac{\partial \delta f}{\partial t} = -X_{\text{eq}} f^*(1 - f^*) \delta f + \sqrt{\frac{f^*(1 - f^*)}{N}} \eta(t), \quad (2.59)$$

where we have defined the equilibrium restoring fitness

$$X_{\text{eq}} \equiv - \left. \frac{\partial s_e(f)}{\partial f} \right|_{f=f^*} = \frac{1}{\beta(1 - \beta)} \left(\frac{(\alpha_2 e^{\Delta X} - \alpha_1)[(1 - \alpha_2)e^{\Delta X} - (1 - \alpha_1)]}{\alpha_1[(1 - \alpha_2)e^{\Delta X} - (1 - \alpha_1)] - (1 - \alpha_1)[\alpha_2 e^{\Delta X} - \alpha_1]} \right)^2. \quad (2.60)$$

In the limit that $\Delta X \ll 1$, this becomes

$$X_{\text{eq}} = \frac{\Delta \alpha^2}{\beta(1 - \beta)} \left[1 + \frac{2\alpha_2 - 1}{\Delta \alpha} \Delta X \right]. \quad (2.61)$$

This model can again be solved using standard methods [33]. The stationary distribution of δf tends toward a normal distribution with mean zero and standard deviation $\sigma_f = 1/\sqrt{2NX_{\text{eq}}}$, which decays on a timescale $\sim 1/X_{\text{eq}}f^*(1 - f^*)$. The quasi-deterministic model is therefore self-consistent provided that

$$\frac{\sigma_f}{f^*(1 - f^*)} = \sqrt{\frac{\Delta \alpha^2}{2N(\beta - \alpha_1)^2(\alpha_2 - \beta)^2}} \ll 1, \quad (2.62)$$

which can be satisfied for sufficiently large N .

The fluctuations in f lead to similar fluctuations in the resource-specific mean fitnesses, \bar{X}_i , whose first order contribution is given by

$$\begin{aligned} \delta \bar{X}_1 &= \frac{\Delta \alpha}{\beta} \left[1 + \frac{\alpha_2}{\Delta \alpha} (e^{\Delta X} - 1) \right] e^{-\bar{X}_1} \delta f, \\ \delta \bar{X}_2 &= -\frac{\Delta \alpha}{1 - \beta} \left[1 - \frac{(1 - \alpha_2)}{\Delta \alpha} (e^{\Delta X} - 1) \right] e^{-\bar{X}_2} \delta f. \end{aligned} \quad (2.63)$$

Competition between three strains

Having characterized the dynamics for a pair of strains, we next consider a scenario in which a third strain is introduced into a stable ecosystem where a pair of strains already coexist. Without loss of generality, we will assume that the third strain is a mutant version of the second strain, with fitness $X_3 = \Delta X + s$ and strategy vector α_3 . From the definition of the model in Eq. (2.1), this mutant strain will have an invasion fitness

$$S_{\text{inv}} = \alpha_3 \left(e^{\Delta X + s - \bar{X}_1} - 1 \right) + (1 - \alpha_3) \left(e^{\Delta X + s - \bar{X}_2} \right) \quad (2.64)$$

where the resource-specific mean fitnesses $\bar{X}_1(t)$ and $\bar{X}_2(t)$ are dictated by the two strain process in Eq. (2.46). If the mutant was actually identical to its parent strain (i.e., if $\alpha_3 = \alpha_2$ and $s = 0$), it should never be favored to invade, since it can at best compete neutrally with its parent. This implies that

$$\alpha_2 \left(e^{\Delta X - \bar{X}_1} - 1 \right) + (1 - \alpha_2) \left(e^{\Delta X - \bar{X}_2} - 1 \right) = 0, \quad (2.65)$$

when averaged over the timescales required for the mutation to invade. Multiplying this expression by e^s and subtracting it from Eq. (2.64), we can then rewrite the general invasion fitness in the form

$$S_{\text{inv}} = (e^s - 1) + (\alpha_3 - \alpha_2) \left(e^{-\bar{X}_1(t)} - e^{-\bar{X}_2(t)} \right) e^{\Delta X + s}, \quad (2.66)$$

We consider the implications of this expression in various special cases below.

No fitness differences

In a completely neutral scenario ($\Delta X = s = 0$), the resource-specific mean fitnesses are solely determined by the fluctuations $\delta \bar{X}_1$ and $\delta \bar{X}_2$ from Eq. (2.63), and Eq. (2.66) reduces to

$$S_{\text{inv}} = (\alpha_3 - \alpha_2) [\delta \bar{X}_2 - \delta \bar{X}_1] = \frac{(\alpha_2 - \alpha_3)(\alpha_2 - \alpha_1)}{\beta(1 - \beta)} \delta f(t). \quad (2.67)$$

Since $\langle \delta f(t) \rangle = 0$, this agrees with the deterministic results of Ref. [51], who found that all further invasion fitnesses vanish in a neutral population when the ecosystem is fully exploited. However, our stochastic analysis shows that fluctuations can induce momentary selection pressures of order

$$\delta S_{\text{inv}} \sim \frac{(\alpha_2 - \alpha_3)(\alpha_2 - \alpha_1)}{\beta(1 - \beta)} \frac{1}{\sqrt{N X_{\text{eq}}}}, \quad (2.68)$$

which can be large compared to $1/N$. However, these momentary selection pressures average out to zero over a timescale $1/X_{\text{eq}} f^*(1 - f^*)$. When N is large, this is much shorter than the timescale $\sim 1/\delta S_{\text{inv}}$ required for the mutant lineage to escape the drift barrier. This shows that internal fluctuations cannot induce anomalous establishment events in our model. To leading order in N , ecological selection pressures vanish in a neutral population when two strains coexist on two substitutable resources.

Pure fitness mutations

In the case where the mutant lineage is created by a pure fitness mutation, $\alpha_3 = \alpha_2$, and the invasion fitness reduces to

$$S_{\text{inv}} = e^s - 1 \approx s, \quad (2.69)$$

which is identical to the standard Wright-Fisher model. This justifies our interpretation of X_μ as a fitness parameter. Eq. (2.69) is a slightly stronger result, since it implies that pure fitness mutations continue to establish at the same rate, regardless of the structure of the ecosystem. When such a mutation establishes, it is guaranteed to displace its parent strain, resulting in a two-strain competition between strain 3 and strain 1, which now differ in fitness by an amount $\Delta X + s$. If $\Delta X + s \geq \Delta X_{\text{max}}$ from Eq. (2.9), then stable coexistence will be disrupted, and strain 3 will take over the entire population. On the other hand, if $\Delta X + s < \Delta X_{\text{max}}$, the mutant will only displace its parent strain, and will be prevented from sweeping through the entire population. Instead, the successful mutation will shift the equilibrium frequency by an amount

$$\begin{aligned} \Delta f &= f^*(\Delta X + s) - f^*(\Delta X), \\ &\approx \frac{\beta(1 - \beta) + (\beta - \alpha_1)(\beta - \alpha_2)}{\Delta \alpha^2} \cdot s, \end{aligned} \quad (2.70)$$

where we have employed the linearized approximation for f^* from Eq. (2.10) in the main text.

Pure strategy mutations

If the mutant lineage is created by a pure strategy mutation ($s = 0$), then the invasion fitness reduces to

$$S_{\text{inv}} = \frac{\alpha_3 - \alpha_2}{\alpha_1 - \alpha_2} (e^{\Delta X} - 1), \quad (2.71)$$

where we have retained only the leading order contribution as $N \rightarrow \infty$. The $\Delta X \ll 1$ limit is listed in Eq. (2.15) in the main text. The interpretation of this expression, and the various scenarios that can arise after establishment, are described in the main text as well.

Evolution of a single-locus ecology

The results above allow us to analyze the effects of further evolution in our consumer resource model. As a first pass, we focus on a simplified scenario, in which strategy mutations switch between two fixed strategy vectors, α_1 and α_2 , and occur at rate U_α . We assume that α_1 and α_2 span β , so that the strains can stably coexist. We also assume that α_1 and α_2 are sufficiently close to β that we can invoke the near-ESS limits of various expressions above. We note that while this assumption is also employed in the adaptive dynamics literature

[21, 10], our model also differs from these results in a key way, as it includes α that go beyond the infinitesimal evolution assumption in adaptive dynamics.

Our model also differs from the canonical adaptive dynamics scenario in that it includes pure fitness mutations, which occur at rate $U_X \rho_X(s)$. We assume that the tails of $\rho_X(s)$ are sufficiently light that the distribution can be approximated by a characteristic beneficial fitness effect [28], which we will also denote by the generic variable s below. Our analysis here will focus on the strong-selection weak mutation (SSWM) regime that arises in the limit that $N \rightarrow \infty$ and $U_\alpha + U_X \rightarrow 0$. The first assumption guarantees that genetic drift is only relevant when mutations are sufficiently rare, so that the establishment process can be modeled by the branching process techniques above. The second assumption guarantees that all mutations establish or go extinct before the next mutation occurs, so that they can be described by the two- and three-strain competition processes above. Violations of this assumption are considered in more detail in a following section.

No strategy mutations

We first consider the dynamics under pure fitness mutations when $U_\alpha/U_X = 0$. We assume that the population has just diversified into a pair of coexisting strains with fitness difference $\Delta X = 0$, and equilibrium frequency f_0^* . Pure fitness mutations will occur in each clade at rate $NU_X f_0^*$ and $NU_X(1 - f_0^*)$, respectively. According to Eq. (2.69), these establish with probability $p_{\text{est}} = 2s$, sweep through their parent clade, and result in a new fitness differential,

$$\Delta X = \begin{cases} +s & \text{with probability } f_0^*, \\ -s & \text{with probability } 1 - f_0^*, \end{cases} \quad (2.72)$$

which depends on the genetic background in which the mutation arose. This fitness differential will lead to a shift in the equilibrium frequency $\Delta f = \pm s/s_c$ described by Eq. (2.11) in the main text. If $\Delta f < -f_0^*$ or $\Delta f > 1 - f_0^*$, then by definition stable coexistence will be disrupted, since the new frequency would fall outside the interval from 0 to 1. [One can also see this directly from the fitness bounds ΔX_{\min} and ΔX_{\max} in Eqs. (2.8) and (2.9).] Since the mutant has already swept through its parent clade, a disruption of coexistence implies that it will take over the entire population. In the near-ESS limit, we can combine these two conditions to obtain a convenient asymptotic condition for s :

$$s \gg s_c f_0^*(1 - f_0^*) \equiv \frac{(\alpha_1 - \beta)(\beta - \alpha_2)}{\beta(1 - \beta)}. \quad (2.73)$$

In this regime, the lifetime of coexistence is of order $\tau_{\text{collapse}} \sim 1/NU_X s$ (the time that it takes for one fitness mutation to occur).

In the opposite regime, when $s \ll s_c f_0^*(1 - f_0^*)$, individual fitness mutations lead to small shifts in f_0^* , and many such mutations must accumulate before stable coexistence is disrupted. In this case, we can model the changing equilibrium frequency using an effective diffusion process. In an interval of time δt , the fitness differential changes by $\delta \Delta X = s(k_2 - k_1)$,

where k_2 and k_1 denote the number of fitness mutations that accumulate in the f^* and $1 - f^*$ backgrounds, respectively. In the weak mutation limit, these occur as a Poisson process with rates $2NU_X s f^* \delta t$ and $2NU_X s (1 - f^*) \delta t$, respectively, so that

$$\langle k_2 - k_1 \rangle = 2NU_X s (2f^* - 1) \delta t, \quad (2.74)$$

$$\text{Var}(k_2 - k_1) = 2NU_X s \delta t. \quad (2.75)$$

The fitness difference ΔX can therefore be described by an effective diffusion process,

$$\frac{\partial \Delta X}{\partial t} = 2NU_X s^2 [2f^*(\Delta X) - 1] + \sqrt{2NU_X s^3} \eta(t), \quad (2.76)$$

where the equilibrium frequency f^* itself depends on ΔX through Eq. (2.10) in the main text. Changing variables from ΔX to f^* , we obtain Eq. (2.12) in the main text. For our detailed calculations, it will be somewhat more convenient to work with the rescaled variables $Y = 2f^* - 1$ and $k = 2NU_X s t$, which yields a related equation

$$\frac{\partial Y}{\partial k} = \frac{2s}{s_c} Y + \sqrt{\frac{4s^2}{s_c^2}} \cdot \eta(\tau). \quad (2.77)$$

This is similar to the equation for the drift-induced fluctuations in Eq. (2.59), except that the bias is now a destabilizing force, rather than a stabilizing one. This reflects the fact that larger clades are more likely to acquire beneficial mutations in the weak mutation limit, which leads to further increases in frequency. We can quantify the strength of this snowballing effect by analyzing the ultimate fixation probability of strain 1 (i.e., the probability that $Y \rightarrow 1$) as a function of the current value of $Y = 2f^* - 1$. Eq. (2.77) implies a corresponding backward equation for the fixation probability

$$\frac{2s}{s_c} Y \frac{\partial P_{\text{fix}}}{\partial Y} + \frac{2s^2}{s_c^2} \frac{\partial P_{\text{fix}}}{\partial Y} = 0, \quad (2.78)$$

whose solution is given by

$$P_{\text{fix}}(f^*) \approx \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{2f^*-1}{\sqrt{s/s_c}}} e^{-\frac{u^2}{2}} du. \quad (2.79)$$

This function undergoes a sharp transition near $2f^* - 1 \sim \sqrt{s/s_c}$. When $|2f^* - 1| \ll \sqrt{s/s_c}$, fixation and extinction of the clade are equally likely, while for $2f^* - 1 \gg \sqrt{s/s_c}$, fixation is virtually guaranteed. This transition has a simple interpretation in terms of the relative strengths of the bias and noise terms in Eq. (2.12): $\sqrt{s/s_c}$ represents a critical frequency difference above which the bias dominates over the noise term. Since $\sqrt{s/s_c}$ is itself a small parameter in the $s \ll s_c$ regime, this implies that the random portion of the clade competition process is confined to frequencies near 50%. Reversals from frequencies near $f^* \approx 0$ or $f^* \approx 1$ are asymptotically unlikely.

To investigate the dynamics of this process, we analyze the mean squared frequency difference $\langle Y^2 \rangle$. Using Eq. (2.77), we can derive a closed moment equation for $\langle Y^2 \rangle$

$$\frac{\partial \langle Y^2 \rangle}{\partial k} = \frac{4s}{s_c} \langle Y^2 \rangle + \frac{4s^2}{s_c^2}, \quad (2.80)$$

whose solution is given by

$$\langle Y(k)^2 \rangle = Y(0)^2 e^{\frac{4sk}{s_c}} + \frac{s}{s_c} \left(e^{\frac{4sk}{s_c}} - 1 \right). \quad (2.81)$$

Solving for k and converting back to units of time, we find that

$$t = \frac{s_c}{8NU_X s^2} \log \left(\frac{\langle Y(t)^2 \rangle + s/s_c}{Y(0)^2 + s/s_c} \right). \quad (2.82)$$

The behavior of this function has a simple heuristic interpretation based on the fundamental timescales of Eq. (2.12). These heuristics follow from standard arguments [7, 14, 6], so we will simply quote the relevant results here, while referring the interested reader to Chapter 1 of Ref. [25] for a more detailed exposition.

Starting from $|2f_0^* - 1| \ll \sqrt{s/s_c}$, the clade frequencies will wander diffusively for a time $\tau_{\text{drift}} \sim \frac{s_c}{NU_X s^2}$ until the frequency difference reaches $\sqrt{s/s_c}$, after which point the major clade deterministically acquires mutations for $\tau_{\text{collapse}} \sim \frac{s_c}{NU_X s^2} \log(s_c/s)$ more generations until it reaches fixation. On the other hand, if the clades start with a frequency difference $|2f_0 - 1| \gg \sqrt{s/s_c}$, then the major clade will deterministically fix within $\sim \frac{s_c}{NU_X s^2} \log \left(\frac{1}{|2f_0^* - 1|} \right)$ generations

Including strategy mutations

We can use the results above to analyze the case where $U_\alpha/U_X > 0$. For very low values of U_α/U_X , strategy mutations will rarely occur before the ecosystem collapses according to the process described above. In this case, the main effect of strategy mutations is to re-diversify a population that consists of a single ecotype. The invasion fitness of such a mutation is therefore given by Eq. (2.4) in the main text, and will vary depending on which ecotype dominates the population.

We can therefore distinguish between two regimes. If $|2f_0^* - 1| \ll \sqrt{s/s_c}$, then both ecotypes are equally likely to fix, and the average invasion fitness is

$$\bar{X}_{\text{inv}} = \frac{(\beta - \alpha_1)(\alpha_2 - \alpha_1)}{\beta(1 - \beta)} + \frac{(\beta - \alpha_2)(\alpha_1 - \alpha_2)}{\beta(1 - \beta)} = s_c, \quad (2.83)$$

This leads to a diversification timescale $\tau_{\text{diversify}} \sim 1/NU_\alpha s_c$, and the diversification-selection balance in Eq. (2.13) in the main text. On the other hand, if $|2f_0^* - 1| \gg \sqrt{s/s_c}$ then the clade with the larger initial frequency will typically be the one that fixes. Without loss of

generality, we will relabel the strains so that f^* always represents this clade. In this scenario, the average invasion fitness is instead given by $\bar{X}_{\text{inv}} \sim s_c(1 - f_0^*)$, which is strictly smaller than s_c . In this case, the diversification-selection balance is given by

$$\frac{\Pr[\mathcal{S} = 2]}{\Pr[\mathcal{S} = 1]} \approx \frac{\tau_{\text{collapse}}}{\tau_{\text{diversify}}} \sim \frac{U_\alpha}{U_X} \left(\frac{s_c}{s}\right)^2 (1 - f_0^*) \log\left(\frac{1}{2f_0^* - 1}\right). \quad (2.84)$$

For still larger values of U_α/U_X , strategy mutations will start to occur before one of the clades has fixed in the population. If the mutation occurs in the fitter clade, it will have an invasion fitness $S_{\text{inv}} = |\Delta X|$, and will reset the fitness difference to zero if it establishes. On the other hand, if the mutation occurs in the less fit clade, it will have a negative invasion fitness and will not be able to establish. Thus, the net effect of these strategy mutations is to set $\Delta X = 0$ at a time-dependent rate

$$\lambda_0(t) = 2NU_\alpha f^* \Delta X \theta(\Delta X) - 2NU_\alpha (1 - f^*) \Delta X \theta(-\Delta X), \quad (2.85)$$

where $\theta(z)$ is the Heaviside step function, and $f^*(t)$ and $\Delta X(t)$ are determined by the effective diffusion process in Eq. (2.12) in the main text. The first successful strategy mutation will occur on a characteristic invasion timescale determined by the implicit relation

$$\int_0^{\tau_{\text{invade}}} \lambda(t) dt \sim 1. \quad (2.86)$$

Since the fitter strain will typically be the most abundant as well, Eq. (2.85) will only differ by a factor of two from the much simpler expression

$$\lambda_0(t) \sim NU_\alpha |\Delta X|. \quad (2.87)$$

Since Eq. (2.86) is only accurate up to an order one factor, we will use this simpler approximation for $\lambda_0(t)$ instead.

Based on these definitions, we can obtain a self-consistent solution to Eq. (2.86) in various regimes. If $U_\alpha \gg U_X$, then strategy mutations will arise much faster than individual mutations. In this case, a lucky fitness mutation will establish in one of the clades after a time of order $1/NU_X s$, so that

$$\int_0^t \lambda_0(t') dt' \sim NU_\alpha s \left(t - \frac{1}{NU_X s}\right). \quad (2.88)$$

This yields an invasion timescale

$$\tau_{\text{invade}} \sim \frac{1}{NU_\alpha s} + \frac{1}{NU_X s} \sim \frac{1}{NU_X s}, \quad (2.89)$$

which is self-consistent provided that $U_\alpha \gg U_X$.

If $\tau_{\text{invade}} \gg 1/NU_X s$, then multiple fitness mutations will accumulate before the first successful strategy mutation arises. If $\tau_{\text{invade}} \ll \tau_{\text{drift}}$ then the fitness differential ΔX wanders diffusively as $|\Delta X| \sim \sqrt{NU_X s^3 t}$, and

$$\int_0^t \lambda(t') dt' \sim NU_\alpha \sqrt{NU_X} (st)^{3/2}. \quad (2.90)$$

This leads to an invasion timescale of order

$$\tau_{\text{invade}} \sim \frac{1}{NU_X s} \left(\frac{U_X}{U_\alpha} \right)^{2/3}, \quad (2.91)$$

which is self consistent provided that $U_\alpha \ll U_X \ll U_\alpha (s_c/s)^{3/2}$.

If $\tau_{\text{invade}} \gg \tau_{\text{drift}}$, or if the initial frequency differential already exceeds the critical value $\sqrt{s/s_c}$, then the successful strategy mutation will occur when ΔX is growing deterministically as

$$\Delta X \sim s_c \sqrt{(2f_0^* - 1)^2 + \frac{s}{s_c} \cdot e^{\frac{4NU_X s^2 t}{s_c}}}, \quad (2.92)$$

so that

$$\int_0^t \lambda_0(t') dt' \sim \frac{U_\alpha s_c^2}{U_X s^2} \sqrt{(2f_0^* - 1)^2 + \frac{s}{s_c} \cdot e^{\frac{4NU_X s^2 t}{s_c}}}. \quad (2.93)$$

If $\tau_{\text{invade}} \ll \tau_{\text{collapse}}$, this leads to an invasion timescale,

$$\tau_{\text{invade}} \sim \frac{s_c}{8NU_X s^2} \log \left(\frac{U_X^2 s^4}{U_\alpha^2 s_c^4 (2f_0^* - 1)^2 + \frac{s}{s_c}} \right), \quad (2.94)$$

which will be self-consistent provided that $U_\alpha \left(\frac{s_c}{s} \right)^{3/2} \ll U_X \ll U_\alpha \left(\frac{s_c}{s} \right)^2$. Finally, for $U_X \gg U_\alpha \left(\frac{s_c}{s} \right)^2$, strategy mutants are sufficiently rare that the ecosystem will typically collapse and re-diversify before invasion can occur. In this case, τ_{invade} formally diverges. The various regimes for τ_{invade} are summarized in Eq. (2.14) in the main text.

The effects of clonal interference

In our analysis above, we have focused on the weak mutation limit, in which only two or three strains exist within the population at any one time. While this enabled many analytical simplifications, it is also known that many microbial populations lie outside this regime. This is particularly true for many microbial evolution experiments in which stable coexistence has been observed to evolve spontaneously. While a thorough analysis of this regime is beyond the scope of the present work, we will summarize the key differences that are likely to arise in the effective diffusion process in Eq. (2.12) in the main text.

Outside of the weak-mutation limit, many established beneficial mutations will be driven to extinction due to clonal interference with other beneficial mutations that happen to segregate at the same time [22]. In the limit that clonal interference is strong ($NU_X \gg 1$), this has two main consequences. First, the rate of adaptive substitution scales much more weakly with N than the linear expectation $NU_X s$ from the SSWM limit. In the case of coexisting strains, this will also apply to the subpopulations Nf^* and $N(1 - f^*)$ that correspond to the two clades. As a result, the bias term in Eq. (2.12) will be significantly reduced (and essentially vanishes in the limit of strong clonal interference). Second, clonal interference causes the rate of adaptation to become more deterministic in addition to reducing it, since it is no longer limited by the supply of beneficial mutations. The dynamics of these fluctuations are poorly understood in the general case, though Ref. [15] has shown that they lead to a long-term diffusion constant,

$$D_X = \left[\frac{s}{\log\left(\frac{s}{U_X}\right)} \right]^3. \quad (2.95)$$

for the total fitness gain in a model similar to ours. Thus, as long as $Nf^*(1 - f^*)$ remains sufficiently large that clonal interference within each clade remains strong, we expect the effective diffusion model in Eq. (2.12) to be better approximated by the limiting form

$$s_c \frac{\partial f}{\partial t} \sim \sqrt{D_X} \cdot \eta(t). \quad (2.96)$$

Due to the weaker bias term, we expect that the relative frequencies of the clades can undergo dramatic reversals before one or the other accumulates a fitness advantage that is large enough it to fix. Interestingly, such reversals have been observed in a long-term experiment in *E. coli* [29]. However, a more thorough analysis of this clonal interference regime remains an interesting avenue for future work.

Competition for many resources

In this section, we show how many of the results derived in the two-resource case can be extended to systems with larger numbers of resources. Most of these results will apply for arbitrary values of \mathcal{R} , but we are particularly interested in the qualitative differences that arise in the many resource limit where $\mathcal{R} \gg 1$.

Invasion of a mutant strain

We begin by considering a mutation that occurs in an ecosystem with an arbitrary number of coexisting strains, with equilibrium resource-specific mean fitnesses, \bar{X}_i . Without loss of generality, we will assume that the mutation occurs in the $\mu = 1$ strain, and leads to a new phenotype $(X_\mu + s, \vec{\alpha}_\mu + \Delta\vec{\alpha})$, where the strategy perturbation must satisfy the normalization constraint $\sum_i \Delta\alpha_i = 0$. The invasion fitness for the resident strain μ must be zero, since it is by definition present at the ecological equilibrium. Using this fact, along with the normalization condition on $\Delta\vec{\alpha}$, one can show that the general invasion fitness for the mutant is given by

$$S_{\text{inv}} \equiv \sum_i (\alpha_{\mu,i} + \Delta\alpha_i) \left[e^{X_\mu + s - \bar{X}_i} - 1 \right], \quad (2.97)$$

$$= (e^s - 1) + e^s \sum_i \Delta\alpha_i \left(e^{X_\mu - \bar{X}_i} - 1 \right), \quad (2.98)$$

which generalizes the two-resource invasion fitness in Eq. (2.66). In the near-ESS limit where s , X_μ , and \bar{X}_i are all small compared to one, this expression reduces to Eq. (2.16) in the main text.

Ecological equilibria

The invasion fitness in Eq. (2.97) depend on the structure of the stable ecosystem through the resource-specific mean fitnesses, \bar{X}_i , which depend on the equilibrium strain frequencies f_μ^* through the definition in Eq. (2.2). Compared to the two-resource case above, it is generally more difficult to calculate the ecological equilibrium for a set of strains when $\mathcal{R} > 2$. Part of this difficulty is caused by the vector nature of the resource space, which can no longer be projected down onto a single scalar dimension. However, this is more than just a book-keeping issue — there are also fundamentally new kinds of ecological equilibria that can arise when $\mathcal{R} > 2$. In a two-resource system, ecological equilibria are either monocultures (with $\mathcal{S} = 1$ resident strains), or else contain the maximum number of coexisting strains permitted by the environment ($\mathcal{S} = 2$). However, when $\mathcal{S} > 2$, one can also have stable coexistence at any intermediate value of $1 < \mathcal{S} < \mathcal{R}$, in addition to the saturated state with $\mathcal{S} = \mathcal{R}$. These two classes of equilibria turn out to have very different properties.

Saturated ecosystems. The saturated stable state ($\mathcal{S} = \mathcal{R}$) is the closest analogue of the two-resource equilibrium that we studied in Appendix 2.6. In this case, we can obtain

an explicit solution for the strain frequencies, f_μ^* , and resource-specific mean fitnesses, \bar{X}_i , attained at equilibrium as a function of the phenotypes $(X_\mu, \vec{\alpha}_\mu)$ of the resident strains. By definition, the per capita growth rate $(\partial_t \log f_\mu)$ of each resident strain must vanish at equilibrium, which yields a system of \mathcal{S} equations for the \mathcal{R} resource-specific mean fitnesses:

$$\sum_i \alpha_{\mu,i} e^{-\bar{X}_i} = e^{-X_\mu}. \quad (2.99)$$

When $k = p$, this system can be inverted to obtain

$$e^{-\bar{X}_i} = \sum_\mu \alpha_{i,\mu}^{-1} e^{-X_\mu}, \quad (2.100)$$

where $\alpha_{i,\mu}^{-1}$ is the left inverse of $\alpha_{\mu,i}$. In the limit that $|X_\mu - X_\nu| \ll 1$, this reduces to Eq. (2.17) in the main text. Using the definition of \bar{X}_i in Eq. (2.2) in the main text, we obtain a second system of \mathcal{R} equations for the k equilibrium frequencies:

$$\beta_i = \sum_\mu \alpha_{\mu,i} e^{X_\mu - \bar{X}_i} f_\mu^*, \quad (2.101)$$

which is the non-neutral generalization of Eq. (2.6) in the main text. Again, when $\mathcal{S} = \mathcal{R}$, we can invert this system to obtain

$$f_\mu^* = e^{-X_\mu} \sum_i \beta_i e^{\bar{X}_i} \alpha_{i,\mu}^{-1} = \sum_i \frac{\beta_i \alpha_{i,\mu}^{-1}}{\sum_\nu \alpha_{i,\nu}^{-1} e^{X_\mu - X_\nu}}, \quad (2.102)$$

since the left and right inverses are equal in this case. In the limit that $|X_\mu - X_\nu| \rightarrow 0$, we obtain the leading order contribution

$$f_\mu^* \approx \sum_i \beta_i \alpha_{i,\mu}^{-1} - \sum_{i,\nu} \beta_i \alpha_{i,\mu}^{-1} \alpha_{i,\nu}^{-1} (X_\mu - X_\nu). \quad (2.103)$$

To gain intuition into these formulae, we consider a set of strains whose resource strategies are a mixture of specialist and generalist components:

$$\alpha_{\mu,i} = (1 - \epsilon) \beta_i + \epsilon \delta_{\mu,i}, \quad (2.104)$$

where $0 \leq \epsilon \leq 1$ provides a measure of the “distance” between the resource strategies. In this case, the inverse matrix has the asymptotic limits

$$\alpha_{i,\mu}^{-1} \sim \begin{cases} \delta_{i,\mu} + \mathcal{O}(1 - \epsilon) & \text{if } 1 - \epsilon \ll 1, \\ \frac{\delta_{i,\mu} - \beta_\mu}{\epsilon} + \mathcal{O}(1) & \text{if } \epsilon \ll 1, \end{cases} \quad (2.105)$$

so that

$$\bar{X}_i \sim \begin{cases} X_i & \text{if } 1 - \epsilon \ll 1, \\ \frac{\sum_{j \neq i} \beta_j (X_i - X_j)}{\epsilon} & \text{if } \epsilon \ll 1, \end{cases} \quad (2.106)$$

and

$$f_\mu^* \sim \begin{cases} \beta_\mu & \text{if } 1 - \epsilon \ll 1, \\ \beta_\mu \left[1 - \frac{\sum_{\nu \neq \mu} \beta_\nu (X_\mu - X_\nu)}{\epsilon^2} \right] & \text{if } \epsilon \ll 1. \end{cases} \quad (2.107)$$

Unsaturated ecosystems. In contrast to the saturated case, when the number of surviving species is less than the number of resources ($\mathcal{S} < \mathcal{R}$) the equations in Eq. (2.99) underdetermine the resource-specific mean fitnesses, \bar{X}_i , so we must invoke the non-linear constraints in Eq. (2.101) to jointly solve for \bar{X}_i and f_μ^* . Alternatively Ref. [59] has shown that the equilibrium values of \bar{X}_i can be obtained from the solution of a convex optimization problem, subject to the constraints in Eq. (2.99). In particular, if we define the transformed variable $h_i = e^{-\bar{X}_i}$, then the equilibrium value of h_i is the solution to the convex optimization problem

$$\vec{h}^* = \operatorname{argmax}_{\vec{h}} \left\{ \sum_i \beta_i \log h_i : \sum_i \alpha_{\mu,i} h_i = e^{-X_\mu} \forall \mu \right\}. \quad (2.108)$$

In fact, this method yields a general solution for the equilibrium value of \bar{X}_i for *any* initial collection of strains, provided that the equality constraints in Eq. (2.108) are replaced by inequalities (\leq). Given the equilibrium values of \bar{X}_i , the surviving species correspond to the indices μ where the equality condition is satisfied. The corresponding values of f_μ^* satisfy the (generally overdetermined) set of equations in Eq. (2.101), which can be inverted using constrained linear regression. We employ this technique to implement the SSWM simulations in Appendix 2.6 below.

We note that since the objective function in Eq. (2.108) depends on β_i , the equilibrium values of \bar{X}_i will also generally depend on the environmental supply vector in an unsaturated ecosystem, in contrast to the β -independent values obtained in the saturated case. Thus, the ecosystem is no longer able to dynamically adjust to “shield” the internal selection pressures from the current state of the environment [59, 51, 3]. Shifts in β_i can therefore lead to new opportunities for evolutionary adaptation.

Evolution in a binary resource usage model

Since there are few empirical constraints on the genetic architecture of resource strategies in the limit of many resources ($\mathcal{R} \gg 1$), we focused on a toy “binary usage” model similar to the one considered by Ref. [59]. In this model, genomes can either encode the ability to utilize a given resource or not (e.g. through the presence or absence of a particular pathway), so that the resource strategy is of the form

$$\alpha_{\mu,i} = \frac{I_{\mu,i}}{\sum_i I_{\mu,i}}. \quad (2.109)$$

where $I_{\mu,i} \in 0,1$ is a binary indicator variables. Individuals can acquire loss-of-function mutations at rate $U_\alpha^- \sum_i I_{\mu,i}$, which cause one of the values of $I_{\mu,i} = 1$ to switch to $I_{\mu,i} = 0$. We also assume that they can acquire gain of function “mutations” (e.g. horizontal acquisition of a gene from the environment) at rate $U_\alpha^+ \mathcal{R}$, which force a randomly chosen uptake rate to the $I_{\mu,i} = 1$ state. Under these assumptions, the pure mutation dynamics will lead to an binomial ensemble of resource strategies, analogous to the one considered by Ref. [59], with a “success probability” of

$$\frac{\langle \sum_i I_{\mu,i} \rangle}{\mathcal{R}} = \frac{U_\alpha^+}{U_\alpha^-}. \quad (2.110)$$

For simplicity, we will assume that the \mathcal{R} resources are all supplied at nearly identical rates. Note that in the completely symmetric case ($\beta_i = 1/\mathcal{R}$), a “generalist” strain with $I_{\mu,i} = 1$ will constitute a marginal evolutionary stable state. To avoid this pathological behavior, we will consider small perturbations around the completely symmetric state:

$$\beta_i = \frac{1}{\mathcal{R}} \left(1 + \epsilon_i - \frac{1}{\mathcal{R}} \sum_j \epsilon_j \right), \quad (2.111)$$

where the ϵ_i are small random perturbations drawn from some distribution, and sorted in descending order ($\epsilon_1 \geq \epsilon_2 \geq \dots \epsilon_{\mathcal{R}}$). For simplicity, we will assume that the ϵ_i are i.i.d. Gaussian variables with scale $\epsilon \ll 1$. The steep tail ensures that the maximum perturbation scales as

$$\epsilon_1 \sim \sqrt{2\epsilon^2 \log \mathcal{R}}, \quad (2.112)$$

and can be bounded to be sufficiently small for a suitable choice of ϵ . Under these assumptions, an ecosystem comprised of a single “generalist” strain will still have nonzero ecological selection pressures encoded by the resource-specific mean fitnesses,

$$\bar{X}_i \approx -\epsilon_i, \quad (2.113)$$

so that some alternate resource strategies will be favored to invade.

By simulating evolutionary dynamics in this model in the weak mutation limit for various values of U_X and s (Appendix 2.6), we find that the long-term structure of the ecosystem tends toward a state in which there is a single generalist strain and $\mathcal{S} - 1$ single loss-of-function variants that have recently descended from this strain (Figs. 2.5 and 2.7-2.9). In the limit that $U_X/U_\alpha \rightarrow 0$, this state must also coincide with a saturated state ($\mathcal{S} = \mathcal{R}$). If we let f_1 denote the frequency of the generalist strain, then the equilibrium frequencies are given by

$$f_1 = 1 - (\mathcal{R} - 1)\epsilon_1, \quad (2.114)$$

$$f_i = \left(1 - \frac{1}{\mathcal{R}} \right) (\epsilon_1 - \epsilon_i), \quad (2.115)$$

where we have assumed that $\mathcal{R}\epsilon_1 \ll 1$. In other words, all resources except the one with the largest value of ϵ_i will have a loss-of-function strain. In the limit that $\epsilon_1\mathcal{R} \ll 1$, the loss-of-function strains will constitute a tiny fraction of the population, and most mutations will arise in the generalist strain. In particular, the accumulation of fitness mutations will cause the fitness of the generalist strain to grow as $X_1 \sim NU_X s^2 t$. We therefore wish to understand when and how this fitness differential drives some of the loss-of-function variants to extinction.

Due to the symmetry of the system, if j strains are driven to extinction, these must be strains with loss-of-function mutations in genes with the next j largest values of ϵ_i , i.e. $i = 2, \dots, j+1$. Let $X_c(j)$ denote the critical value of X_1 required for these j strains to go extinct. Expanding Eq. (2.1) in the main text to lowest order in X_1 , $(1 - f_1)$, $1/\mathcal{R}$, and ϵ , the equilibrium frequencies satisfy

$$R - 1 = \sum_i (1 - \delta_{\mu,i}) \left[\frac{\mathcal{R} - 1}{\mathcal{R}} (1 - X_1 f_1) + \epsilon_i + \frac{f_1}{\mathcal{R}} + f_i \delta_{i>j+1} \right], \quad (2.116)$$

or

$$f_i \approx \frac{1 - f_1}{\mathcal{R}} - \epsilon_i - \mathcal{R} X_1 f_1. \quad (2.117)$$

To self consistently solve for $1 - f_1$, we sum over $i = j+2, \dots, \mathcal{R}$ to obtain:

$$1 - f_1 \approx \mathcal{R} \cdot \frac{1}{j+1} \sum_{k=1}^{j+1} \epsilon_k - \frac{\mathcal{R}^2 (\mathcal{R} - 1 - j)}{j+1} X_1. \quad (2.118)$$

Substituting this back into our expression for f_i , we obtain:

$$f_i \approx \frac{1}{j+1} \sum_{k=1}^{j+1} \epsilon_k - \epsilon_i - \frac{\mathcal{R}^2}{j+1} X_1. \quad (2.119)$$

We can then self-consistently solve for j by setting $f_{j+1} = 0$. For example, if $j = 1$, then we have

$$X_c(1) \sim \frac{\epsilon_1 - \epsilon_2}{\mathcal{R}^2}. \quad (2.120)$$

This will be a quenched random variable, since we have assumed that the ϵ_i are randomly distributed. Given our Gaussian assumption, the typical value of $\epsilon_1 - \epsilon_2$ will occur for

$$\epsilon_1 - \epsilon_2 \sim \frac{\epsilon \log \log \mathcal{R}}{\sqrt{\log \mathcal{R}}}, \quad (2.121)$$

which yields

$$X_c(1) \sim \epsilon \cdot \frac{\log \log \mathcal{R}}{\mathcal{R}^2 \sqrt{\log \mathcal{R}}}. \quad (2.122)$$

On the other hand, if $j \sim \mathcal{R}$, we have

$$X_c(j) \sim \epsilon \cdot \frac{1}{\mathcal{R}}. \quad (2.123)$$

These two fitness scales are separated by a gap of order

$$\frac{X_c(\mathcal{R})}{X_c(1)} \sim \frac{\mathcal{R}\sqrt{\log \mathcal{R}}}{\log \log \mathcal{R}}, \quad (2.124)$$

which grows increasingly large as $\mathcal{R} \gg 1$.

We can use these results to derive heuristic expressions for the number of species \mathcal{S} at steady state as a function of U_α/U_X . We first consider the limit where $\mathcal{S} \ll R$. As mentioned above, the generalist strain comprises the vast majority of the population, so that to a first approximation, we can assume that all fitness and strategy mutations occur on this genetic background. Furthermore, since $\mathcal{S} \ll \mathcal{R}$, most loss of function mutations will target a resource i that does not already have a loss-of-function variant, where the resource-specific mean fitness is given by

$$\bar{X}_i \approx \log \left(\frac{\alpha_{1,i} e^{X_1}}{\beta_i} \right) \approx X_1 - \epsilon_i. \quad (2.125)$$

According to Eq. (2.16), the invasion fitness for a loss-of-function variant that targets resource i is given by

$$S_{\text{inv}} \sim \frac{-\epsilon_i}{\mathcal{R}}. \quad (2.126)$$

Since these loss of function mutations are produced from the generalist background at rate NU_α per resource, the number of coexisting strains increases at rate

$$\frac{d\mathcal{S}}{dt} = \sum_i NU_\alpha \cdot \frac{|\epsilon_i|}{R} \theta(-\epsilon_i) \sim NU_\alpha \epsilon. \quad (2.127)$$

in the absence of fitness mutations.

However, as we mentioned above, the accumulation of fitness mutations will cause the fitness of the generalist strain to grow as $X_1(t) \sim NU_X s^2 t$. Since loss-of-function variants do not acquire further fitness mutations of their own, their fitness is frozen at whatever fitness the generalist strain had at the time that the mutation arose. The fitness difference between the mutant and the generalist therefore grows with time until it reaches a critical value $X_c(\mathcal{R}) \sim \epsilon/\mathcal{R}$, at which point the loss-of-function variant is driven to extinction. This gives rise to two characteristic dynamical regimes depending on whether $X_c(\mathcal{R})$ is larger or smaller than the effect s of a typical fitness mutation.

If $s \ll X_c(\mathcal{R})$, then the generalist lineage must acquire multiple fitness mutations to drive one of the loss-of-function variants to extinction. To a first approximation, the fitness

difference between the generalist and the j th most-recently created loss-of-function variant in this regime is given by

$$\Delta X_j \sim \frac{j}{NU_\alpha \epsilon} \cdot NU_X s^2. \quad (2.128)$$

The number of coexisting ecotypes \mathcal{S} at steady-state is therefore determined by the relation $\Delta X_{\mathcal{S}} \sim X_c(\mathcal{R})$, which reflects a balance between the elimination of the oldest loss-of-function variant due to the accumulation of fitness mutations and the production of new loss-of-function variants through strategy mutations. Solving for \mathcal{S} , we obtain the scaling relation,

$$\mathcal{S} \sim \frac{1}{R} \frac{U_\alpha}{U_X} \left(\frac{\epsilon}{s} \right)^2, \quad (2.129)$$

listed in Eq. (2.19) in the main text.

On the other hand, if $s \gg X_c(\mathcal{R})$, then a single fitness mutation in the generalist strain is sufficient to drive loss-of-function variants to extinction. Before this mutation arises, all the loss-of-function variants will share the same fitness difference ($\Delta X_j = 0$), and this value suddenly shifts to $\gg X_c(\mathcal{R})$ once the successful fitness mutation occurs, driving all of the existing loss-of-function variants to extinction. This leads to oscillations in \mathcal{S} in the time between successive fitness mutations, which range from $\mathcal{S}_{\min} \sim 1$ immediately after the fitness mutation arises, to a maximum value of $\mathcal{S}_{\max} \sim U_\alpha \epsilon / U_X s$ right before the next mutation arises. Since the loss-of-function variants accumulate linearly with time, this leads to a time-averaged value

$$\mathcal{S} \sim \frac{U_\alpha \epsilon}{U_X s}. \quad (2.130)$$

Both expressions should remain valid up to the point where there is an appreciable probability that new loss-of-function mutations target a resource that already has pre-existing variant ($\mathcal{S} \sim \mathcal{R}$). However, there can still be a broad intermediate regime between this point and the point where the ecosystem is completely saturated ($\mathcal{R} - \mathcal{S} \lesssim 1$). The saturated state will coincide with the evolutionary steady-state if the generalist strain is able to seed fitter loss-of-function variants into all the relevant resource dimensions before $X_1(t)$ increases to the point $X_c(1)$, where the first strains start to go extinct. Once again, there are two characteristic timescales depending on whether $X_c(1)$ is large or small compared to s .

If $s \ll X_c(1)$, then the generalist lineage must acquire multiple fitness mutations to before the first loss-of-function variants are driven to extinction. This will happen over a timescale,

$$T_{\text{collapse}} \sim \frac{X_c(1)}{NU_X s^2} \sim \frac{\epsilon}{NU_X s^2} \cdot \frac{\log \log \mathcal{R}}{\mathcal{R}^2 \sqrt{\log \mathcal{R}}}. \quad (2.131)$$

During this time, loss-of-function mutations will occur in the generalist background at rate $NU_\alpha \mathcal{R}$ and will establish with probability $\sim X_1(t)$. Since the loss-of-function mutations are

chosen randomly, $\mathcal{R} \log \mathcal{R}$ such establishments are required to cover the total number of resource dimensions with high probability [12]. This requires a timescale T_{div} that satisfies

$$\int_0^{T_{\text{div}}} NU_{\alpha} \mathcal{R} NU_x s^2 t \sim \mathcal{R} \log \mathcal{R}, \quad (2.132)$$

or

$$T_{\text{div}} \sim \sqrt{\frac{\log \mathcal{R}}{NU_{\alpha} NU_x s^2}}. \quad (2.133)$$

The ecosystem will remain saturated if $T_{\text{collapse}} \gg T_{\text{div}}$, which leads to the condition

$$U_X \ll \frac{U_{\alpha}}{\mathcal{R}^4 \log^2 \mathcal{R}} \left(\frac{\epsilon}{s}\right)^2. \quad (2.134)$$

We can compare this point to the transition to $\mathcal{R} - \mathcal{S} \sim \mathcal{R}$ from Eq. (2.19), which shows that these regimes are separated by a gap of order

$$\frac{U_X(\mathcal{R} - \mathcal{S} \sim 1)}{U_X(\mathcal{R} - \mathcal{S} \sim \mathcal{R})} = (\mathcal{R} \log \mathcal{R})^2, \quad (2.135)$$

while

$$\frac{U_X(\mathcal{R} - \mathcal{S} \sim \mathcal{R})}{U_X(\mathcal{S} \sim 1)} = \mathcal{R}^2. \quad (2.136)$$

Limits on the number of utilized resources

The fragility of the diversification-selection balance in Appendix 2.6 can be attributed in large part to the emergence of a fit generalist strain that is able to utilize all of the available resources. In practice, however, there might be biological constraints or other costs that limit the number of resources that a given strain can metabolize. This leads us to consider an extension of our binary resource usage model, in which the maximum number of utilized resources is capped at some value $\mathcal{R}_c \ll \mathcal{R}$. In this way, we can consider complex ecosystems ($\mathcal{R} \rightarrow \infty$) while restricting the metabolic repertoire of any given strain. A full analysis of this model is beyond the scope of the present work. Instead, we will outline a heuristic calculation that suggests that diversification-selection balance at large \mathcal{R} is achieved for substantially higher values of \mathcal{S} than in Appendix 2.6 above.

We first note that when $\mathcal{R}_c \ll \mathcal{R}$, multiple strains are required to cover the available resources. The minimum possible number of strains is $\mathcal{S} \sim \mathcal{R}/\mathcal{R}_c$, which is achieved when each of the strains specializes on a disjoint subset of \mathcal{R}_c resources. To lowest order in ϵ , the frequencies of these strains are given by

$$f_{\mu} \approx \frac{\mathcal{R}_c}{\mathcal{R}}. \quad (2.137)$$

With the same genetic architecture of strategy mutations that we assumed above, this state will form the basis of the new diversification-selection balance. Generalizing our analysis above, this state will consist of $\mathcal{R}/\mathcal{R}_c$ independent copies of the diversification-selection balance in Appendix 2.6, except with $\mathcal{R} \rightarrow \mathcal{R}_c$. The strains that utilize \mathcal{R}_c resources will be prevented from branching into new resources because of the maximum resource capacity. Meanwhile, single loss-of-function mutants on these backgrounds will be too small to acquire a gain-of-function mutation before their parent acquires enough fitness differences to drive them to extinction.

However, this behavior is strongly dependent on the specific genetic architecture that we assumed, as well as our focus on the SSWM limit. In larger populations, there may be a substantial probability for strains to acquire multiple strategy mutations in a short period of time, which would allow them to break out of their resource neighborhood. To mimic this effect in the SSWM limit, we can introduce a new rate $U_\alpha^{(2)} \ll U_\alpha$ to represent the probability that two strategy mutations arise in the same lineage in a single generation. In particular, we will use this new rate to model *resource swap* events, in which one of the currently utilized resources is deleted and replaced with a randomly drawn resource. As above, we will assume that this mutation rate scales with the number of utilized resources, so that the net rate is given by $U_\alpha^{(2)} k_\mu$, where $k_\mu = 1/\sum_i \alpha_{\mu,i}^2$.

In this augmented model, if we start from the set of $\mathcal{R}/\mathcal{R}_c$ disjoint strains, then the fitnesses of these strains will wander diffusively as $X_\mu \sim NU_X f_\mu s^2 t \pm \sqrt{NU_X f_\mu s^3 t}$, so that the typical fitness differences between a pair of strains is of order $\Delta X(t) \sim \sqrt{NU_X f_\mu s^3 t}$. These fitness differences will create a selection pressure for strategy mutations that swap a resource from one of the fitter strains with a resource from one of the less fit strains. Such a mutation will have an invasion fitness

$$S_{\text{inv}} = \frac{e^{\Delta X(t)} - 1}{\mathcal{R}_c} \approx \frac{\Delta X(t)}{\mathcal{R}_c}. \quad (2.138)$$

Successful swap mutations will be produced on a timescale $\tau_{\text{diversify}}$ that satisfies

$$\int_0^{\tau_{\text{diversify}}} NU_\alpha^{(2)} \mathcal{R}_c \cdot \frac{\sqrt{NU_X f_\mu s^3 t}}{\mathcal{R}_c} dt \sim 1. \quad (2.139)$$

Solving for $\tau_{\text{diversify}}$, we obtain

$$\tau_{\text{diversify}} \sim \frac{1}{s} (NU_\alpha^{(2)})^{-2/3} \left(NU_X \cdot \frac{\mathcal{R}_c}{\mathcal{R}} \right)^{-1/3}, \quad (2.140)$$

$$\Delta X(\tau_{\text{diversify}}) \sim s \left(\frac{U_X \mathcal{R}_c}{U_\alpha^{(2)} \mathcal{R}} \right)^{1/3}. \quad (2.141)$$

Once the successful swap mutation invades, it will create a new ecotype that coexists with the parent clade, as well as the ecotype that currently utilizes the new resource. To

lowest order in $1/\mathcal{R}_c$, the equilibrium frequency is given by

$$f_\nu^* \sim \frac{1 - e^{-\Delta X(t)} \mathcal{R}_c}{2} \frac{\mathcal{R}_c}{\mathcal{R}}. \quad (2.142)$$

When $\Delta X \ll 1$, this frequency will be small compared to the other dominant ecotypes. As above, fitness mutations will therefore preferentially accumulate in the dominant ecotypes, causing the fitness advantage, $\Delta X(t)$, of the swap mutant to decrease over time at rate $NU_X f_\mu^* s^2$. After a time of order

$$\tau_{\text{collapse}} \sim \frac{\Delta X(\tau_{\text{diversify}})}{NU_X s^2 \cdot \frac{\mathcal{R}_c}{\mathcal{R}}} \sim \frac{1}{s} \left(\frac{U_X \mathcal{R}_c}{U_\alpha^{(2)} \mathcal{R}} \right)^{1/3} \frac{\mathcal{R}}{NU_X \mathcal{R}_c}, \quad (2.143)$$

the fitness of the less fit ecotype will have caught up to the swap mutant, and the latter will be driven to extinction. The ratio between τ_{collapse} and $\tau_{\text{diversify}}$ is therefore given by

$$\frac{\tau_{\text{collapse}}}{\tau_{\text{diversify}}} \sim \left(\frac{U_\alpha^{(2)} \mathcal{R}}{U_X \mathcal{R}_c} \right)^{1/3}. \quad (2.144)$$

If this ratio is sufficiently large, then new swap mutants will typically establish before the fitness differences drive any of the existing swap mutants to extinction. For fixed \mathcal{R}_c , this will become increasingly true $\mathcal{R} \rightarrow \infty$.

On the other hand, if $\tau_{\text{collapse}} \ll \tau_{\text{diversify}}$, then a typical resource swap mutation will be driven to extinction before the next arises. However, because the timing of the swap mutations is a random process, anomalously late mutations may occur for which $\Delta X(t) \sim \mathcal{O}(1)$. In this case, the swap mutant is no longer rare compared to its parent, and there is strong selection pressure for loss- (and later gain-) of-function mutations to arise in this mutant background.

Together, these arguments suggest that the simplest generalization of the steady-state in Appendix 2.6 will generally be unstable whenever we impose a cap on the number of utilized resources, and that the corresponding diversification-balance will be attained for much higher values of \mathcal{S} than we would expect based on our previous analysis. Further analysis of these dynamics are left for future work.

Connections to adaptive dynamics

Our model shares certain features associated with the traditional models studied in adaptive dynamics [21, 10], though it also differs from these models in several key ways. In this section, we attempt to make this connection more explicit, using the notation and terminology employed in the adaptive dynamics literature. As adaptive dynamics relies on the weak mutation limit, we will confine our discussion to this regime as well.

Two resources, no fitness differences

For simplicity, we will start by considering the two-resource case in the absence of fitness differences, where individuals are described by a scalar resource phenotype α . To make the connection with adaptive dynamics explicit, we will define a rescaled trait,

$$x = \frac{\alpha - \beta}{\sqrt{\beta(1 - \beta)}}, \quad (2.145)$$

such that $x \rightarrow 0$ as $\alpha \rightarrow \beta$. Following Ref. [21], we then let $s_x(y)$ denote the invasion fitness of a mutant of phenotype y in a monomorphic population of phenotype x . In the neighborhood of $x \rightarrow 0$, Eq. (2.4) shows that $s_x(y)$ takes on a simple quadratic form

$$s_x(y) = (x - y)x. \quad (2.146)$$

Under the standard adaptive dynamics assumption that y is infinitesimally close to x , the phenotypes will evolve in the direction of the fitness gradient,

$$D(x) = \left. \frac{\partial s_x(y)}{\partial y} \right|_{y=x} = -x. \quad (2.147)$$

This gradient vanishes for $x = 0$ which shows that $x^* = 0$ (or $\alpha = \beta$) is an evolutionarily singular strategy. This strategy is *convergence stable*, in that infinitesimal mutations drive the population toward $x = 0$ when $|x| > 0$. However, it is only marginally ESS-stable, since $\partial^2 s_x(y)/\partial y^2 = 0$ at $x = 0$. The second derivative classification in Ref. [21] also shows that stable dimorphisms can coexist in the neighborhood of x^* .

These two features combine to make our evolutionarily singular strategy behave as *both* an evolutionarily stable strategy (ESS) and an evolutionary branching point. On the one hand, $x^* = 0$ resembles an ESS because no mutant strains are favored to invade once the population reaches x^* . On the other hand, $x^* = 0$ resembles an evolutionary branching point because the population will typically branch into a stable dimorphism once $x - x^*$ approaches the typical spacing between mutants. Thus, in practice, the population will always branch before it reaches the ESS, even if this is excluded under truly infinitesimal evolution. However, unlike a traditional branching point where $\partial^2 s_x^*(y)/\partial^2 y > 0$, there is no further selection to drive the branched phenotypes x_1 and x_2 away from each other once branching has occurred. We showed in the text that this can be viewed as a generic feature of a saturated ecosystem (where $\mathcal{S} = \mathcal{R}$) when there are no overall fitness differences between strains.

Resource continuum, no fitness differences

One might ask why the evolutionarily singular strategy is so peculiar in our model, given that consumer-resource theory is often touted as an example of evolutionary branching points in the adaptive dynamics literature [1]. The key difference is that in this existing literature, the trait x does not usually refer to the uptake rate of a single resource, but instead is used to parameterize an entire curve of resource uptake rates for a continuum of different resources. To choose a simple example, one might imagine that the resources denote seeds of different sizes, which are indexed by a continuous parameter z . The function $\beta(z)$ then represents the distribution of seed sizes supplied by the environment, which is often assumed to have a Gaussian form

$$\beta(z) = \frac{e^{-\frac{z^2}{2}}}{\sqrt{2\pi}}, \quad (2.148)$$

centered at some special value $z = 0$. Individual uptake rates are often assumed to have a similar Gaussian shape

$$\alpha(z|x) = \frac{e^{-\frac{(z-x)^2}{2\sigma^2}}}{\sqrt{2\pi\sigma^2}}, \quad (2.149)$$

with a preferred value of $z = x$ and a characteristic width σ . The trait x is then subject to further evolution, rather than the individual uptake rates $\alpha(z)$. Substituting these functions into Eq. (2.1) (with $X_\mu = 0$), the invasion fitness for phenotype y in a monomorphic population of phenotype x is given by

$$s_x(y) = \exp \left(\frac{1}{\sigma^2} \left[\frac{x^2}{2} \left(1 + \frac{1}{\sigma^2} \right) - \frac{y^2}{2} \left(1 - \frac{1}{\sigma^2} \right) - \frac{2xy}{\sigma^2} \right] \right) - 1. \quad (2.150)$$

The fitness gradient $\partial s_x(y)/\partial y$ vanishes when $x = 0$, showing that $x^* = 0$ is an evolutionarily singular strategy, as anticipated. The second derivative is given by

$$\frac{\partial^2 s_x^*(y)}{\partial y^2} = \frac{1 - \sigma^2}{\sigma^4}. \quad (2.151)$$

For $\sigma > 1$, x^* is a true ESS, while for $\sigma < 1$, x^* is a true evolutionary branching point. In the neighborhood of x^* , selection will act to drive the phenotypes further apart from each other after branching has occurred.

We can understand this behavior using the intuition developed in the main text. Since there are an infinite number of resources in this model, the ecosystem is certainly not saturated when $\mathcal{S} = 2$. Thus, we can expect much of the selection pressure to focus on bringing the population averaged uptake rate $\bar{\alpha}(z)$ closer to the environmental supply rate $\beta(z)$. When the niche width σ is larger than the range of resources supplied by the environment, the best way to do this is with a single strain centered at $x = 0$. Branching is therefore not favored.

On the other hand, if σ is smaller than the range of supplied resources, then the ecosystem as a whole can match the environmental supply rate better if there are two strains centered at intermediate locations on the real axis ($|x - y| > 0$).

In this way, we see that the $\mathcal{R} = 2$ resource case, far from being pathological, serves as a basic building block that allows us to understand more complex scenarios that are often considered in the literature. It also illustrates how the genetic architecture of the uptake rates (in this case, whether the $\alpha(z)$ can evolve independently or are restricted to the Gaussian family) can play a key role in determining the emergent dynamics of the model.

Directional selection as an intermediate asymptotic

We now return to the two-resource case above and examine how changes in the overall fitness (X) alter the adaptive dynamics analogy. Individuals are now described by a two-dimensional phenotype, (X, α) . Generalizing our analysis above, we will now define a two-dimensional trait space:

$$x_1 = \frac{\alpha - \beta}{\sqrt{\beta(1 - \beta)}}, \quad x_2 = X. \quad (2.152)$$

In this notation, the invasion fitness in Eq. (2.7) becomes

$$s_{x_1, x_2}(y_1, y_2) = (y_2 - x_2) + (x_1 - y_1)x_1, \quad (2.153)$$

whose fitness gradient is given by

$$\nabla_y s_{x_1, x_2}(y_1, y_2) = (-x_1, 1). \quad (2.154)$$

As expected, the overall fitness dimension always selects for phenotypes that increase X , regardless of the value of α . As a consequence, there are no longer any evolutionarily singular strategies in this model, so the formal classification such points in the adaptive dynamics framework does not apply any more. Nevertheless, we have seen that behaviors very similar to evolutionary branching still occur in our model if we project down onto the x_1 coordinate. Furthermore, the old evolutionarily singular strategy at $x_1^* = 0$ continues to play a key role in these dynamics. The major difference is that these ecologically stable polymorphisms are now only quasi-stable under evolutionary perturbations, as our analysis in the main text shows that further fitness evolution can drive one of the ecotypes to extinction (Fig. 2.3B). This behavior is consistent with observations from laboratory evolution experiments [29].

Although we have motivated this behavior with the abstract notion of overall fitness, our analysis suggests that similar behavior will generically arise in multi-dimensional phenotype spaces whenever one of the traits (i) approaches its marginal branching point x_i^* , while at least one of the other traits (j) remains far from x_j^* (Fig. 2.6). Previous work suggests that such highly asymmetric approaches to a stationary point may be common feature of gradient descent dynamics in high dimensional spaces [37]. This suggests that the competition between resource and strategy mutations can be viewed as a more general intermediate




Figure 2.6: **An intermediate asymptotic of adaptive dynamics.** In a multidimensional phenotype space, a population that is far from the evolutionarily singular strategy can display the quasi-stable branching behavior analyzed in the main text if one of the trait dimensions (x_1) is close to the singular coordinate (x_1^*). In the specific context of our consumer resource model, x_1 corresponds to the resource uptake strategy (α), while x_2 corresponds to the overall fitness (X).

asymptotic that describes the process of ecological diversification during the asymptotically long times required to approach an evolutionarily singular strategy in a high dimensional trait space.

Of course, our simplified consumer-resource model is peculiar in that it contains just a single marginal ESS ($\vec{\alpha} = \vec{\beta}$) even for large \mathcal{R} . For generic high-dimensional landscapes, in contrast, previous work suggests that the relevant critical points will often be more akin to saddle points [37]. In this case, we conjecture that Fig. 2.6 would describe the approach to one such saddle point, until the positive eigenvalues combine to drive the population away from the current critical point and toward a saddle of lower index. This hypothesis suggests that repeated bouts of diversification and ecosystem collapse could occur even in glassy regimes where the population never reaches a true ESS. A detailed understanding of these dynamics remains an interesting topic for future work.

Simulations

Individual-based simulations

The simulations in Fig. 2.1 were carried out using an individual-based, discrete generation algorithm similar to the one employed in Refs. [28, 27] for a single resource. Each simulation starts with a clonal population of N individuals, and in each subsequent generation, the population undergoes a selection step followed by a mutation step. At each step, we keep track of the number of individuals n_μ with a given strategy vector $\alpha_{\mu,i}$ and overall fitness X_μ .

In the selection step, each lineage n_μ is assigned a new size from a Poisson distribution with mean

$$\lambda_\mu = C \left(\sum_{i=1}^{\mathcal{R}} \alpha_{\mu,i} e^{X_\mu - \bar{X}_i} \right), \quad (2.155)$$

where

$$\bar{X}_i = \log \left[\sum_{\mu} \frac{\alpha_{\mu,i} e^{X_\mu}}{\beta_i} \cdot \left(\frac{n_\mu}{\sum_{\nu} n_\nu} \right) \right], \quad (2.156)$$

and $C = N / \sum \lambda_\mu$ is a normalization constant chosen to ensure that the total population size remains near $N \pm \mathcal{O}(\sqrt{N})$.

In the mutation step, the new lineage size is pruned into multiple sublineages representing different mutations that occur on the original lineage background. With probability U_X , an individual founds a new sublineage ν is founded with fitness $X_\nu = X_\mu + s$, where s is drawn from the distribution of fitness effects $\rho_X(s)$. With probability U_α , an individual founds a new sublineage with a strategy vector $\vec{\alpha}_\nu$ drawn from the distribution $\rho_\alpha(\vec{\alpha}'|\vec{\alpha})$.

The simulations in Fig. 2.1 were carried out for $\mathcal{R} = 2$ with $\beta = 0.5$. We utilized a Gaussian distribution of fitness effects, $\rho_X(s) \propto \exp(-s^2/2s_0^2)$, for the pure fitness mutations. The distribution of strategy mutations, $\rho_\alpha(\alpha'|\alpha)$, was taken to be a beta distribution with mean α and coefficient of variation $\text{Var}(\alpha')/E(\alpha')^2 = 0.05$, but with α' rounded to the nearest value of $1/5, \dots, 4/5$. The initial resource strategy for the ancestral population was chosen uniformly at random from these discrete values.

Each simulation was performed for a total of 60,000 generations. Every 500 generations, we simulated a round of “metagenomic sequencing”. We calculated the population frequencies of all mutations present in the population, and reported these values after binomial resampling at a depth of $D = 1000$.

A copy of our implementation in C++ is available on Github (https://github.com/benjaminhgood/consumer_resource_simulations).

SSWM simulations

To simulate the long-term dynamics of the binary usage model in Fig. 2.4 (Appendix 2.6), we use an optimized simulation algorithm that is specifically designed for the strong-selection, weak mutation (SSWM) regime. Similar to traditional SSWM algorithms in population genetics [44], this algorithm gains an efficiency advantage by simulating only successful invasion events. In our case, however, the successful invasion events can now lead to non-trivial ecological equilibria, in addition to simple fixation.

Our simulations start with a collection of strains $(X_\mu, \vec{\alpha}_\mu)$ at time $t = 0$. To assess convergence to diversification-selection balance, we performed simulations for two initial conditions: (i) a single generalist strain with $\alpha_{\mu,i} = 1/\mathcal{R}$ (ii) a collection of \mathcal{R} specialist strains with $\alpha_{\mu,i} = \delta_{\mu,i}$ and X_μ drawn from a Gaussian distribution with variance $\sigma = 10^{-7}$. Figs 2.8 and 2.9 show a comparison of these two initial conditions for $\mathcal{R} = 10$. Since the agreement is generally good, we utilized the more rapidly converging generalist initial conditions for the main simulations in Figs 2.4 and 2.7.

After drawing the initial condition, we first calculate the ecological equilibrium, f_μ^* , for this collection of strains using the convex optimization procedure in Appendix 2.6, using the MOSEK software package [4]. This algorithm yields the equilibrium values of $h_i^* = e^{-\bar{X}_i}$ and the set of ecotypes Σ^* that survive at equilibrium. Within this subset, the equilibrium frequencies are obtained from the solution of the linear system in Eq. (2.101), which will be overdetermined when $\mathcal{S} < \mathcal{R}$. We obtain a solution to this system by solving the constrained least squares problem,

$$\vec{f}^* = \operatorname{argmin}_{\vec{f}} \left\{ \sum_i \left| \sum_{\mu \in \Sigma^*} \alpha_{\mu,i} e^{X_\mu} f_\mu - \frac{\beta_i}{h_i^*} \right|^2 : \sum_{\mu \in \Sigma^*} f_\mu = 1 \right\}, \quad (2.157)$$

using the SciPy library [32].

Once the initial ecological equilibrium is obtained, the simulation proceeds via a series of virtual timesteps, each of which represents the successful invasion of a single mutation. In each step, we first enumerate the set of fitness and strategy mutants that are generated from mutations on each of the current strains μ , and calculate their corresponding invasion fitness from Eq. (2.97). We use these values to calculate the net rate of successful invasions from each mutation type. We assume that fitness mutations confer a characteristic fitness benefit s , so that the rate of successful fitness mutations in strain μ is given by

$$R_\mu^X = NU_X f_\mu^* (e^s - 1). \quad (2.158)$$

The rate of successful loss-of-function mutations for resource i is given by

$$R_{\mu,i}^- = \max \left\{ 0, NU_\alpha f_\mu^* \cdot \sum_{j \neq i} \frac{\alpha_{\mu,j} (h_j^* - h_i^*)}{k_\mu - 1} \cdot \theta(\alpha_{\mu,i} - \delta) \right\}, \quad (2.159)$$

where $k_\mu = 1/\sum_i \alpha_{\mu,i}^2$ is the current number of resources utilized by strain μ , $\theta(z)$ is the Heavisde step function, and δ is an infinitesimally small positive number so that the step function is well-defined. The rate of successful gain-of-function mutations is given by an analogous expression,

$$R_{\mu,i}^+ = \max \left\{ 0, NU_\alpha f_\mu^* \cdot \sum_{j \neq i} \frac{\alpha_{\mu,j}(h_i^* - h_j^*)}{k_\mu + 1} \cdot \theta(\delta - \alpha_{\mu,i}) \right\}. \quad (2.160)$$

Since these successful invasion events arise as a compound Poisson process, the time T_{est} to the next successful invasion event is exponentially distributed with rate

$$R_{\text{tot}} = \sum_\mu \left[R_\mu^X + \sum_i (R_{\mu,i}^+ + R_{\mu,i}^-) \right]. \quad (2.161)$$

Using the Poisson thinning property, the identity of the invading mutation is chosen at random from the enumerated list with probability proportional to its corresponding R -value. Once the identity of the invading strain is determined, we find the new ecological equilibrium \vec{h}^* and \vec{f}^* using the constrained procedure above. By assumption, the time to reach this new equilibrium is negligible compared to T_{est} in the SSWM limit. The current time t is then incremented by T_{est} , and the process repeats itself.

We repeated this process for a total of M successful invasion steps until the ecosystem converged to diversification-selection balance ($M \sim 100,000$). The simulations in Figs. 2.5 and 2.7-2.9 were carried out for $\epsilon = 10^{-3}$ and $s = 10^{-7}$, scanning through different values of U_X/U_α .

A copy of our implementation in Python is available on Github (<http://github.com/StephenMartis/consumer-resource-many-resources>).




Figure 2.7: **Ecological structure at diversification-selection balance in a binary usage model.** (Top) For each of the simulated populations in Fig. 2.4, the fraction of the population occupied by the generalist ecotype, $\alpha_{\mu,i} = 1/\mathcal{R}$. (Bottom) For the same populations, the frequency-weighted average of $k_\mu = 1/\sum_i \alpha_{\mu,i}$ (a measure of the number of utilized resources) for the remaining non-generalist ecotypes. A value of $\mathcal{R} - \langle k_\mu \rangle = 1$ indicates that the rest of the population consists of single loss-of-function mutants that descend from the generalist background.




Figure 2.8: **Approach to diversification-selection balance from different initial conditions.** An analogous version of Fig. 2.4 comparing specialist and generalist initial conditions for $\mathcal{R} = 10$.

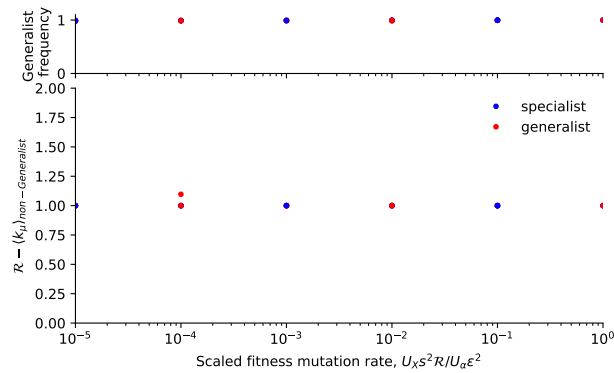


Figure 2.9: **Long-term ecological structure from different initial conditions.** An analogous version of Fig. 2.7 comparing specialist and generalist initial conditions for $\mathcal{R} = 10$.

2.7 Bibliography

- [1] M. Ackermann and M. Doebeli. Evolution of niche width and adaptive diversification. *Evolution*, 58(12):2599–2612, 2004.
- [2] G. Ackland and I. Gallagher. Stabilization of large generalized Lotka-Volterra foodwebs by evolutionary feedback. *Phys Rev Lett*, 93(15):158701, 2004.
- [3] 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(3):033406, 2018.
- [4] E. D. Andersen and K. D. Andersen. The mosek interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In *High performance optimization*, pages 197–232. Springer, 2000.
- [5] M. L. Bendall, S. L. Stevens, L.-K. Chan, S. Malfatti, P. Schwientek, J. Tremblay, W. Schackwitz, J. Martin, A. Pati, B. Bushnell, et al. Genome-wide selective sweeps and gene-specific sweeps in natural bacterial populations. *ISME J*, 10(7):1589–1601, 2016.
- [6] I. Cvijović, B. H. Good, E. R. Jerison, and M. M. Desai. Fate of a mutation in a fluctuating environment. *Proceedings of the National Academy of Sciences*, 112(36):E5021–E5028, 2015.
- [7] M. M. Desai and D. S. Fisher. Beneficial mutation selection balance and the effect of genetic linkage on positive selection. *Genetics*, 176:1759–1798, 2007.
- [8] U. Dieckmann, P. Marrow, and R. Law. Evolutionary cycling in predator-prey interactions: population dynamics and the red queen. *Journal of theoretical biology*, 176(1):91–102, 1995.
- [9] M. Doebeli. A model for the evolutionary dynamics of cross-feeding polymorphisms in microorganisms. *Population Ecology*, 44(2):59–70, 2002.
- [10] M. Doebeli. *Adaptive Diversification*. Princeton University Press, 2011.
- [11] W. J. Ewens. *Mathematical Population Genetics*. Springer-Verlag, New York, second edition, 2004.
- [12] W. Feller. *An introduction to probability theory and its applications*, volume 2. John Wiley & Sons, 2008.
- [13] S. E. Finkel and R. Kolter. Evolution of microbial diversity during prolonged starvation. *Proceedings of the National Academy of Sciences*, 96(7):4023–4027, 1999.

- [14] D. S. Fisher. Leading the dog of selection by its mutational nose. *Proc Natl Acad Sci USA*, 108:2633–2634, 2011.
- [15] D. S. Fisher. Asexual evolution waves: fluctuations and universality. *J Stat Mech*, 2013: P01011, 2013.
- [16] K. R. Foster, J. Schluter, K. Z. Coyte, and S. Rakoff-Nahoum. The evolution of the host microbiome as an ecosystem on a leash. *Nature*, 548(7665):43, 2017.
- [17] E. M. Frenkel, M. J. McDonald, J. D. Van Dyken, K. Kosheleva, G. I. Lang, and M. M. Desai. Crowded growth leads to the spontaneous evolution of semistable coexistence in laboratory yeast populations. *Proc Natl Acad Sci USA*, 112(36):11306–11311, 2015.
- [18] M. L. Friesen, G. Saxer, M. Travisano, and M. Doebeli. Experimental evidence for sympatric ecological diversification due to frequency-dependent competition in *Escherichia coli*. *Evolution*, 58(2):245–260, 2004.
- [19] C. Gardiner. *Handbook of Stochastic Methods*. Springer, New York, 1985.
- [20] N. R. Garud, B. H. Good, O. Hallatschek, and K. S. Pollard. Evolutionary dynamics of bacteria in the gut microbiome within and across hosts. *bioRxiv*, 2017. doi: 10.1101/210955.
- [21] S. A. Geritz, G. Mesze, J. A. Metz, et al. Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary ecology*, 12(1): 35–57, 1998.
- [22] P. Gerrish and R. Lenski. The fate of competing beneficial mutations in an asexual population. *Genetica*, 127:127–144, 1998.
- [23] J. H. Gillespie. *Population Genetics: A concise guide*. Johns Hopkins University Press, Baltimore, MD, second edition, 1998.
- [24] J. H. Gillespie. Genetic drift in an infinite population: The pseudohitchhiking model. *Genetics*, 155:909–919, 2000.
- [25] B. H. Good. *Molecular Evolution in Rapidly Evolving Populations*. PhD thesis, 2016.
- [26] B. H. Good and M. M. Desai. Fluctuations in fitness distributions and the effects of weak linked selection on sequence evolution. *Theor Pop Biol*, 85:86–102, 2013.
- [27] B. H. Good and M. M. Desai. The impact of macroscopic epistasis on long-term evolutionary dynamics. *Genetics*, 199(1):177–190, 2015.
- [28] B. H. Good, I. M. Rouzine, D. J. Balick, O. Hallatschek, and M. M. Desai. Distribution of fixed beneficial mutations and the rate of adaptation in asexual populations. *Proc Natl Acad Sci USA*, 109:4950–4955, 2012.

- [29] B. H. Good, M. J. McDonald, J. E. Barrick, R. E. Lenski, and M. M. Desai. The dynamics of molecular evolution over 60,000 generations. *Nature*, 551(7678):45, 2017.
- [30] R. B. Helling, C. N. Vargas, and J. Adams. Evolution of *Escherichia coli* during growth in a constant environment. *Genetics*, 116(3):349–358, 1987.
- [31] M. D. Herron and M. Doebeli. Parallel evolutionary dynamics of adaptive diversification in *Escherichia coli*. *PLoS Biology*, 11(2):e1001490, 2013.
- [32] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL <http://www.scipy.org/>.
- [33] M. Kardar. *Statistical physics of fields*. Cambridge University Press, 2007.
- [34] N. Kashtan, S. E. Roggensack, S. Rodrigue, J. W. Thompson, S. J. Biller, A. Coe, H. Ding, P. Marttinen, R. R. Malmstrom, R. Stocker, et al. Single-cell genomics reveals hundreds of coexisting subpopulations in wild prochlorococcus. *Science*, 344(6182):416–420, 2014.
- [35] D. G. Kendall. On the generalized "birth-and-death" process. *The annals of mathematical statistics*, 19(1):1–15, 1948.
- [36] M. Kondoh. Foraging adaptation and the relationship between food-web complexity and stability. *Science*, 299(5611):1388–1391, 2003.
- [37] J. Kurchan and L. Laloux. Phase space geometry and slow dynamics. *Journal of Physics A: Mathematical and General*, 29(9):1929, 1996.
- [38] G. I. Lang, A. W. Murray, and D. Botstein. The cost of gene expression underlies a fitness trade-off in yeast. *Proceedings of the National Academy of Sciences*, 106(14):5755–5760, 2009.
- [39] 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 forty evolving yeast populations. *Nature*, 500(7464):571–574, 2013.
- [40] S. F. Levy, N. Ziv, and M. L. Siegal. Bet hedging in yeast by heterogeneous, age-correlated expression of a stress protectant. *PLoS biology*, 10(5):e1001325, 2012.
- [41] R. Mac Arthur. Species packing, and what competition minimizes. *Proceedings of the National Academy of Sciences*, 64(4):1369–1371, 1969.
- [42] R. H. MacArthur and E. O. Wilson. *The theory of island biogeography*. Princeton university press, 1967.

- [43] J. Mallet. Hybridization, ecological races and the nature of species: empirical evidence for the ease of speciation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1506):2971–2986, 2008.
- [44] D. M. McCandlish and A. Stolzhus. Modeling evolution using the probability of fixation: history and implications. *Quarterly Review of Biology*, 89:225–252, 2014.
- [45] M. J. McDonald, S. M. Gehrig, P. L. Meintjes, X.-X. Zhang, and P. B. Rainey. Adaptive divergence in experimental populations of *Pseudomonas fluorescens*. iv. genetic constraints guide evolutionary trajectories in a parallel adaptive radiation. *Genetics*, 183(3):1041–1053, 2009.
- [46] J. O. McInerney, A. McNally, and M. J. O’Connell. Why prokaryotes have pangenomes. *Nature microbiology*, 2(4):17040, 2017.
- [47] R. A. Neher. Genetic draft, selective interference, and population genetics of rapid adaptation. *Annual Review of Ecology, Evolution, and Systematics*, 44:195–215, 2013.
- [48] P. Nosil, L. J. Harmon, and O. Seehausen. Ecological explanations for (incomplete) speciation. *Trends in ecology & evolution*, 24(3):145–156, 2009.
- [49] J. Plucain, T. Hindré, M. Le Gac, O. Tenaillon, S. Cruveiller, C. Médigue, N. Leiby, W. R. Harcombe, C. J. Marx, R. E. Lenski, et al. Epistasis and allele specificity in the emergence of a stable polymorphism in *Escherichia coli*. *Science*, 343(6177):1366–1369, 2014.
- [50] S. R. Poltak and V. S. Cooper. Ecological succession in long-term experimentally evolved biofilms produces synergistic communities. *ISME J*, 5(3):369–378, 2011.
- [51] A. Posfai, T. Tallefumier, and N. S. Wingreen. Metabolic trade-offs promote diversity in a model ecosystem. *Phys Rev Lett*, 118(2):028103, 2017.
- [52] P. B. Rainey and M. Travisano. Adaptive radiation in a heterogeneous environment. *Nature*, 394(6688):69–72, 1998.
- [53] D. E. Rozen and R. E. Lenski. Long-term experimental evolution in *Escherichia coli*. viii. dynamics of a balanced polymorphism. *American Naturalist*, 155(1):24–35, 2000.
- [54] B. J. Shapiro, J.-B. Leducq, and J. Mallet. What is speciation? *PLoS genetics*, 12(3):e1005860, 2016.
- [55] N. Shores, M. Hegreness, and R. Kishony. Evolution exacerbates the paradox of the plankton. *Proc Natl Acad Sci USA*, 105(34):12365–12369, 2008.
- [56] J. M. Smith. The genetics of stasis and punctuation. *Annual review of genetics*, 17(1):11–25, 1983.

- [57] A. Sousa, R. S. Ramiro, J. Barroso-Batista, D. Güleresi, M. Lourenço, and I. Gordo. Recurrent reverse evolution maintains polymorphism after strong bottlenecks in commensal gut bacteria. *Molecular biology and evolution*, 34(11):2879–2892, 2017.
- [58] L. Sun, H. K. Alexander, B. Bogos, D. J. Kiviet, M. Ackermann, and S. Bonhoeffer. Effective polyploidy causes phenotypic delay and influences bacterial evolvability. *PLoS biology*, 16(2):e2004644, 2018.
- [59] M. Tikhonov and R. Monasson. Collective phase in resource competition in a highly diverse ecosystem. *Phys Rev Lett*, 118(4):048103, 2017.
- [60] M. Tikhonov and R. Monasson. Innovation rather than improvement: a solvable high-dimensional model highlights the limitations of scalar fitness. *Journal of Statistical Physics*, pages 1–31, 2017.
- [61] D. Tilman. *Resource competition and community structure*. Princeton University Press, 1982.
- [62] V. Torsvik and L. Øvreås. Microbial diversity and function in soil: from genes to ecosystems. *Current opinion in microbiology*, 5(3):240–245, 2002.
- [63] C. C. Traverse, L. M. Mayo-Smith, S. R. Poltak, and V. S. Cooper. Tangled bank of experimentally evolved burkholderia biofilms reflects selection during chronic infections. *Proc Natl Acad Sci USA*, 110(3):E250–E259, 2013.
- [64] L. Van Valen. A new evolutionary law. *Evol Theory*, 1:1–30, 1973.
- [65] K. Vetsigian. Diverse modes of eco-evolutionary dynamics in communities of antibiotic-producing microorganisms. *Nature Ecology & Evolution*, 1:0189, 2017.
- [66] S. Zhao, T. D. Lieberman, M. Poyet, M. Groussin, S. M. Gibbons, R. J. Xavier, and E. J. Alm. Adaptive evolution within the gut microbiome of individual people. *bioRxiv*, page 208009, 2017.

Chapter 3

Quenched disorder and directional selection lead to punctuated equilibrium in the competition for substitutable resources

Preface

We start from the results from the previous chapter and focus specifically on the high-dimensional limit, where types are competing for many resources. In particular, let's focus on the emergence of a generalist 'steady-state' and the diversification-selection balance. This is not necessarily the case for different choices of genetic architecture, although we only showed this for a peculiar special case in which compound mutational events were allowed. In the current chapter, I seek to explore a class of plausible alternate genetic architectures by appealing to an ensemble approach. In this approach I define a class of disordered models, randomly drawn from a distribution and make statements at the level of the distribution, rather than any particular instantiation. I achieve this by adding a small perturbation to the simple 'binary usage' genetic architecture: that resource strategies have intrinsic random fitness effects. This ends up strongly affecting the eco-evolutionary dynamics, resulting in a type of 'phase' behavior. In addition, I recast the Lyapunov function of the consumer-resource model into a form that allows for a simpler interpretation of the eco-evolutionary process.

The remainder of this chapter will be submitted for publication in an edited form.

3.1 Abstract

Quenched disorder in the form of rugged fitness landscapes can strongly impact the mode and tempo of evolutionary dynamics, as well as the observed distribution of genotypes in a population. However, in natural and laboratory settings, evolution and complex ecological dynamics can play out in tandem. When this is the case, populations can split into coexisting ‘ecotypes’ with different ecological characteristics and varying degrees of overlap between their niches. It is much less clear how ruggedness might impact this *eco-evolutionary* scenario or affect distributions of ecotypes and their characteristics. To address this, we extend a simple model of eco-evolutionary dynamics based on the competition for substitutable resources by incorporating quenched disorder. We have previously demonstrated that without disorder, fit generalists tend to dominate the population at long times [11]. However, weak disorder is sufficient to destabilize the generalist fixed point, giving rise to an extended regime in which the dynamics can be summarized by a cycle consisting of short bursts of diversification and long periods spent waiting for fitness mutations. Furthermore, in this ‘punctuated equilibrium’ phase, the composition of ecotype strategies continually changes. We postulate that at long times the dynamics are ergodic in metabolic strategy space. Finally, we suggest how these model dynamics might be relevant in real ecological communities.

3.2 Introduction

In natural settings, microbes often exist in large, diverse communities, competing for a common pool of resources. As they compete and grow, they are subject to evolutionary forces such as selection, genetic drift and recombination on timescales that are observable. Studies have demonstrated that even in controlled laboratory experiments initialized with clonal populations, random mutations can lead to stable diversification into coexisting types [28, 3]. After diversification, there is continued adaptation within diversified lineages [10], which can perturb the equilibria at which coexistence is maintained [30]. Therefore, it seems that evolution can drive fluctuations in microbial ecosystems, potentially perturbing their stability and affecting their composition.

Given this observation, coupled with the principle of competitive exclusion and the stochastic effects inherent to evolutionary processes, it might be surprising then that extensive genomic diversity, even at the sub-species level, persists in natural microbial systems [20]. One might expect that an exceptionally fit generalist type would evolve by chance and sweep out any diversity within a set of closely related lineages, but this seems not to be the case in key natural examples [14, 1, 25]. The persistence of fine-scale diversity presents an additional theoretical challenge – how do seemingly discrete clades persist over tens of millions of generations in the face of adaptation? Why do we not instead observe more of a continuum of recently diverged types?

Much theoretical work has attempted to model persistent diversity by focusing primarily on ecological considerations. Models based on the classical consumer-resource and Lotka-

Volterra frameworks show that diversity can persist in the many-ecotype limit, given constraints on the distribution of intertype interactions [5], spatial structure [27, 2], time dependent environmental fluctuations [13], the absence of fitness differences [29], as well as other more particular mechanisms (e.g. so-called ‘Kill-the-winner’ dynamics) [32]. Crucially, many of these coexistence mechanisms neglect that continued evolution, in addition to assembly and competition, can play a major role in shaping observed microbial community structure. Furthermore, since these models do not incorporate evolution or the notion of inheritance, they are intrinsically unable to comment on the generation of the long-lived microdiversity that has been observed in natural systems.

In recent work [11], we explicitly modeled the eco-evolutionary process by allowing for evolutionary dynamics to occur on the background of MacArthur’s consumer-resource model with linear metabolic trade-offs [18, 29, 33]. We split evolutionary change into two components: 1) changes in metabolic strategies (diversity-generating mutations) 2) changes in fitness (selective mutations). We showed that in the context of this model, the generation of diversity and the selective forces eliminating it can balance each other, leading to a non-equilibrium steady state in which the relative rate of directional selection controls the number of co-existing species at long times [11]. However, in the long time limit, the community can be extremely simple, even in regimes where the rate of strategy mutations is very high. The final stable community consists of a dominant generalist type which seeds closely related types that only ever achieve low frequencies. Because this state is stable to the evolutionary dynamics (see Supplement), additional structure like spatially separated demes cannot readily facilitate the coexistence of ecotypes with diverged strategies at long times.

Key to this behavior is the linear budget trade-off we considered, which has been considered in other contexts [29, 33]. Trade-offs between phenotypic traits have been extensively studied in the ecology literature and are expected to play a major part in ecosystem composition and stability [7, 35]. Recent work has identified the persistence of variants with trade-offs in closely related strains [4]. Modeling work has tried to understand the underlying effect of metabolism on ecological structure [21] and the effect of nonlinear tradeoffs on coexistence and the stability of generalists [26, 17]. A recent simulation study looked at evolution with nonlinear tradeoffs in a consumer resource model on 3 resources [6] and observed various limiting behaviors depending on the shape of the constraint surface. Crucially, the trade-offs considered were smooth, and directional selection was not considered. In the evolution literature, rugged fitness landscapes have also been extensively studied in purely competitive dynamics without any sort of additional ecological structure [22, 23]. However, it remains an open question how such ruggedness might affect the dynamics of systems evolving on a high dimensional phenotypic space.

Here, we propose a relatively simple extension of the evolving consumer-resource model in which we incorporate a ‘rugged’ (or disordered) metabolic budget. We show that even weak disorder leads to a transition in which the generalist-dominated diversification-selection balance phase is destabilized. The generalist state gives way to a ‘punctuated equilibrium’ phase in which ecotypes are not localized in strategy space over long times. In purely

competitive models of recombinant evolution, it has been demonstrated that ruggedness can lead to clustering of alleles into discrete clones [22]. In our eco-evolutionary model, ruggedness can lead to long-lived clustering of resource strategies that bears qualitative similarity to the ‘genomic backbones’ observed in natural populations [14]. Finally, we show that in the punctuated equilibrium phase, since there is no single attracting state, a simple metapopulation inevitably exceeds the bounds on diversity set by competitive exclusion.

3.3 Consumer resource model with quenched disorder

We construct a model of \mathcal{S} ecotypes competing for \mathcal{R} substitutable resources (e.g. different carbon sources) in a well-mixed ‘chemostat.’ Nutrients are supplied to the chemostat at normalized rates β_i , such that $\sum_i \beta_i = 1$. For convenience, we have $\beta_i \sim 1/\mathcal{R} + \epsilon_i$ with $\epsilon_i \sim \mathcal{N}(0, \epsilon)$. Generally we work in the ‘near uniform’ supply limit where $\epsilon \ll 1/\mathcal{R}$, although many of our qualitative observations will hold for less uniform supply vectors.

Each ecotype’s phenotype is parametrized by a metabolic strategy vector $\alpha_{\mu,i}$ and a scalar ‘energy budget’ e^{X_μ} . The strategy vector is normalized $\sum_i \alpha_{\mu,i} = 1$, lending to its interpretation as a vector of allocations in a budget. The continuous space of possible strategies can be represented as an \mathcal{R} -dimensional simplex due to this constraint. The total energy allocation of the ecotype, the ecotype’s capacity to convert resources into biomass, can be interpreted as a type of scalar fitness. From here on, we will refer to it as such. A separation of timescales between the rate of consumption of nutrients and the growth rate of the species is assumed so that the resources are held at steady state. See SI for more details about the model and some natural interpretations.

With these considerations in place, the population dynamics of the species’ frequencies, f_μ , are given by the following set of \mathcal{S} differential equations:

$$\frac{df_\mu}{dt} = \left(\sum_{i=1}^{\mathcal{R}} \frac{\beta_i \alpha_{\mu,i} e^{X_\mu}}{\sum_\nu \alpha_{\nu,i} e^{X_\nu} f_\nu} - 1 \right) f_\mu \quad (3.1)$$

These population dynamics possess a Lyapunov function, as noted in many previous studies [33, 11], which is maximized along a valid dynamical trajectory. Importantly, this steady state can be found numerically using standard techniques from convex optimization. The equilibrium value of the Lyapunov function has a useful interpretation:

$$H(\vec{f}^*) = \underbrace{\overline{X^*}}_{\text{directional selection}} - \underbrace{\sum_{i=1}^{\mathcal{R}} \beta_i \log \frac{\beta_i}{D_i}}_{\text{ecology}}$$

where we have defined the equilibrium average fitness:

$$\overline{X^*} \equiv \log \langle e^X \rangle = \log \left(\sum_\nu e^{X_\nu} f_\nu^* \right)$$

and the normalized demand:

$$D_i \equiv \sum_{\nu} \alpha_{\nu,i} e^{X_{\nu} - \bar{X}^*} f_{\nu}^*$$

So the Lyapunov function seeks to maximize the mean fitness, while matching the resource supply distribution with the appropriately normalized demand distribution. This quantity $H(\vec{f}^*)$ will increase along any incremental evolutionary trajectory. We discuss this in more detail in the Supplementary Information.

Our rugged consumer resource model only slightly alters the model structure by dictating that the fitness can be split into an additive component and a disordered component that depends on the strategy vector:

$$X_{\mu} = X_{\mu}^{\text{additive}} + X_{\mu}^{\text{disorder}}(\vec{\alpha}_{\mu})$$

The biological motivation for this term is that different ecotypes or species might pay different fitness costs or benefits by arranging their metabolic strategy in specific ways. This can be interpreted as a type of *epistasis*, or non-additive allelic contributions to a phenotype. In this case, the strategy components non-additively combine to contribute to the total budget of a strain. For an assembly process, this has been reported not to significantly change the steady state properties of the system (as described in [33, 29]) – however it will markedly change the evolutionary dynamics. In the SI, we discuss the consequences of this sort of quenched disorder for purely ecological processes like community assembly, but for the rest of the main text, we will be concerned with the setting in which ecotypes are allowed to evolve.

In defining the disordered term, we primarily work with a ‘pairwise’ model in which $X_{\mu}^{\text{disorder}}$ is a function of the pairwise presence/absence of components of the strategy vector:

$$X_{\mu}^{\text{disorder}} = \sum_{i,j=1}^{\mathcal{R}} J_{ij}(\alpha_{\mu,i}; \alpha_{\mu,j})$$

This parametrization is commonly used in the population genetics literature [23], is computationally simple to implement, and is sufficient to possess qualitative properties (e.g. a multi-peaked landscape) we are interested in presently. However, the choice of the particular disorder model should not change the qualitative aspects of our work. The pairwise model and alternatives are discussed in the Supplementary Information.

3.4 Evolution in the binary resource mutational architecture

In order to model the evolutionary process, we have to choose a genetic architecture. This choice is not particularly well-constrained by experimental data, so we opt for a choice that is

maximally tractable while still reasonably well-motivated. Specifically, we allow for constant improvements in fitness (e.g. the ‘staircase model’):

$$X_\mu \rightarrow X_\mu + s$$

and the gain and loss of metabolic function (e.g. ‘bitflip’ mutations) accompanied by disorder effects:

$$\begin{aligned}\vec{\alpha}_\mu &\rightarrow \vec{\alpha}'_\mu = \frac{\vec{\alpha}_\mu \pm \|\vec{\alpha}_\mu\|_2^2 \hat{e}_i}{1 \pm \|\vec{\alpha}_\mu\|_2^2} \\ X_\mu &\rightarrow X_\mu - X^{disorder}(\alpha_\mu) + X^{disorder}(\alpha'_\mu)\end{aligned}$$

where \hat{e}_i is the i^{th} unit vector. Note that $\|\vec{\alpha}_\mu\|_2$ is the L2 norm, while the strategy vector is normalized according to the L1 norm. This form of the strategy mutation is such that normalization is maintained:

$$\sum_{i=1}^{\mathcal{R}} \alpha_{\mu,i} = \sum_{i=1}^{\mathcal{R}} \alpha'_{\mu,i} = 1$$

In this picture, allowed strategies fall on facets of the \mathcal{R} -simplex.

This partition between two mutational ‘types’ captures the notion that only a fraction of a species’ genome contributes to its metabolism, while many mutational loci (including synonymous mutations and ones occurring in non-coding regions [31]) can have significant fitness effects. The staircase mutational model is assumed for its relative simplicity, and for the fact that the distribution of fitness effects can be characterized by a single typical effect size in many simple population genetic models of directional selection [9]. The binary resource architecture serves as a minimal model of gain and loss of function of different metabolic pathways. Microbial genomes are typically very dynamic, mutating rapidly and exchanging genetic material through horizontal gene transfer [20, 8], so we expect this to be a good ‘zeroth order’ approximation.

We work in the strong selection, weak mutation regime, neglecting concurrent mutations and clonal interference, so that mutations arrive at rates proportional to the frequency of the background they appear on and to their *invasion fitness*. The invasion fitness is the growth rate of a mutant when infinitesimally rare. The invasion fitness of a mutant on background μ can be computed as a function of the existing population equilibrium:

$$S_{\text{inv}} = \sum_{i=1}^{\mathcal{R}} e^{X_\mu^{\text{additive}} + X_{\text{mut}}^{\text{disorder}}} \alpha_{\text{mut},i} \frac{\beta_i}{\sum_{\nu} e^{X_\nu} \alpha_{\nu i} f_\nu^*} - 1$$

When this quantity is positive, the mutant can invade, when it is zero or negative, in our infinite population size model, the mutant cannot invade. However, in finite populations and populations undergoing rapid adaptation (as is the case in many microbial populations), such deleterious mutations might be relevant mutations that can reach substantial fractions of the population. Understanding the impact of such mutations is an interesting avenue for future work.

We studied the binary resource model with a ‘flat’ fitness landscape ($X^{\text{disorder}} = 0$) in detail in previous work [11]. At long times, this model was found to reach a non-equilibrium steady state in which a highly fit generalist type, which can consume every resource ($\vec{\alpha} = \vec{1}/\mathcal{R}$), dominates the population. This type, in turn seeds single loss of function variants which persist at low frequencies of order ϵ for time scales of order $\sim \epsilon/U_X s^2$ (see Figure 3.1). These knockouts are wiped out when the generalist acquires sufficiently many fitness mutations, which is increasingly likely in a higher dimensional model since the knockout type frequencies become negligible. However, this picture qualitatively changes with the addition of a disordered fitness component.




Figure 3.1: Generalist type with two knockout types in the competition for three resources. The generalist is the type located at the center of the strategy simplex. With disordered fitness, some fraction of knockouts will have higher fitness than the generalist type and so when they are competing, they will take up a disproportionate share of the total population. With high enough knockout fitness, subsequent mutations are much more likely to occur on these backgrounds, leading to the inevitable extinction of the generalist.

3.5 Results

Fragile generalists in the large \mathcal{R} limit

We start with with a single generalist strain on an ecosystem with \mathcal{R} resources and consider its stability to the mutational process. We simplify our analysis in this section, restricting genotype space to the generalist and single loss-of-function (or knockout) variants (KOVs) with randomly drawn fitness. Stability results derived from this single KOV model can be thought of as rough bounds for the full binary resource model in the sense that if the generalist is unstable in the single KOV model, it will also be unstable in the full model.

In the simplified scenario, KOVs with positive invasion fitness are able to invade. Without ruggedness, these variants only persist at low frequencies of order $f_k \sim \epsilon$, while the generalist remains at frequency $f_g \sim 1$. Therefore any subsequent fitness mutations are more likely to occur on the generalist background, which in turn drives the knockouts to extinction – the system is ‘localized’ to the generalist state. Strategy evolution halts when $\mathcal{R} - 1$ single KO types have invaded since any subsequent mutations have neutral invasion fitness (see Supplement or [11] for details).

However, with the inclusion of disorder, a KOV can be accompanied by a fitness gain, which is potentially substantial. In turn, this fitness gain can result in a much larger f_k (see Figure 3.1 for schematic). The generalist can tolerate this decline in frequency, so long as the fitness differences are not large enough that subsequent mutations are more likely to happen on KOV backgrounds. So there is a critical strength of disorder, σ_c , above which a runaway effect leads to the eventual extinction of the generalist with high probability. We calculate σ_c for two edge cases, $U_X \gg U_\alpha$ and $U_X \ll U_\alpha$.

For $U_X \gg U_\alpha$, we can derive an upper bound for σ_c relatively simply. In this scenario, at time $t = 0$ a strategy mutation seeds a single knockout. We denote the generalist frequency f_g and the knockout frequency $f_k = 1 - f_g$. In the disordered model, there is an *intrinsic* fitness difference between the generalist and the knockout when it is seeded, which we denote $\Delta X(0) = X_g(0) - X_k(0)$. Since there is a fitness differential already, the knockout can have a (substantially) higher frequency than it would in the neutral model.

After the knockout is seeded there will be a long time spent waiting for the next strategy mutation. During this time, fitness mutations will continue to accrue on the two backgrounds and shift the fitness differential, $\Delta X(t) = X_g(t) - X_k(t)$, which will in turn perturb the equilibrium frequencies. The critical point is the fitness difference required so that a single knockout has a higher frequency than a generalist at $t = 0$. In this case subsequent mutations will accrue on the knockout background, eventually driving the generalist to extinction.

We can calculate the critical value by starting with Eq. 3.1 for the two types. Setting the time derivative of Eq. 3.1 to zero, we can solve for the generalist frequency at equilibrium, f_g . This is a function of difference between the generalist and knockout fitness, $\Delta X = X_g - X_{\text{KO}}$, and the supply corresponding to the knocked out resource, β_k :

$$f_g = \frac{\beta_k}{1 - \frac{\mathcal{R}-1}{\mathcal{R}}e^{\Delta X}} \quad (3.2)$$

We set $f_g = 1/2$ to get a condition for the critical knockout fitness a generalist can tolerate:

$$X_k^c = \log \left[\frac{\mathcal{R} - 1}{(1 - 2\beta_k)\mathcal{R}} \right] \quad (3.3)$$

In the many resource setting ($\mathcal{R} \gg 1$), the highest fitness single knockouts will be the ones most likely to invade. Formally in the limit \mathcal{R} , since $\frac{\mathcal{R}-1}{\mathcal{R}} \rightarrow 1$ and $\beta_k \rightarrow 0$, even a small positive fitness effect is sufficient for a single knockout to have a higher frequency than the generalist.

When the disordered fitness is drawn from a distribution, the scaling of this critical fitness with the number of resources is given by the appropriate extremal criterion. For disordered fitness drawn from a Gaussian distribution with scale σ , we have the following scaling for the critical strength of disorder:

$$\sigma_c \lesssim \frac{\log \left[\frac{\mathcal{R}-1}{(1-2\beta_i)\mathcal{R}} \right]}{\sqrt{2 \log \mathcal{R}}} \sim \frac{1}{\mathcal{R} \sqrt{\log \mathcal{R}}} \quad (3.4)$$

which matches numerical results in Fig. 3.2. Importantly since σ_c scales inversely with the number of resources, generalists competing for many resources can be particularly fragile to fitness differences. However, it should be emphasized that σ_c is potentially substantially lower in the full model that allows for multiple knockouts (since many of these will also have fitness advantages).

Next, we consider the opposite dynamical limit, $U_X \ll U_\alpha$. In this case, strategy mutations accrue quickly while fitness mutations accrue slowly. It is difficult to exactly compute σ_c in this case, but once again, we can compute an upper bound. Given a realization of disorder, we can ask what the frequency of the generalist is in the presence of all single knockouts with positive invasion fitness. We can also calculate this equilibrium exactly for arbitrary knockout fitness. We arrive at the phenomenological scaling ansatz:

$$\sigma_c \sim \mathcal{R}^{-2}$$

which agrees very well with numerics (Fig. 3.2).

Importantly we can notice that, in the limit of infinite \mathcal{R} , both of these upper bounds scale down to zero so that the generalist state is extremely fragile. However, in the limit of large but *finite* \mathcal{R} , we see that a large gap opens between the critical values. These points must connect with each other so that there must be a critical line across the phase diagram of the disordered model (see dashed line Fig. 3.2). However, if we turn from this simplified model to the full binary resource architecture where multiple knockouts are allowed, σ_c can be much lower than the upper bound of the single knockout model (see red line Fig. 3.2). This is due to the additional entropy due to double, triple, and higher order KOVs.

Pure strategy evolution results in degenerate ‘shielded’ equilibria

Whereas previously we discussed the stability of the single generalist type on the boundaries of the phase diagram, we now turn the question of what ecology might emerge after the generalist goes extinct. We start in the $\sigma > \sigma_c$ region of the phase diagram, so that generalists are assured to go extinct. For a given realization of random disorder and the environmental supply $\vec{\beta}$, we can observe different dynamical behaviors in the different limits of the phase diagram.

We start by examining the $U_X \gg U_\alpha$ limit. Above the critical disorder line, we are also above the variation in resource supply ϵ . We term this the *weak ecology* regime. Since $\sigma > \epsilon$,

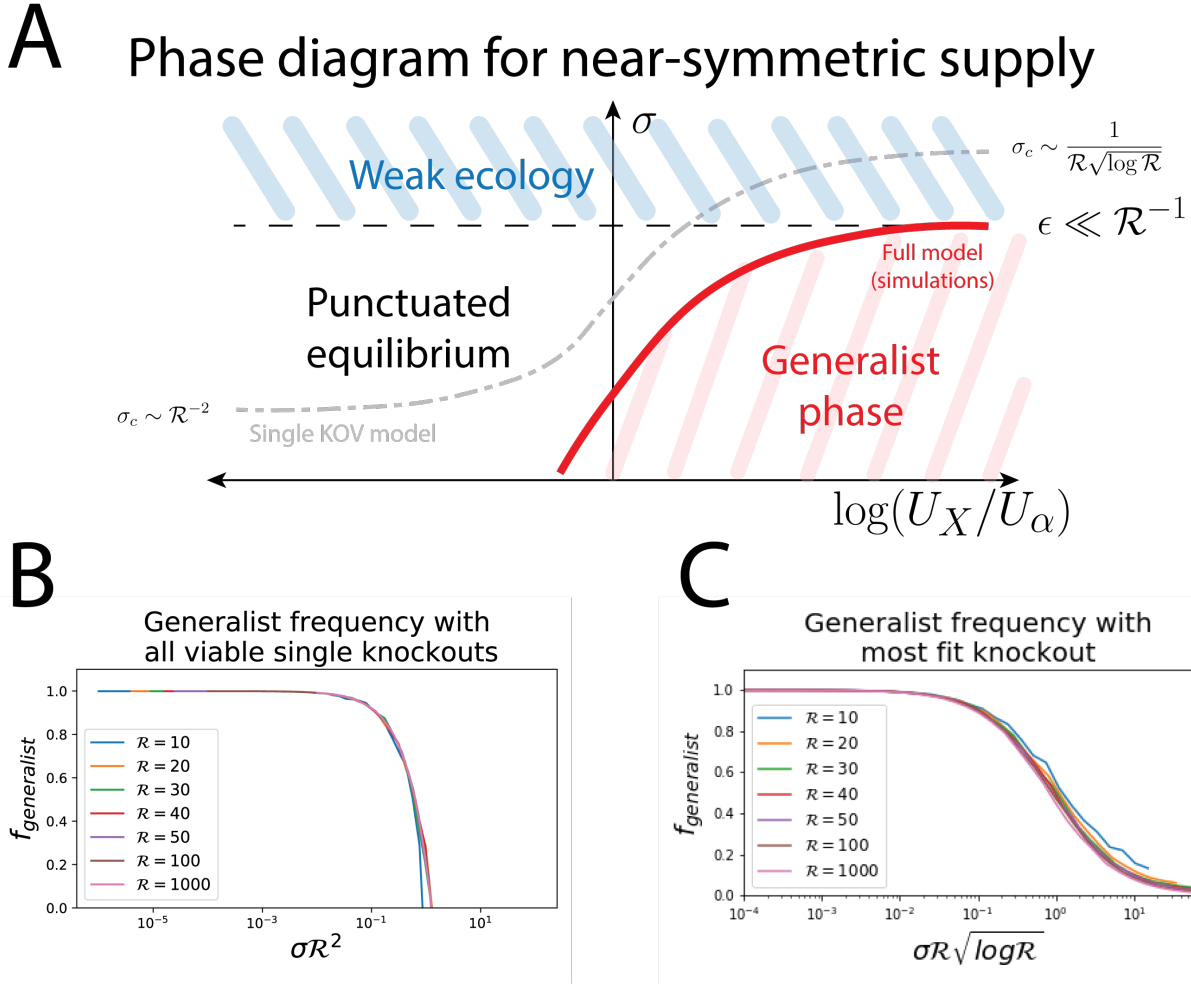


Figure 3.2: **A** Sketch of the eco-evolutionary phase diagram. The red region indicates the phase in which the ecosystem is generalist dominated. For sufficiently high σ and low U_X/U_α , this gives way to a dynamic ‘punctuated equilibrium’ phase which we discuss later in the text. We can estimate the boundary between the generalist phase and the punctuated equilibrium phase by considering a simplified model in which there is competition only between the generalist and single knockouts (the grey dashed line). However, the phase boundary in the full model (the red line) can be substantially lower as seen in simulations. **B** The scaling of the generalist frequency with all single knockouts gives the asymptotic limit of the grey dashed line for slow fitness evolution ($U_X/U_\alpha \ll 1$). **C** The scaling of the generalist frequency with the most fit single knockout gives the asymptotic limit of the grey dashed line for fast fitness evolution ($U_X/U_\alpha \gg 1$).

fitness differences will swamp any ecological differences so that the evolutionary process will favor the mean fitness component of the Lyapunov function, whereas the ecological component will be subleading in the invasion fitness. So the dynamics will select for the highest fitness types that span the resource space. This results in an effective niche partitioning of resource space, whereby strains with relatively high fitness will fix on partially disjoint sets of resources. This is analogous to a fitness peak in an adaptive walk on a rugged landscape [15], except that multiple strains will coexist at the peak. Because the resources are unevenly partitioned, continued fitness mutation accrues differentially across the different strains. Eventually fitness differences can get sufficiently large that one niche invades the others, but this may take exponentially many mutations or requires unrealistically large fitness effect sizes and cannot readily be accessed with simulations.

In the other limit, for sufficiently small U_X , we see that strategy evolution proceeds and eventually halts while the ecosystem waits for the next fitness mutation. We call this static ecosystem a *strategy uninhabitable state* (SUS), in which no local strategy mutants (e.g. single gain or loss-of-function variants) have a positive invasion fitness. With a fixed disorder realization, if we run the dynamics many times, we see that different replicas can reach different SUSs. This type of ergodicity breaking is characteristic of dynamics on rugged landscapes and corresponds to different replicas getting trapped in different local optima of the eco-evolutionary process. See Fig. 3.3 for an illustration of this phenomenon.

We also see that for increasing σ , the fractional filling of a SUS, \mathcal{S}/\mathcal{R} where \mathcal{S} is the number of strains at equilibrium, decays from 1 to approximately $2/\mathcal{R}$ (since the generalist cannot take over for $U_X = 0$, $\epsilon \neq 0$). While it is difficult to make analytical progress on characterizing general properties of a SUS, we can still make some progress with simulations to understand their qualitative properties. Specifically, we can characterize a SUS by looking at the statistics of the equilibrium *resource availabilities*, defined as the ratio between the supply and demand of resources:

$$A_i^* = \frac{\beta_i}{\sum_{\nu} \alpha_{\nu,i} e^{X_{\nu}} f_{\nu}^*}$$

We see that this grows approximately linearly with the strength of disorder, consistent with a ‘shielded’ state as defined in [33]. Such shielded states are robust to small perturbations in the supply $\vec{\beta}$ by dynamically adjusting to minimize ecological selection pressures.

Fitness evolution can create opportunity for transitions between equilibria

Now that we have established a qualitative sketch of the phase diagram and the behavior at the boundaries, we can turn to analyzing the properties of the dynamical phase where the fitness mutation rate becomes comparable to the strategy mutation rate, $U_X/U_{\alpha} \lesssim 1$. We use ratio of the fitness mutation rate and the strategy mutation rate, U_X/U_{α} , as a control parameter. We also fix the rate of overall mutations so that N total mutations (of both

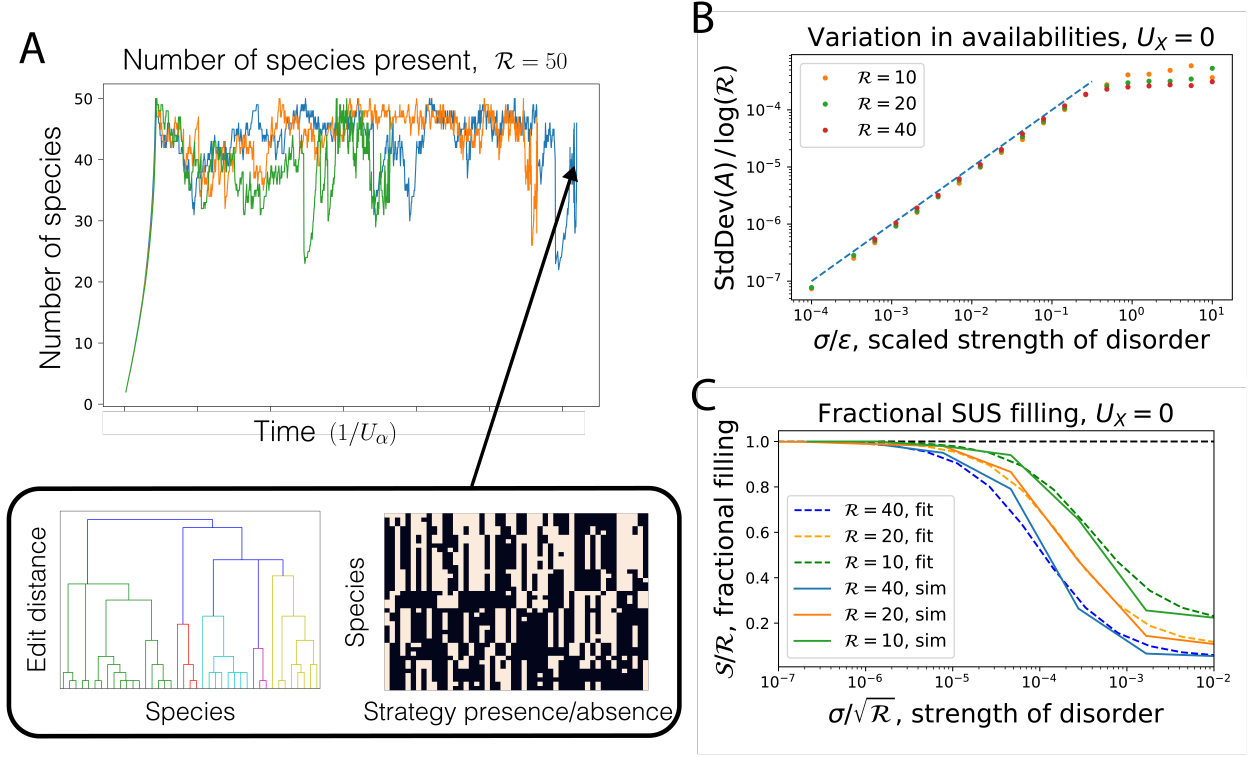


Figure 3.3: **A** Three independent time series of the number of species initialized from the generalist state, fixed realization of disorder, with $\sigma = 10^{-6}$, $s = 10^{-6}$, $\epsilon = 10^{-3}$. Only bitflip strategy mutations are allowed. The number of species increases as loss of function mutants invade, and the generalist is lost rapidly. Eventually evolution halts and no more strategy mutants can invade. Bottom panel shows details of the final SUS state from the blue trajectory. Species cluster, but are not merely related by single knockouts within clusters. **B** Variance in the resource availabilities is grows linearly with the strength of disorder until it exceeds the standard deviation in the resource supply ϵ . The dashed line is a guide to the eye, $\text{StdDev}(A) \sim \sigma$. **C** Fractional filling of the final SUS state, defined as the number of ecotypes divided by the number of resources. Fits are logistic functions. Simulation curves are averaged over ten realizations. The critical disorder strength σ_c (at which the generalist goes extinct) falls far to the left of the range of the plot.

fitness and strategy type) occur in a fixed time T_{tot} . The mutations do not necessarily arrive at evenly spaced intervals since the invasion fitnesses of the different mutational types are dynamic quantities.

We can understand the dynamics in the regime by analyzing the behavior of the Lyapunov function, H . We see that for a wide range of U_X/U_α that H grows approximately linearly in time when averaged over many mutations. This can be understood by considering limiting

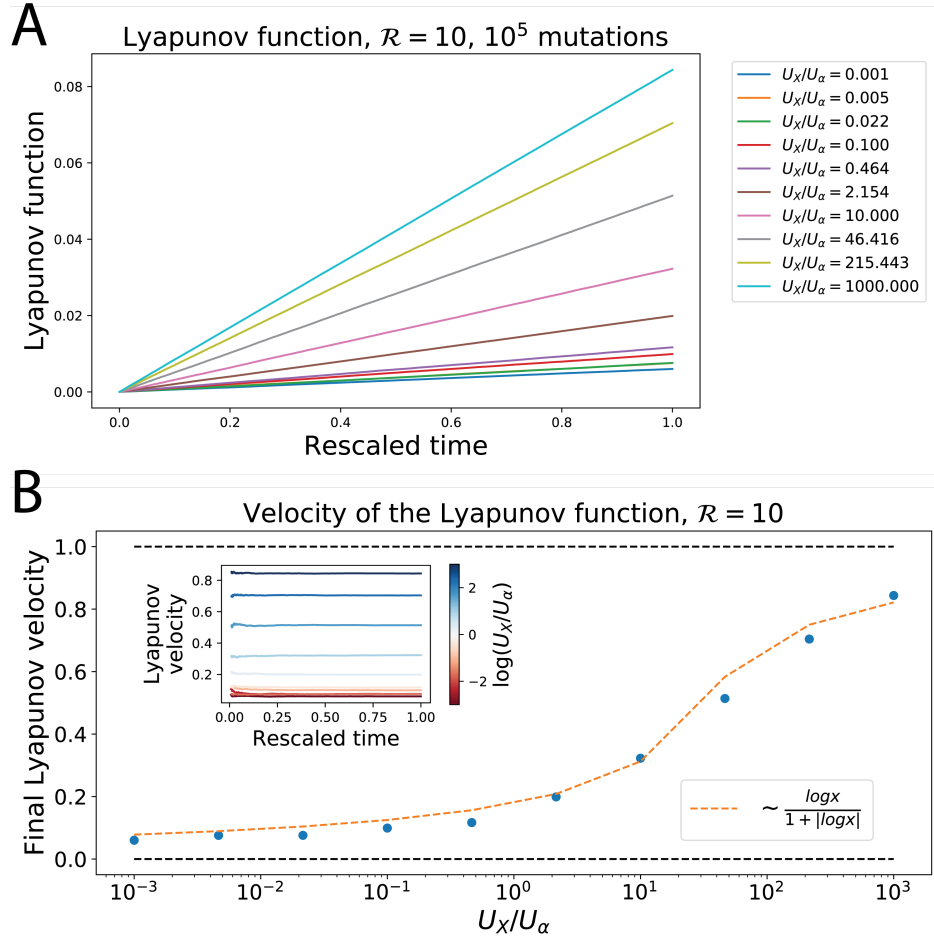


Figure 3.4: **Long-time dynamics of the Lyapunov function.** **A.** The Lyapunov function evolves approximately linearly at long times. Simulations are for $\sigma = 10^{-6}$. **B.** The velocity of Lyapunov function grows with U_X/U_α from 0 to U_Xs (which the y-axis of panel B has been normalized by). A fit with a generalized sigmoid function is provided to guide the eye and to indicate that the decay to the limiting values can be quite slow (e.g. much slower than logistic).

cases. In the limit $U_\alpha = 0$, this can be understood as mutations continually accruing on a single generalist background – this is population genetics picture of sequential adaptation, with a maximal slope given by U_Xs . In the limit of $U_X = 0$, the system reaches an SUS and evolution halts, so that the slope goes to zero. That the Lyapunov dynamics are also approximately linear for intermediate values is akin to the diversification-selection balance described in [11].

However, this is not the complete story since the short time dynamics can be quite differ-

ent. These short time dynamics end up being crucial to the character of the ecosystem and its composition, leading to distinct ‘phases,’ which we will describe. In the limit $U_X/U_\alpha \gg 1$, we have established that the ecosystem is pinned to the generalist state. We have also established that in the regime $U_X/U_\alpha \ll 1$, there is sufficiently rapid strategy evolution that the ecosystem will reach an SUS. Upon reaching the SUS, evolution will effectively halt until a fitness mutation arrives. In this sense, evolution is fitness mutation-limited since there are no local ecological opportunities available. However, when a fitness mutation eventually arrives, this will increase the mean fitness, but at the cost of slightly misaligning the supply and demand distributions. This will start to create new opportunity for strategy mutations. Eventually, sufficiently many fitness mutations will accrue on different strategy backgrounds that a new strategy is able to invade.

These dynamics can be tracked by considering the ecological component of the Lyapunov function alone:

$$D_{KL}(\beta||D) \equiv \sum_i \beta_i \log \left(\frac{\beta_i}{D_i} \right)$$

The form of this term is the Kullback-Leibler divergence between the supply of resources $\vec{\beta}$ and the appropriately normalized equilibrium demand, \vec{D} . By the form of the Lyapunov function, the eco-evolutionary process seeks to minimize this divergence while also maximizing the mean fitness. When fitness mutations are abundant, the optimal way to do this is to have mutations accrue on a single generalist background. When fitness mutations are rare, ecological invasions occur in between the arrival of fitness mutations. However, even if fitness mutations occur slowly, they are crucial for allowing transitions between SUSs. In this sense the ecosystem will cycle through SUSs in a way that is analogous to Gould’s qualitative picture of ‘punctuated equilibrium’ [12]. The dynamics of the Lyapunov function in turn, look ‘stuttered,’ as opposed to the generalist phase where H increases linearly even on shorter times. We plot the short term dynamics of H in the generalist phase and in the punctuated equilibrium phase in Fig. 3.5.

We can identify this punctuated equilibrium phase by appealing to the time-averaged distribution of strategies. At each time point, we store the frequency distribution of strategies. We then time average this distribution over $N \gg 1$ mutational arrivals:

$$\rho_{TA}(\vec{\alpha}) \equiv \frac{1}{T_{tot}} \sum_{i=1}^N t_i f_i(\vec{\alpha})$$

where t_i is the time between mutation i and mutation $i+1$ and where $f_i(\vec{\alpha})$ is the frequency of ecotype strategy $\vec{\alpha}$ after the introduction of mutation i . This time averaged distribution gives the probability that one will observe a given species when sampling the ecosystem over long times. In the generalist phase, the ecosystem will be trapped in or near a single state, and so this distribution will be sharply peaked. However, in the punctuated equilibrium phase, the time-averaged distribution will smear out over all possible strategies as the system travels between SUSs.

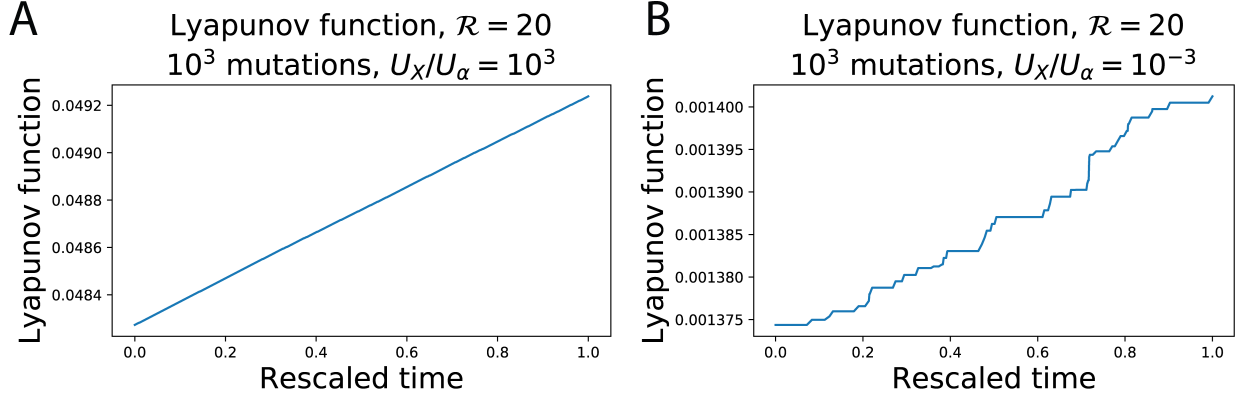


Figure 3.5: **Short time dynamics of the Lyapunov function.** **A.** Lyapunov function grows approximately linearly over short times for $U_X/U_\alpha \gg 1$ as mutations primarily accrue on a generalist background. **B.** However, for $U_X/U_\alpha \ll 1$, the Lyapunov function increases in a more ‘stuttered’ manner since ecological mutations dominate. On longer timescales, H evolves approximately linearly for both cases.

We define an order parameter to distinguish between these two phases which we call the ‘effective number of species:’

$$N_{eff} = \exp \left\{ - \sum_{\vec{\alpha}} \rho_{TA}(\vec{\alpha}) \ln \rho_{TA}(\vec{\alpha}) \right\}$$

Intuitively, this measure tells us the number of uniformly distributed species that would match the entropy of the time-averaged distribution. For a sharply peaked distribution, this number is close to one and for a broader distribution it can be significantly higher (up to $2^{\mathcal{R}}$). This order parameter is plotted for different mutation rates and strengths of disorder in the phase diagram in Fig. 3.6. In addition, we see that the qualitative dynamics of the ecological divergence, $D_{KL}(\beta||D)$, strongly differs between the generalist and punctuated equilibrium phases.

3.6 Discussion

We have mapped out the dynamical phase diagram of an evolving consumer-resource model with disordered fitness. Leveraging the structure of the model, and specifically the Lyapunov function, we are able to identify a fitness-dominated generalist phase and an ecology-dominated punctuated equilibrium phase. In previous work, we had identified the underlying model structure as a fairly generic limit of resource competition [11]. As such, we expect our qualitative conclusions should be relevant for this broader class of models. In addition,

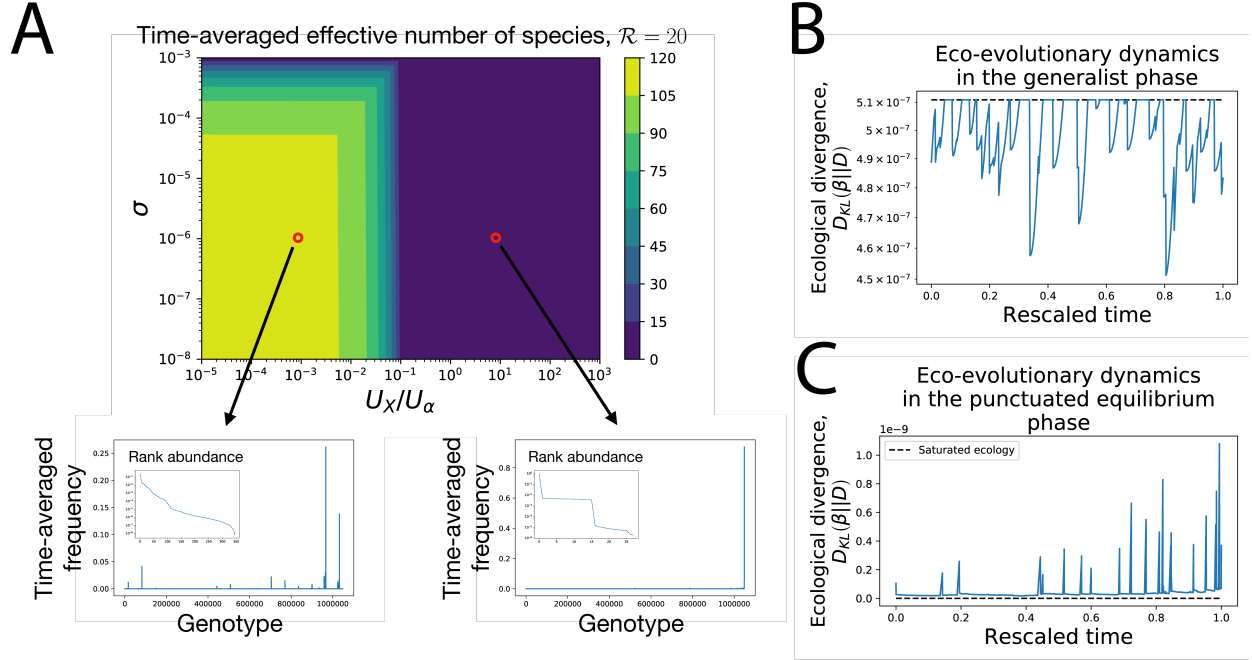


Figure 3.6: **Eco-evolutionary ‘phases.’** **A.** Phase diagram plotting the effective number of species (the exponential of the time-averaged entropy) as a function of the strength of disorder, σ , and the rescaled fitness mutation rate. Note that the transition extends down to extremely low σ . Also note that as σ approaches the scale of the typical niche size the transition disappears. This is a result of $\mathcal{O}(1)$ types having extremely high fitness that cannot be destabilized on simulation timescales. Individual simulations were run for 10^7 mutations. Points in the phase diagram were averaged over 20 simulations each. **B** Ecological divergence in the generalist phase is pinned near its maximal value, which is the divergence of a single generalist type $\max D_{KL} = \sum_i \beta_i \log \mathcal{R} \beta_i$ (which we refer to as a ‘generalist ecology’). **C** In the punctuated equilibrium phase, the ecological divergence stays much closer to zero (which is a ‘saturated ecology’) for long stretches of time, with sharp, stuttered excursions away due to rare fitness mutations. Note the separation of scales between the values of the divergence in **B** and **C**. The lower values in **C** approach the bounds of numerical tolerance in our simulations (10^{-10}).

in real systems, the ability to perform many different complex tasks might result in effective quenched disorder like that which we have proposed here (e.g. trade-off or constraint surfaces might have some complicated rugged functional form). We have demonstrated that such disorder becomes a more and more relevant factor as the dimension of resource space grows.

In particular, since the model might describe ecosystems on intermediate to longer timescales, we will speculate about the relation between our results and long-term labo-

ratory evolution experiments. In laboratory evolution experiments, there is rapid fitness evolution for several thousand generations and approximate clonality is maintained. However, it has been demonstrated that the rate of fitness adaptation can slow down significantly over longer time horizons [37]. As this occurs, diversification is readily observed [10]. In the context of our model, it is expected that an initially clonal population would diversify when substitutable resources are available, especially as the rate of fitness adaptation decreases. Quantitatively, the approximate timescale of this diversification should give an estimate of the phenomenological rate parameter U_α . However, this is merely suggestive and would require extensive experimental work to say anything conclusive.

Our model's domain of applicability is not simply limited to microbial populations competing for resources. For instance this work might be relevant in modeling the evolution of the adaptive immune response over long times. In the adaptive immune system, there is a co-evolutionary process between host immune cells and the pathogens that infect those hosts [24]. It is desirable for the host to obtain 'generalist' or broadly neutralizing antibodies (BNAbs), which are able to bind to many different antigens thus allowing for broad and robust immunity [36]. Understanding the stability of generalists is extremely relevant for understanding the likelihood that BNAbs both arise and persist. Since static consumer resource models have been extensively used to model the statistical properties of immune repertoires [19], this seems like a natural space to apply insights from our dynamical model.

We foresee many potential generalizations of our model to include more relevant ecological and evolutionary processes. For instance, a dynamical model of metabolic rewiring, in which the topology of the disordered landscape can itself evolve, or full explorations of clonal interference regimes are interesting avenues for future work. A full exploration of how the model generalizes to spatially structured environments or to environments with fluctuating resources would also be interesting, especially in light of how the additional timescales would interfere with or amplify the punctuated equilibrium phase.

3.7 Supporting Information (SI)

Consumer resource model

There are a few independent ways to ‘derive’ the consumer resource model equations (Eq (1) in the main text). We reproduce two of these arguments here for the sake of completeness (from [33, 29] and also [11]). In both of these derivations, we imagine that we are concerned with a chemostat with \mathcal{R} *substitutable* resources. These resources may be different sugars or other nutrients providing a single limiting element. We assume that there is outflow from the chemostat at rate δ . We measure time in units of δ so that we can, WLOG, set $\delta = 1$.

Phenomenological derivation: supply and demand

Without any nutrients or growth, a population of type μ is lost at rate $\delta = 1$. In order to be maintained in the chemostat, this type must consume sufficient nutrients to overcome this loss rate. In the presence of nutrients, the growth of type μ can be summarized as:

$$\frac{dn_\mu}{dt} = \left(\sum_{i=1}^{\mathcal{R}} E_{\mu i} A_i - 1 \right) n_\mu$$

We refer to $E_{\mu i}$ as the (unnormalized) strategy vector for type μ . This can be interpreted, loosely speaking, as the amount of “effort” type μ puts in to obtaining resource i . The A_i is defined as the “availability” of resource i given the current composition of the ecosystem.

There are certain physically motivated properties that A_i ought to have which we can use to constrain its functional form. We can consider the total demand for resource i ,

$$D_i = \sum_{\nu=1}^S E_{\nu i} n_\nu$$

We also consider the supply rate of resource i (in units of biomass/time), which is assumed to be constant over time, S_i . Two reasonable assumptions are that: 1) as the supply of a resource increases, its availability should increase 2) as demand for a resource increases, its availability should decrease. A simple function that encapsulates this behavior is simply the ratio between the two:

$$A_i = \frac{S_i}{D_i} = \frac{S_i}{\sum_{\nu} E_{\nu i} n_\nu}$$

so that the dynamics for the strains is:

$$\frac{dn_\mu}{dt} = \left(\sum_{i=1}^{\mathcal{R}} \frac{E_{\mu i} S_i}{\sum_{\nu} E_{\nu i} n_\nu} - 1 \right) n_\mu$$

In turn, we can then separate the “effort” into two components:

$$E_{\mu i} \equiv e^{X_\mu} \alpha_{\mu i}$$

a magnitude, e^{X_μ} , and a normalized ‘strategy vector,’ $\alpha_{\mu i}$, where:

$$\sum_i \alpha_{\mu i} = 1$$

Furthermore, we can readily demonstrate that the total population size in the chemostat is constant at steady state and given by $\sum_i S_i$. Normalizing by the total population size and passing to equations for the frequencies of different strains, we recover our equation from the main text:

$$\frac{df_\mu}{dt} = \left(\sum_{i=1}^{\mathcal{R}} \frac{e^{X_\mu} \alpha_{\mu i} \beta_i}{\sum_\nu e^{X_\nu} \alpha_{\nu i} f_\nu} - 1 \right) f_\mu$$

‘Microscopic’ derivation: quasi-steady state resource dynamics

We can also recover this phenomenological derivation from a more microscopic consideration of the dynamical equations describing resources and strains in a chemostat. We start with the equations for strains:

$$\frac{dn_\mu}{dt} = (g_\mu(\vec{c}) - 1) n_\mu$$

which include a growth rate, g_μ , that depends on the resource concentrations, $\vec{c} = (c_1, c_2, \dots, c_{\mathcal{R}})$. In addition we assume that the growth rate of strain μ is given by the weighted sum:

$$g_\mu(\vec{c}) = \sum_i E_{\mu i} r_i(c_i)$$

where $r_i(c_i)$ is the uptake rate of resource i and $E_{\mu i}$ is an enzyme or energy budget by which strain μ commits different portions of its budget to different resource uptake pathways. The uptake rate is assumed to be fixed across types by appealing to qualitative conservation of pathways across closely related types and can take any reasonable functional form.

In addition, we define the dynamics of the resources:

$$\frac{dc_i}{dt} = S_i - \left(\sum_\nu k_{\nu i} E_{\nu i} n_\nu \right) r_i(c_i)$$

where $k_{\nu i}$ is a conversion factor between biomass produced and nutrients taken up. For simplicity we assume that $k_{\nu i} = k_i$, so that nutrients are converted into biomass at the same fraction across strains. We have also assumed that the vast majority of nutrients supplied are consumed (so the outflow rate is negligible). If consumption is sufficiently fast then, measured on the outflow timescale, the nutrient concentration is approximately constant ($dc_i/dt \approx 0$) so that:

$$r_i(c_i) \approx \frac{S_i/k_i}{\sum_\nu E_{\nu i} n_\nu}$$

This gives us:

$$\frac{dn_\mu}{dt} = \left(\sum_{i=1}^{\mathcal{R}} \frac{E_{\mu i} S_i / k_i}{\sum_\nu E_{\nu i} n_\nu} - 1 \right) n_\mu$$

Therefore from these microscopic considerations, we get the same dynamical equations for the strain frequencies (except with renormalized supply vector S_i/k_i).

Lyapunov function

As described in other contexts, these consumer-resource dynamics possess a Lyapunov function. This function can be written in the form:

$$H(\vec{f}) = 1 - \sum_\nu f_\nu - \sum_i \beta_i \log \left(\frac{\beta_i}{\sum_\nu e^{X_\nu} \alpha_{\nu i} f_\nu} \right)$$

which differs slightly from the versions presented in [33, 11] by the added constant. The 1 makes the role of the terms in H more apparent, but does not change the properties of the function.

This function is convex and bounded from above and increases over any trajectory of the dynamics and as such indicates that there is a unique equilibrium. Any valid trajectory will maximize this function subject to the constraints that $\sum_\nu f_\nu = 1$ and $f_\nu \geq 0$. Looking at the form of the Lyapunov function we can also understand the qualitative character of the dynamics. The first two terms can be interpreted as enforcing a constraint that the frequencies are normalized and always simply sum to zero at equilibrium. The third term encodes the population mean fitness and the Kullback-Leibler divergence between the supply vector and the ‘normalized’ demand (which we will define shortly).

It is illustrative to consider the optimum of H for a generalist type with a fitness X_g . The value of the Lyapunov function achieved by this generalist at equilibrium is given by:

$$H(f_g^* = 1) = - \sum_i \beta_i \log \mathcal{R} \beta_i + X_g$$

The first term encodes the mismatch between the generalist strategy and the environmental supply:

$$D_{\text{KL}}(\vec{\beta} || \vec{\alpha}_g) = \sum_i \beta_i \log \mathcal{R} \beta_i$$

The second term is simply the generalist fitness. So as the generalist accrues fitness mutations (and is replaced by subsequently more fit types, the Lyapunov function increases. Therefore fitness mutations are always favored to invade. Furthermore, ecological opportunity is mediated by differences between supply and demand in the form of the Kullback-Leibler divergence. Both of these concepts can be generalized to more complex community cases, which we shall proceed to do.

One thing we can notice is that although the dynamics are unique up to a fitness translation, the Lyapunov function is not. The absolute fitness of the strains enters into the value

it takes at its optimum. To make clear the effects of fitness vs. strategies, we can consider an ecosystem that maximizes H , given by the set of triples $\{(f_\mu^*, \vec{\alpha}_\mu, X_\mu)\}$. We can separate out the mean fitness:

$$e^{\bar{X}^*} \equiv \sum_{\nu} e^{X_\nu} f_\nu^*$$

Note that this quantity is averaged over ecotypes and not resources (i.e. it does not have an index corresponding to a strategy dimension). As such, it is different from the ‘resource-specific mean fitness’ which we defined in a previous publication [11]. We define the normalized weights:

$$w_\nu^* \equiv e^{X_\nu - \bar{X}^*} f_\nu^*$$

where we have used asterisks to denote equilibrium quantities. Making the substitutions, we get:

$$H(\vec{f}^*) = \bar{X}^* - \sum_i \beta_i \log \frac{\beta_i}{\sum_{\nu} w_\nu^* \alpha_{\nu i}} = \bar{X}^* - \sum_i \beta_i \log \frac{\beta_i}{D_i}$$

where we note that the normalized demand is in fact a valid distribution (e.g. $\sum_i D_i = \sum_i \sum_{\nu} w_\nu^* \alpha_{\nu i} = 1$) so that the second term is precisely a Kullback-Leibler divergence. Therefore, the ecological process seeks to optimize the Lyapunov function while balancing two things simultaneously, maximizing the mean fitness of the resident strains while minimizing the Kullback-Leibler divergence (or the relative entropy [16]) between the supply and the normalized demand.

We can understand the entropy term by looking at the ensemble of all possible communities with fixed mean fitness, \bar{X}^* . Over this ensemble of communities, the optimal ones are the ones that are arranged such that demand matches the normalized supply exactly:

$$\beta_i = \sum_{\nu} e^{X_\nu - \bar{X}^*} \alpha_{\nu i} f_\nu^*.$$

so that the optimal H is just the mean fitness:

$$H(\vec{f}^*) = \bar{X}^*$$

This defines a manifold of optimal communities, which does not strictly constrain the number of coexisting types at equilibrium. It is always possible to define this manifold when there are no fitness differences between the strains present. It is a strictly infinite dimensional manifold of linear combinations of points in the $(\mathcal{R} - 1)$ -simplex. However, there is a more stringent requirement when there are fitness terms, e^{X_ν} . We consider this and describe the condition geometrically in the next section.

It has been observed that for assembled communities, there are two classes of ecological steady states that can be obtained by the dynamics: ‘shielded’ and ‘vulnerable’ ecosystems. Shielded ecosystems are ones in which purely ecological selection pressures vanish (e.g. the Kullback Leibler divergence above vanishes because the normalized demand exactly matches the supply). Vulnerable ecosystems are ones in which there can still be substantial ecological selection pressure due to a mismatch in the supply and the demand.

Eco-evolutionary dynamics of the Lyapunov function

We have discussed what happens to the Lyapunov function for a single optimum, but in an evolutionary scenario, the value of the Lyapunov function will always increase. This can be understood by introducing a type at low frequency, which leaves the Lyapunov function unchanged. The subsequent dynamics following the same set of equations, so that the type will only invade and reach finite abundance if its presence increases the Lyapunov function. In the process of invasion, a new type will either invade and replace zero or more strains. There are no ‘ecological suicide’-like effects where invasion causes extinctions of both residents and the invader or ‘rock-paper-scissor’ dynamics.

We can get a handle on the evolutionary process in the high dimensional limit by considering a state in which the demand almost exactly matches the supply. In this state, the equilibrium value of the Lyapunov function becomes:

$$H(\vec{f}^*) = \bar{X} - \sum_i \beta_i \log \frac{\beta_i}{\sum_\nu e^{X_\nu - \bar{X}^*} f_\nu^* \alpha_{\nu i}} \approx \bar{X}$$

From this state we can analyze changes in the Lyapunov function due to changes in the fitness and the strategy of a type ν (e.g. the gradient flow). To compute these quantities, we imagine a situation in which we suddenly infinitesimally perturb one of the strain’s phenotypes, *without running the dynamics*. This results in an infinitesimal change in the Lyapunov function without distorting the frequencies.

For a small fitness change, the infinitesimal change in the Lyapunov function is:

$$\delta H = \delta X_\nu f_\nu^* \left(1 + \vec{\beta} \cdot \vec{\alpha}_\nu \right)$$

So that we can see that strategies closest to the supply vector (encoded by the dot product) increase H the most. Note that the maximum such perturbation can differ from the maximum ‘invasion fitness’ which governs the evolutionary process. Therefore, while evolution increases the Lyapunov function over a trajectory, it does not necessarily follow the steepest direction of the gradient.

Similarly, for strategy mutations, we have that an infinitesimal strategy perturbation gives:

$$\delta H = e^{X_\nu - \bar{X}^*} f_\nu^* \delta \vec{\alpha}_\nu \cdot \vec{\beta}$$

where, once again the strength (and now sign) of the perturbation depends on its relation to the supply vector. Note that the strength of the perturbation also depends on the fitness of the strain background. Once again, this differs from the direction taken by the evolutionary gradient.

Now let us consider evolution on a fixed realization of disorder. This results in infinitesimal changes of the form:

$$\delta H = f_\nu^* \left[\delta X_\nu + \left(\delta X_\nu \vec{\alpha}_\nu + e^{X_\nu - \bar{X}^*} \delta \vec{\alpha}_\nu \right) \cdot \vec{\beta} \right]$$

Taken together, we can see that for this joint process (by which strategy and fitness change simultaneously) on a finite strategy space, the evolutionary process will terminate in the absence of pure fitness evolution, but potentially at a local optimum of the Lyapunov function. Furthermore, we can assert that in the presence of fitness mutations, there is no longer a strict optimum. Rather, optima of the disordered landscape are now all saddles with an unstable direction along the fitness axis.

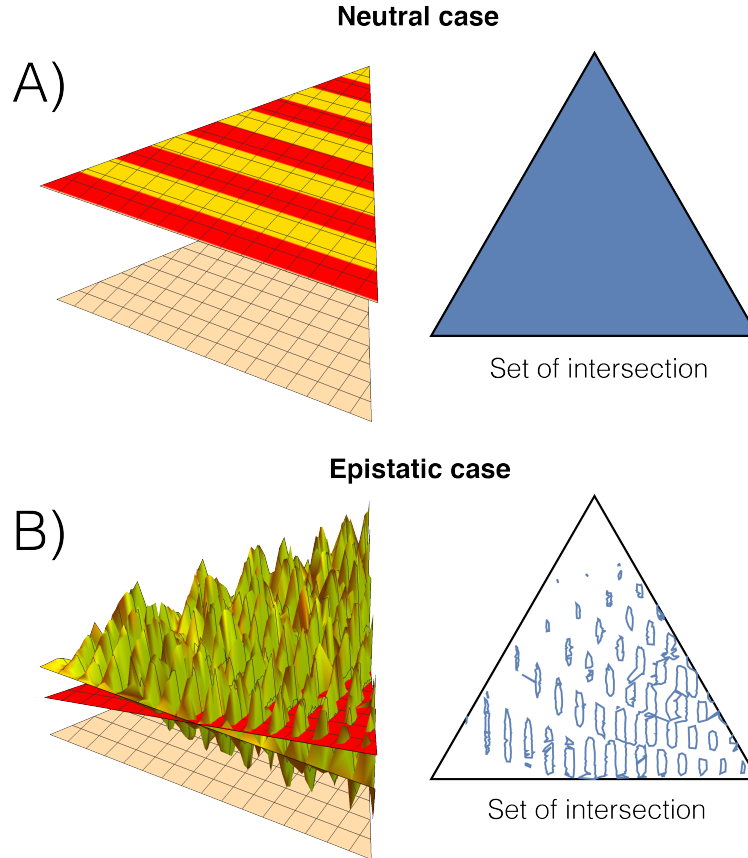


Figure 3.7: The necessary coexistence condition for a set of ecotype strategies in the consumer resource model can be understood geometrically. The set of allowed strategies can be found by projecting the intersection of a plane with the fitness function onto the strategy simplex. In the neutral case, this fitness function is a plane, so the intersection is the whole of the simplex. The intersection is almost certainly a set of vanishing measure for a sufficiently rugged function (as illustrated above) so that randomly chosen species are not likely to coexist at high diversity. However, there are many possible intersecting planes that can be chosen based on the additive fitness and strategies so that many different combinations of species can be made to coexist. A modified version of the convex hull condition from [29] must also be satisfied for sufficiency and is discussed in the SI.

High diversity states in the consumer-resource model with rugged metabolic tradeoffs

The consumer resource model, as formulated, possesses the interesting property that when $X_\mu = X_\nu \ \forall \mu, \nu$, many more species can coexist than the number of available resources, seemingly breaking the constraints set by competitive exclusion [29]. The condition for this anomalously high coexistence is that the nutrient supply $\vec{\beta}$ be within the convex hull of the set metabolic strategy vectors of all the competing types $\{\vec{\alpha}_\mu\}$. Randomly drawn fitness differences break this symmetry, and the limit set by competitive exclusion, $\mathcal{S} \leq \mathcal{R}$, is restored.

The number of coexisting species \mathcal{S} can be understood simply from the condition that each such species must have a net growth rate of 0 at equilibrium. This can be translated into the condition for the resource availabilities \vec{A} s.t. $A_i \equiv \beta_i / \sum_\nu \alpha_{\nu i} e^{X_\nu} f_\nu$ that:

$$\vec{\alpha}_\mu \cdot \vec{A} = e^{-X_\mu} \ \forall \mu$$

When all the X_μ are equal (or equivalently $X_\mu = 0$), this condition can be satisfied for $\mathcal{S} > \mathcal{R}$ species. This can be understood by noticing that when fitness is 0 across ecotypes, the equation is solved by $\vec{A} = \vec{1}$ (or, geometrically, these equations define an $(\mathcal{R} - 1)$ -dimensional hyperplane). Since for every resource dimension, the availabilities $A_i = 1$, this condition can be interpreted as the fact that each resource supply is perfectly matched by the uptake of species present.

In order to remain self-consistent, an additional constraint must be met, namely the so-called ‘convex hull constraint’ [29]. This constraint specifies that for the above solution to be valid, the supply vector $\vec{\beta}$ must fall within the convex hull of the set of resource strategies $\{\vec{\alpha}_\mu\}$. This can once again be understood geometrically. If each resource supply β_i is equal to a convex combination of strategies, $\sum_\mu \alpha_{\mu i} f_\mu$, then it must be within the convex hull of $\{\vec{\alpha}_\mu\}$ by definition.

When there is random variation in the X_μ , this co-planarity can no longer be maintained for $\mathcal{S} > \mathcal{R}$. However, with an disordered fitness term, where the $X_\mu(\vec{\alpha}_\mu)$ are *quenched* random functions of the $\vec{\alpha}_\mu$, the strategies $\vec{\alpha}_\mu$ enter on both sides of the set of (now nonlinear) equations defining the ecological equilibrium:

$$\vec{\alpha}_\mu \cdot \vec{A} = e^{-X_\mu(\alpha_\mu)} \ \forall \mu$$

This defines the intersection (up to an arbitrary translation in fitness) of a plane and a potentially very complicated manifold (see Fig. 3.7 for a schematic). Since a solution to this equation is likely to be a set of vanishing measure, it is improbable for *randomly* chosen species to coexist beyond the limits of competitive exclusion. However, *particular* sets of many coexisting species can be found by solving for the intersection. For instance, in the case where fitness is split into an additive and disordered component, one can trivially set $X_{\mu, \text{additive}} = -X_{\mu, \text{disorder}}(\vec{\alpha}_\mu)$ to get back the effectively neutral model (i.e. a flat plane) and break competitive exclusion as before. More generally, when it is possible to *tune* X_μ and $\vec{\alpha}_\mu$

(for instance by genetic manipulations or by evolution), one can find many allowed equilibria with $\mathcal{S} > \mathcal{R}$, with relatively few constraints on the resource strategies of ecotypes. Moreover, for random disorder on a continuous strategy space, there are likely to be many possible equilibria with $\mathcal{S} > \mathcal{R}$ which are fundamentally different from the effectively neutral case, since there are possible ‘tilted’ hyperplanes such that $\vec{A} \neq \vec{1}$. In the cases where $\vec{A} \neq \vec{1}$, self-consistency requires a ‘tilted’ convex hull condition:

$$\frac{\beta_i}{A_i^*} \in \text{conv}(\{e^{X_\mu(\alpha_\mu)} \alpha_{\mu i}\})$$

We can get a handle on this by working in a slightly different space. At equilibrium, types have ‘weighted strategies’ given by:

$$\tilde{\alpha}_{\mu i} \equiv w_\mu = \frac{\alpha_{\mu i} e^{X_\mu}}{\sum_\nu \alpha_{\nu i} e^{X_\nu} f_\nu^*}$$

These must satisfy:

$$\vec{w}_\mu^* \cdot \vec{\beta} = 1 \quad \forall \mu$$

so that in order to coexist, the weighted strategies must fall on the hyperplane in \mathcal{R} -dimensional real space that is perpendicular to the resource vector $\vec{\beta}$.

However, despite the possibility of diverse $\mathcal{S} > \mathcal{R}$ states including ones that might be highly skewed due to fitness differences, all such states are unstable to additional fitness differences imparted by continued evolution (as is the case in the neutral model). Once again this can be understood geometrically. If a type increases its fitness, this will tilt the hyperplane that defines coexistence, so that some (potentially several) types will be excluded from the new coexistence hyperplane that is reached by the fixed point of the dynamics unless all types match the fitness gain exactly. However, since mutations occur on a background at a rate proportional to that background’s frequency, this sort of exact matching dynamic is not likely with the evolutionary process we consider (e.g. fitness differences will always tend to break coexistence).

Evolving consumer-resource model

In this section we discuss the details of the evolutionary process that proceeds on the background of the consumer-resource model. We work in the strong selection, weak mutation (SSWM) limit.

We start with a resident population at equilibrium so that:

$$\sum_i \alpha_{\mu i} e^{X_\mu} A_i^* = 1 \quad \forall \mu$$

which determines the set of strain equilibrium frequencies $\{f_\mu^*\}$. We start by examining the behavior of a mutant type occurring on background ν such that:

$$(X_\nu, \vec{\alpha}_\nu) \rightarrow (X_\nu + \delta X, \vec{\alpha}_\nu + \delta \vec{\alpha})$$

Whether or not this mutant will invade is determined by its invasion fitness, or its growth rate when rare:

$$S_{\text{inv}} = \sum_i (\alpha_{\nu i} + \delta \alpha_i) e^{X_\nu + \delta X} A_i^* - 1 \approx \delta X + e^{X_\nu} \delta \vec{\alpha} \cdot \vec{A}^*$$

In the SSWM limit, invasion only occurs for mutations with positive invasion fitness. Note that this quantity is different from the gradient of the Lyapunov function, so while positive invasion fitness perturbations will increase the value of the Lyapunov function at equilibrium, the most likely evolutionary path will not necessarily follow the steepest gradient in Lyapunov function space.

Since we are in the SSWM limit, mutations occur as a compound Poisson process. The probability that a mutation will occur on a background μ is given by its frequency at equilibrium f_μ^* . In turn the probability that a mutation will *establish*, or survive genetic drift, is taken to be proportional to its invasion fitness, S_{inv} . We split the mutational process into two subprocesses, fitness mutation and strategy mutation.

Fitness mutations occur at overall rate U_X and are assumed to have effect size s drawn from some distribution of fitness effects (DFE) $\rho(s)$. Strategy mutations occur at rate U_α . Mutant strategies $\delta \vec{\alpha}$ are drawn from a distribution of strategy effects (DSE) $\rho(\delta \vec{\alpha})$ that maintains the normalization constraint on the strategy phenotype and does not force any components to be negative. The binary use architecture is just a special case of this process. The DFE is just a delta function with a positive value s and the DSE is a uniform distribution over the nearest neighbors on a $2^{\mathcal{R}}$ -dimensional hypercube (genotype space).

The different mutants arrive at rates proportional to their invasion fitness. Successful fitness mutants arrive at the following rate on background μ :

$$R_\mu^X = N U_X f_\mu^* (e^s - 1)$$

The rate of successful gain (+) and loss (−) of function mutants for resource i on background μ is given by:

$$R_{\mu,i}^\pm = \max \left\{ 0, N U_\alpha f_\mu^* \left(\frac{\vec{\alpha}_\mu \cdot \vec{h}^* \pm h_i \alpha_{\mu,i}}{1 \pm k_\mu^{-1}} - 1 \right) \Theta(\mp \alpha_{\mu,i}) \right\}$$

where $k_\mu = 1/\sum_i \alpha_{\mu,i}^2$ is the number of resources consumed by μ and where Θ is Heaviside step function ensuring that the strategy component is there to lose or gain. Since these mutations arrive as a compound Poisson process, we can use the properties of such processes to get the rate of the next arrival:

$$R_{\text{tot}} = \sum_{\mu} \left[R_{\mu}^X + \sum_i (R_{\mu,i}^+ + R_{\mu,i}^-) \right]$$

A particular invader is chosen randomly according to the weights given by their individual rates. After a mutation invades, the ecosystem equilibrates to a new steady state. This equilibration process is assumed to occur much faster than the mutation arrival process, which is valid in the limit where the population size is very large.

Model simulations

The model is initialized with metabolism matrix J , which is symmetric and drawn from $J_{ij} \sim \text{Normal}(0, \sigma)$. The initial condition is that in which the system is seeded with a single generalist strain. This is similar to the situation typical of laboratory evolution experiments, in which well-adapted strains are propagated in consistent environmental conditions for many generations. We also started with random initial conditions (i.e. assembled communities) but this did not change our qualitative results.

Starting from our initial generalist strain, mutations arrive according to the dynamics described above. Once a mutant arrives, the new ecological equilibrium is found by solving for the resource availabilities via the following convex optimization problem:

$$\vec{A}^* = \text{argmax}_{\vec{A}} \left\{ \sum_{i=1}^{\mathcal{R}} \beta_i \log A_i \text{ s.t. } \sum_{i=1}^{\mathcal{R}} \alpha_{\mu,i} A_i = e^{-X_\mu} \forall \mu \right\}$$

This solution was found using the Mosek solver package. The availabilities can then be inverted to find the frequencies using constrained linear regression. This process is repeated many times and the state of the system is saved at each step, along with the arrival time of each mutant.

Disordered fitness

We consider cases in which the resource strategies encode “rugged” fitness benefits and costs, calling this the disorder term of the fitness. We can formally expand the disorder term:

$$X_\mu = X_\mu^{\text{additive}} + X^{\text{disorder}}(\vec{\alpha}_\mu) = X_\mu^{\text{additive}} + \sum_i h_i \gamma_{\mu i} + \sum_{i,j} J_{ij} \gamma_{\mu i} \gamma_{\mu j} + \dots$$

where we rescale the strategy vector components to take on values of ± 1 :

$$\gamma_{\mu i} = 2 \cdot \|\vec{\alpha}_\mu\|_2^{-2} \alpha_{\mu i} - 1$$

This is simply a Fourier expansion of the function X_μ^{disorder} , and is completely general. We choose to truncate the infinite sum at quadratic order, since this is the minimal order at which the strategy landscape becomes ‘rugged.’ We neglect first order terms (i.e. set $h_i = 0$) since these do not contribute to the ruggedness. Although the generalist can still lose stability when only additive terms are present, the strategy landscape will have a single fitness peak, which does not result in qualitatively new dynamics in the slow fitness regime. We draw the coupling coefficients J_{ij} from a Gaussian distribution with zero mean and width $\frac{\sigma}{\sqrt{2\mathcal{R}(\mathcal{R}-1)}}$, with parameter σ quantifying the strength of disorder.

We can also consider the fully uncorrelated disorder model, in which a random fitness is assigned to an entire strategy, instead of being limited to pairs of components. This is equivalent to taking the full sum above, which in some cases can be more theoretically tractable than the pairwise model and acts as a useful limiting case. In our model, the function X^{disorder} takes on independent random values for each strategy:

$$X^{\text{disorder}}(\vec{\alpha}) \sim \mathcal{N}(0, \sigma)$$

Note, there is no need to normalize random draws as in the case of the pairwise model.

These are two ‘extremal’ models of rugged landscapes, respectively the maximally correlated and maximally uncorrelated cases. The core qualitative result of our model – that the generalist is unstable – will be independent of which extreme we choose. This is true so long as there is a finite collection of single knockouts with sufficiently high fitness advantages that further mutation does not happen on the generalist background.

Frequency of the generalist with knockouts

Single knockout and upper bound on σ_c

We have a generalist type which obtains a loss of function mutant on resource k . We have that the ecological equilibrium is given by:

$$\frac{\beta_k}{f_g} + \sum_{i \neq k} \frac{\beta_i}{f_g + \frac{\alpha_{ki}}{\alpha_{gi}} e^{X_k - X_g} (1 - f_g)} = 1$$

Alternately, we have the following equivalent condition from the dynamical equation for the knockout type:

$$\sum_{i \neq k} \frac{\beta_i}{(1 - f_g) + \frac{\alpha_{gi}}{\alpha_{ki}} e^{X_g - X_k} f_g} = 1$$

We have that the strategy vector components for the generalist and the loss of function variant are given, respectively, by $\alpha_{gi} = 1/\mathcal{R}$ and $\alpha_{ki} = 1/(\mathcal{R} - 1)$. We are free to set the fitness scale due to the symmetry of the model, so WLOG we set $X_g = 0$. So the equilibrium condition is given by:

$$\sum_{j \neq i} \frac{\beta_j}{(1 - f_g) + \frac{\mathcal{R}-1}{\mathcal{R}} e^{X_k} f_g} = \frac{1 - \beta_i}{(1 - f_g) + \frac{\mathcal{R}-1}{\mathcal{R}} e^{-X_k} f_g} = 1$$

Therefore the frequency of the generalist is given by:

$$f_g = \frac{\beta_i}{1 - \frac{\mathcal{R}-1}{\mathcal{R}} e^{-X_k}}$$

From this expression we can notice a couple of things. First of all, for neutral mutations, the frequency of the generalist is given by:

$$f_g = \mathcal{R}\beta_i$$

so that in order for a knockout to persist, it must happen on a resource less abundant than average (i.e. so that $f_g < 1$). If the generalist has a fitness advantage over the knockout, we see that the fitness difference that will no longer allow for diversification is given by:

$$X_g^c = \log \left(\frac{\mathcal{R}(1 - \beta_i)}{\mathcal{R} - 1} \right)$$

We also notice that for a highly fit knockout, the frequency of the generalist is just:

$$f_g = \beta_i$$

so that a generalist type cannot go extinct from a single knockout alone.

Instead continued mutation must occur for the generalist to go extinct. There are two subsequent paths to extinction. First is that further knockout types occur on the generalist background. We discuss this case in the next section since it is more complicated to analyze. The second, simpler case, is that fitness mutations accrue on the knockout background and drive the generalist to extinction.

This will happen when the generalist frequency is less than $1/2$. This requires that the fitness of the knockout mutant be at least:

$$X_k^c = \log \left[\frac{\mathcal{R} - 1}{(1 - 2\beta_i)\mathcal{R}} \right]$$

In the many resource limit, this reduces to:

$$X_k \gtrsim 2\beta_i$$

In the RPEM, the fitness difference between the loss of function type and the generalist will be the sum of $\mathcal{R} - 1$ IID Gaussian variables with variance $\frac{\sigma^2}{\mathcal{R}(\mathcal{R}-1)}$ so that it will be distributed as a Gaussian with mean zero and variance σ^2/\mathcal{R} . In this model σ^2 is the variance in fitness of the possible genotype pool. The maximum of all the possible single loss of function types will scale as:

$$\max(X_l - X_g) \sim \sqrt{2\sigma^2 \log \mathcal{R}}$$

The critical strength of disorder such that the first loss of function type has a higher frequency than the generalist is given by:

$$\sigma^* \sim \frac{\log \left[\frac{\mathcal{R}-1}{(1-2\beta_i)\mathcal{R}} \right]}{\sqrt{2 \log \mathcal{R}}} \sim \frac{1}{\mathcal{R}\sqrt{\log \mathcal{R}}}$$

In the large \mathcal{R} limit, this critical disorder upper bound gets small, so that small disorder effects can strongly impact large ecosystems that evolve.

Frequency of the generalist with two knockouts

We have that there are three nonlinear equations that define the equilibrium frequencies when we have a generalist with two single knockouts. They are given by:

$$\begin{aligned} \frac{\beta_{k_1}}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} e^{X_{k_2}} f_{k_2}} + \frac{\beta_{k_2}}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} e^{X_{k_1}} f_{k_1}} + \frac{1 - \beta_{k_1} - \beta_{k_2}}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} (e^{X_{k_1}} f_{k_1} + e^{X_{k_2}} f_{k_2})} &= 1 \\ \frac{\beta_{k_2}}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} e^{X_{k_1}} f_{k_1}} + \frac{1 - \beta_{k_1} - \beta_{k_2}}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} (e^{X_{k_1}} f_{k_1} + e^{X_{k_2}} f_{k_2})} &= \frac{\mathcal{R} - 1}{\mathcal{R}} e^{-X_{k_1}} \\ \frac{\beta_{k_1}}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} e^{X_{k_2}} f_{k_2}} + \frac{1 - \beta_{k_1} - \beta_{k_2}}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} (e^{X_{k_1}} f_{k_1} + e^{X_{k_2}} f_{k_2})} &= \frac{\mathcal{R} - 1}{\mathcal{R}} e^{-X_{k_2}} \end{aligned}$$

where we use the labels k_1 and k_2 to indicate both the knockout strains and the resource dimensions that they *do not* utilize.

We can define the auxiliary variables:

$$\begin{aligned} x &\equiv \frac{1}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} e^{X_{k_2}} f_{k_2}} \\ y &\equiv \frac{1}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} e^{X_{k_1}} f_{k_1}} \\ z &\equiv \frac{1}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} (e^{X_{k_1}} f_{k_1} + e^{X_{k_2}} f_{k_2})} \end{aligned}$$

so that the equilibrium conditions are now:

$$\begin{aligned} \beta_{k_1} x + \beta_{k_2} y + (1 - \beta_{k_1} - \beta_{k_2}) z &= 1 \\ \beta_{k_2} y + (1 - \beta_{k_1} - \beta_{k_2}) z &= \frac{\mathcal{R}-1}{\mathcal{R}} e^{-X_{k_1}} \\ \beta_{k_1} x + (1 - \beta_{k_1} - \beta_{k_2}) z &= \frac{\mathcal{R}-1}{\mathcal{R}} e^{-X_{k_2}} \end{aligned}$$

We can solve for x , y and z to obtain:

$$\begin{aligned} x &= \frac{1 - \frac{\mathcal{R}-1}{\mathcal{R}} e^{-X_{k_1}}}{\beta_{k_1}} \\ y &= \frac{1 - \frac{\mathcal{R}-1}{\mathcal{R}} e^{-X_{k_2}}}{\beta_{k_2}} \\ z &= \frac{\frac{\mathcal{R}-1}{\mathcal{R}} (e^{-X_{k_1}} + e^{-X_{k_2}}) - 1}{1 - \beta_{k_1} - \beta_{k_2}} \end{aligned}$$

Finally, we notice that the frequency of the generalist is simply:

$$f_g = x^{-1} + y^{-1} - z^{-1}$$

We see that the generalist goes extinct around a similar point in this case with a similar scaling given by the $f_g = 1/2$ condition. This is because in the many resource limit, the top two knockouts (with highest invasion fitness) should only weakly differ.

Frequency of the generalist with n knockouts

We examine the case where we have a generalist with $n \leq \mathcal{R} - 1$ single knockouts. As in the two knockout case, we define auxiliary variables:

$$\begin{aligned} x_k &\equiv \frac{1}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} \sum_{i \neq k} e^{X_i} f_i} \\ z &\equiv \frac{1}{f_g + \frac{\mathcal{R}}{\mathcal{R}-1} \sum_{i=1}^n e^{X_i} f_i} \end{aligned}$$




Figure 3.8: Scaling in the frequency of the generalist when in competition with 2 most fit single knockouts. The extinction of the generalist scales with approximately $1/\mathcal{R}\sqrt{\log \mathcal{R}}$

so that we have:

$$x_k = \frac{1 - \frac{\mathcal{R}-1}{\mathcal{R}}e^{-X_k}}{\beta_k}$$

$$z = \frac{1 - \sum_{i=1}^n \left(1 - \frac{\mathcal{R}-1}{\mathcal{R}}e^{-X_i}\right)}{1 - \sum_{i=1}^n \beta_i}$$

with the frequency of the generalist given by:

$$f_g = \sum_{k=1}^n x_k^{-1} - (n-1)z^{-1}$$

Scaling with $\mathcal{R} - 1$ knockouts

In the main text we looking at the one and maximal number of knockouts with Gaussian pairwise disorder. However, since some of the knockouts have fitness disadvantages, not all can invade. To see how the stability of the generalist scales with even more knockout types, we look at the case with the maximal number of single knockouts ($\mathcal{R} - 1$). This corresponds to a maximally unstable case for the generalist when there is quenched disorder and only single knockouts are allowed to invade. To ensure that all knockouts have beneficial fitness effects, we choose the distribution of knockout fitnesses to be exponential with scale σ for simplicity.

When we look at this case, we see that the generalist frequency rapidly goes to a limiting form as a function of the renormalized fitness scale, $\tilde{\sigma} = \sigma\mathcal{R}^3$. The generalist goes extinct at critical value $\tilde{\sigma}_c = 1$, so that in the competition for more resources, the generalist will go extinct for much lower bare σ_c . In the limit, this critical scale goes to zero.

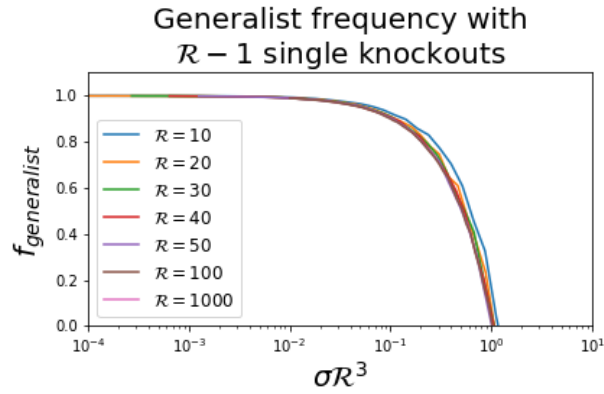


Figure 3.9: Phenomenological scaling in the frequency of the generalist when in competition with knockouts with random fitness advantages. The generalist goes extinct for smaller and smaller fitness differences as the number of resources goes to infinity. The frequency curve scales even more extremely than the single knockout case.

Relation to Tikhonov and Monasson (2018)

There are some key differences between the evolutionary model we propose and that proposed in Tikhonov and Monasson (2018) (hereafter referred to as TM18) [34]. The model proposed therein is a related evolutionary model. However, there are several important differences. One important caveat is that in TM18 mutational steps do not have a sense of locality. In other words, a mutant with a strategy can mutate into any other strategy with any fitness in a single step. We don't expect that this is realistic in a microbial population since getting any particular genomic rearrangement from an evolutionary process seems extremely unlikely.

In addition, mutations are not partitioned into strategy and fitness types. In our model, this separation, allows us to tease apart the different roles of strategy and fitness, leading us to different conclusions. In the TM18 model, since strategy and fitness mutations occur at the same time, this obscures the crucial role that fitness mutations play in breaking 'gridlocks' in the ecological dynamics.

Furthermore, in TM18 there is only a weak sense of inheritance (e.g. linkage). This leads to very qualitatively different dynamics and conclusions from their model and from ours. This highlights the importance of constraining the types of mutational kernels that might be relevant in laboratory and natural populations.

3.8 Bibliography

- [1] S. G. Acinas, V. Klepac-Ceraj, D. E. Hunt, C. Pharino, I. Ceraj, D. L. Distel, and M. F. Polz. Fine-scale phylogenetic architecture of a complex bacterial community. *Nature*, 430(6999):551–554, 2004.
- [2] J. Allen, W. Schaffer, and D. Rosko. Chaos reduces species extinction by amplifying local population noise. *Nature*, 364(6434):229–232, 1993.
- [3] Z. D. Blount, C. Z. Borland, and R. E. Lenski. Historical contingency and the evolution of a key innovation in an experimental population of *escherichia coli*. *Proceedings of the National Academy of Sciences*, 105(23):7899–7906, 2008.
- [4] J. Boocock, M. J. Sadhu, J. S. Bloom, and L. Kruglyak. Ancient balancing selection maintains incompatible versions of a conserved metabolic pathway in yeast. *bioRxiv*, 2019. doi: 10.1101/829325.
- [5] G. Bunin. Ecological communities with lotka-volterra dynamics. *Physical Review E*, 95(4):042414, 2017.
- [6] R. A. Caetano, Y. Ispolatov, and M. Doebeli. Evolution of diversity in metabolic strategies. *bioRxiv*, 2020.
- [7] P. Chesson. Mechanisms of maintenance of species diversity. *Annual review of Ecology and Systematics*, 31(1):343–366, 2000.
- [8] N. Frazão, A. Sousa, M. Lässig, and I. Gordo. Horizontal gene transfer overrides mutation in *escherichia coli* colonizing the mammalian gut. *Proceedings of the National Academy of Sciences*, 116(36):17906–17915, 2019.
- [9] B. H. Good, I. M. Rouzine, D. J. Balick, O. Hallatschek, and M. M. Desai. Distribution of fixed beneficial mutations and the rate of adaptation in asexual populations. *Proceedings of the National Academy of Sciences*, 109(13):4950–4955, 2012.
- [10] B. H. Good, M. J. McDonald, J. E. Barrick, R. E. Lenski, and M. M. Desai. The dynamics of molecular evolution over 60,000 generations. *Nature*, 551(7678):45–50, 2017.
- [11] B. H. Good, S. Martis, and O. Hallatschek. Directional selection limits diversification and promotes ecological tinkering in the competition for many resources. *PNAS*, 2018.
- [12] S. J. Gould and N. Eldredge. Punctuated equilibria: the tempo and mode of evolution reconsidered. *Paleobiology*, 3(2):115–151, 1977.
- [13] G. E. Hutchinson. The paradox of the plankton. *The American Naturalist*, 95(882):137–145, 1961.

- [14] N. Kashtan, S. E. Roggensack, S. Rodrigue, J. W. Thompson, S. J. Biller, A. Coe, H. Ding, P. Marttinen, R. R. Malmstrom, R. Stocker, et al. Single-cell genomics reveals hundreds of coexisting subpopulations in wild *prochlorococcus*. *Science*, 344(6182): 416–420, 2014.
- [15] S. Kauffman and S. Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of theoretical Biology*, 128(1):11–45, 1987.
- [16] Kullback-Leibler divergence. Kullback-Leibler divergence — Wikipedia, the free encyclopedia, 2021. URL https://en.wikipedia.org/wiki/Kullback%E2%80%93Leibler_divergence. [Online; accessed 10-August-2021].
- [17] J. Ma and S. A. Levin. The evolution of resource adaptation: how generalist and specialist consumers evolve. *Bulletin of mathematical biology*, 68(5):1111–1123, 2006.
- [18] R. MacArthur. Species packing and competitive equilibrium for many species. *Theoretical population biology*, 1(1):1–11, 1970.
- [19] A. Mayer, V. Balasubramanian, T. Mora, and A. M. Walczak. How a well-adapted immune system is organized. *Proceedings of the National Academy of Sciences*, 112(19): 5950–5955, 2015.
- [20] J. O. McInerney, A. McNally, and M. J. O’connell. Why prokaryotes have pangenomes. *Nature microbiology*, 2(4):1–5, 2017.
- [21] M. E. Muscarella and J. P. O’Dwyer. Species dynamics and interactions via metabolically informed consumer-resource models. *bioRxiv*, 2019. doi: 10.1101/518449.
- [22] R. A. Neher and B. I. Shraiman. Competition between recombination and epistasis can cause a transition from allele to genotype selection. *Proceedings of the National Academy of Sciences*, 106(16):6866–6871, 2009.
- [23] R. A. Neher and B. I. Shraiman. Statistical genetics and evolution of quantitative traits. *Reviews of Modern Physics*, 83(4):1283, 2011.
- [24] A. Nourmohammad, J. Otwinowski, and J. B. Plotkin. Host-pathogen coevolution and the emergence of broadly neutralizing antibodies in chronic infections. *PLoS genetics*, 12(7):e1006171, 2016.
- [25] J. Oh, A. L. Byrd, M. Park, H. H. Kong, J. A. Segre, N. C. S. Program, et al. Temporal stability of the human skin microbiome. *Cell*, 165(4):854–866, 2016.
- [26] P. A. Orlando and S. R. Hall. How do generalist consumers coexist over evolutionary time? an explanation with nutrition and tradeoffs. *Theoretical ecology*, 8(3):383–398, 2015.

- [27] M. T. Pearce, A. Agarwala, and D. S. Fisher. Stabilization of extensive fine-scale diversity by ecologically driven spatiotemporal chaos. *Proceedings of the National Academy of Sciences*, 2020.
- [28] J. Plucain, T. Hindré, M. Le Gac, O. Tenaillon, S. Cruveiller, C. Médigue, N. Leiby, W. R. Harcombe, C. J. Marx, R. E. Lenski, et al. Epistasis and allele specificity in the emergence of a stable polymorphism in *escherichia coli*. *Science*, 343(6177):1366–1369, 2014.
- [29] A. Posfai, T. Tallefumier, and N. S. Wingreen. Metabolic trade-offs promote diversity in a model ecosystem. *Physical review letters*, 118(2):028103, 2017.
- [30] D. E. Rozen and R. E. Lenski. Long-term experimental evolution in *escherichia coli*. viii. dynamics of a balanced polymorphism. *The American Naturalist*, 155(1):24–35, 2000.
- [31] R. She and D. F. Jarosz. Mapping causal variants with single-nucleotide resolution reveals biochemical drivers of phenotypic change. *Cell*, 172(3):478–490, 2018.
- [32] T. F. Thingstad. Elements of a theory for the mechanisms controlling abundance, diversity, and biogeochemical role of lytic bacterial viruses in aquatic systems. *Limnology and Oceanography*, 45(6):1320–1328, 2000.
- [33] M. Tikhonov and R. Monasson. Collective phase in resource competition in a highly diverse ecosystem. *Physical review letters*, 118(4):048103, 2017.
- [34] M. Tikhonov and R. Monasson. Innovation rather than improvement: A solvable high-dimensional model highlights the limitations of scalar fitness. *Journal of Statistical Physics*, 172(1):74–104, 2018.
- [35] D. Tilman. Constraints and tradeoffs: toward a predictive theory of competition and succession. *Oikos*, pages 3–15, 1990.
- [36] S. Wang, J. Mata-Fink, B. Kriegsman, M. Hanson, D. J. Irvine, H. N. Eisen, D. R. Burton, K. D. Wittrup, M. Kardar, and A. K. Chakraborty. Manipulating the selection forces during affinity maturation to generate cross-reactive hiv antibodies. *Cell*, 160(4):785–797, 2015.
- [37] M. J. Wiser, N. Ribeck, and R. E. Lenski. Long-term dynamics of adaptation in asexual populations. *Science*, 342(6164):1364–1367, 2013.

Chapter 4

Eco-evolutionary feedback can stabilize multi-strain predator-prey communities

Preface

In the previous two chapters there is a significant restriction in our analysis – we specifically considered the strong selection, weak mutation regime in order to maintain some degree of tractability. While this can certainly be a relevant regime in real systems, it is unlikely to be the relevant regime in microbial systems, for instance. In order to move away from this, I opted to analyze a different model system – one with predator-prey interactions. Instead of starting from a mean-field description and backing out the relevant eco-evolutionary dynamics in an ad-hoc manner, I started from a fully stochastic eco-evolutionary model and showed that the mean-field description is the relevant one when the population size and genome length become large. I also suggest that the model can be mapped onto a well-known physics model of non-equilibrium phase transitions. This leads to predictions about the stability of *evolving* predator-prey ecosystems, which, according to canonical results, should be unstable [21] and even chaotic [29] in the high diversity limit. However, in the evolving case, highly diverse ecosystems are more likely to be stable over long times.

The remainder of this chapter will be submitted for publication in edited form.

4.1 Abstract

Ecological models with random interactions have provided insight into the problem of diversity, particularly showing that high variance in the distribution of interaction rates can lead to instability, chaos and extinction. However, these models have traditionally neglected evolution, which is central to the generation of biological variation and can act on timescales comparable to ecological change. Here we demonstrate that when an unstable stochastic predator-prey system is coupled to high-dimensional evolutionary dynamics, highly variable interactions counter-intuitively stabilize the population, delaying extinction and increasing the total population size. Using both stochastic and deterministic simulations and theory based on the statistical physics of disordered systems, we show that this stabilizing effect is driven by an eco-evolutionary feedback loop which causes the population size to grow as a function of the variance of the interaction rates. We demonstrate that the stable regime corresponds with the clonal interference regime of population genetics. Importantly, we expect qualitative aspects of our results to generalize to other evolving complex systems.

4.2 Main

Almost 50 years ago, the late Sir Robert May posed the question "Will a large complex system be stable?" [21] He answered his own question with a resounding "No." His work mathematically formalized a core problem in theoretical ecology – how can there be so much diversity (down to the lowest taxonomic levels [17]) and how might it persist in the face of strong competition, stochasticity, and apparent chaos? In this work, we show that evolutionary dynamics can act as a key stabilizing component in model ecosystems. Importantly, since evolution is coupled to the birth of new individuals, it can act on the timescales comparable to ecological change. When evolution is sufficiently fast, evolutionary and ecological processes enter a stabilizing feedback loop that is driven by underlying phenotypic diversity.

This work enters into a long line of theoretical work to characterize the stability of large, complex ecological models (which are often cast as variants of the generalized Lotka-Volterra equation). We will partially summarize some of this work. About 50 years ago, May and others recognized that large Lotka-Volterra systems (with or without randomly drawn coefficients) readily admit chaotic dynamics [21, 34], which would potentially lead to cascades of extinctions. May's key result is that, for a broad class of disordered ecological models, the variance of the distribution of interaction rates controls the stability of the dynamics – the more variable the interactions, the less stable an ecosystem might be expected to be [21].

Since then, there have been many attempts to propose mechanisms that might stabilize ecological models with highly variable interactions. These mechanisms, while plausible, are typically finely tuned, or have particular structural requirements, and as such are unlikely to be completely general. Among those proposed are the 'Kill-the-Winner' hypothesis [35], spatio-temporally fluctuating environmental conditions [15], and higher-than-pairwise interactions [3]. Most recently, stable spatio-temporal chaos was proposed as a diversity-

preserving mechanism [29, 30]. Intriguingly, there has been some work to suggest that evolution might stabilize small and relatively simple ecological models [32, 36, 37, 1, 18], while other work suggests the opposite [33]. To the best of our knowledge, no work has connected evolution to the ‘large, complex’ case proposed by May or proposed an adequate theory of how the variation generated by evolution interacts with simple ecological interactions.

In order to address this gap, we developed a fully stochastic eco-evolutionary predator-prey model with randomly drawn interactions. The underlying ecological dynamics would be unstable on its own [29]. We coupled the ecological dynamics to an evolutionary process on a high-dimensional ‘genotype’ space. Using efficient stochastic simulations and theory based on the statistical physics of disordered systems, we show that in this context, evolution is able to stabilize the dynamics, but only for sufficiently *high* variance interactions – directly contradicting May’s classic result. We show that potentially complex underlying strain dynamics undergird this population-level stability. We also demonstrate that the transition to stability coincides with the ‘clonal-interference’ regime of population genetics. Finally, we discuss how our qualitative results might generalize to other eco-evolutionary scenarios.

4.3 Model description

We consider a system of many predators and prey in which every predator-prey pair has a randomly drawn interaction rate. We assume nutrient saturation and well-mixed conditions so that there is no explicit population size control for the predator or prey population. We do not expect that the addition of a carrying capacity will change our qualitative results. We also assume that the predation process is ‘lossy’ so that predators do not always reproduce when they consume prey. We make this choice so that the interaction matrix is not anti-symmetric. Anti-symmetric Lotka-Volterra models, which possess a conserved quantity, are stably chaotic at the mean field level [11], while non-anti-symmetric models exhibit diverging chaos [29]. Non-anti-symmetric interactions would also be generated by specifically biological phenomena like phage burst size, in which a single predation event results in the production of many phage particles. We do not attempt to explicitly model the phenomenon of burst size, opting instead to focus on more generic aspects of non-anti-symmetric predator-prey interactions.

The anti-symmetric component of the interaction between prey i and predator k is denoted a_{ik} . The a_{ik} are drawn from a distribution with mean a and standard deviation σ_a . The ‘lossy’ component is denoted s_{ik} and is drawn from a distribution with mean s and standard deviation σ_s . In the simulations in the main text, for ease of presentation we set $\sigma_a = \sigma_s = \sigma$ and $a = s = 10^{-3}$. In the most general case, we assume that each prey has a birth rate, b_i and each predator has a death rate, d_k . In our simulations we set $b_i = d_k = 1$. We discuss modeling choices and variations of the model in the Appendix.

In addition, every predator and every prey is indexed by bitstrings of length L_p and L_b , respectively, so that there are 2^{L_p} predator and 2^{L_b} prey types in the model. In all that follows we consider $L_p = L_b = L$. We refer to the different values of the bitstrings as




Figure 4.1: A schematic biological interpretation of the population dynamic model we employ. A) The systems we consider are well-mixed and consist of two ‘species’ which have many strains. B) Interactions between the two types of species are mediated by bitstring ‘genotypes.’ Interaction rates are drawn from a distribution. C) Mutations to the bitstring genotype only occur during events that result in the birth of a new individual. The probability of a mutation given an event is p .

‘genotypes.’ For each birth event (of predator and prey), there is a probability, p , that a mutation occurs (with probability $1 - p$ that a birth is clonal). We allow for mutations of ‘bitflip’ type so that only L mutant genotypes are accessible from a given genotype i (see Fig 4.1). We denote this set of accessible genotypes \mathcal{S}_i . We describe the reactions in our model in detail in Table 4.1.

Reaction	Description
$X_i \xrightarrow{(1-p)b_i} 2X_i$	Clonal prey birth
$Y_k \xrightarrow{d_k} \emptyset$	Predator death
$X_i + Y_k \xrightarrow{s_{ik}} Y_k$	‘Lossy’ predation (no birth)
$X_i + Y_k \xrightarrow{(1-p)a_{ik}} 2Y_k$	Predation with clonal predator birth
$X_i \xrightarrow{pb_i/L} X_i + X_j \quad \forall j \in \mathcal{S}_i$	Prey birth with random mutation
$X_i + Y_k \xrightarrow{pa_{ik}/L} Y_k + Y_l \quad \forall l \in \mathcal{S}_k$	Predation with predator birth and random mutation

Table 4.1: The set of reactions in our stochastic predator prey model.

These stochastic dynamics lead to the following mean-field equations:

$$\begin{aligned}
\frac{dX_i}{dt} &= X_i \left(b_i - \sum_k (a_{ik} + s_{ik}) Y_k \right) + \frac{p}{L} \sum_{j \in \mathcal{S}_i} (b_j X_j - b_i X_i) \\
\frac{dY_k}{dt} &= Y_k \left(-d_k + \sum_i a_{ik} X_i \right) + \frac{p}{L} \sum_i \sum_{l \in \mathcal{S}_k} (a_{il} X_i Y_l - a_{ik} X_i Y_k)
\end{aligned}$$

The first terms on the right hand side of each equation are identical to the terms of a standard mean-field predator-prey model, while the subsequent terms are novel and due to mutation. These mutation terms take the form of diffusion operators on a hypercube (the graph connecting the bitstring genotypes).

Above, we have assumed that the mutation probability, p , is constant across both predator and prey genotypes and that accessible types are uniformly accessible, although both of these assumptions may be relaxed. Importantly, mutation is *explicitly coupled* to births (which in the case of the predator types only come about due to ecological interactions). An independent mutation rate, μ , is only consistent when p is very small so that evolutionary processes are slow and act on statistically steady ecological timescales. This is the realm of classical population genetics and so-called adaptive dynamics. We discuss models in which mutation is not correlated with birth (the ‘ μ -model’, as opposed to our ‘ p -model’) in the Supplement. The contents of this paper primarily deal with steady regimes in which evolutionary timescales are faster than ecological extinction timescales.

The form of the interactions and the genotypes can be motivated biologically. Often interactions between cell types are mediated by binding between proteins. For instance, we can imagine a phage-bacteria system in which the phage must recognize and bind to a receptor on the surface of a bacterium in order to continue its life cycle. Or we might consider an adaptive immune cell recognizing antigen. The bitstrings then represent an

encoding of these binding sites. The reaction rates vary because different pairs of binding site sequences have different binding affinities. Importantly, these rates are fixed by the physics of the interactions between different pairs of protein sequences. We assume that the different genotypes are closely related enough that these binding affinities are drawn from a single distribution with a well-defined mean and variance. However, single mutations can vastly change the binding affinity of proteins, which has been observed in many years of mutagenesis experiments [5, 19].

Explicitly modeling phenotypic correlations between closely related genotypes is possible within our framework. Correlations between a_{ik} and s_{ik} are almost certainly bound to exist in real systems and constrain the evolutionary landscape. We can view our uncorrelated interaction landscape as a correlated one that has been coarse-grained over functional types that are sufficiently closely accessible by mutation, but not so coarse-grained that the sign of interactions can change. As such, we expect many of our results to hold when there are still sufficiently large phenotypic differences between genotypically close phenotypes. We discuss some such cases in the Supplement. The primary result is that correlations, while more difficult to analyze and to simulate, only mildly affect our observations.

4.4 Population stability with strong selection, strong mutation

We simulated the stochastic dynamics outlined above and plotted the total prey and predator population sizes:

$$X_{tot} = \sum_i X_i, \quad Y_{tot} = \sum_k Y_k$$

Representative simulation trajectories of the total population sizes for $\sigma/a = 10^3$ and $\sigma/a = 0$ ($p = 0.1$ fixed) are shown in Fig. 4.2. In the neutral case, $\sigma = 0$, either the predator or prey population goes extinct on simulation relevant timescales. When we increase sigma, we find the surprising result that population-level fluctuations are reduced. We find that σ several orders of magnitude greater than the mean interaction strength damps population oscillations in both stochastic and deterministic mean-field simulations. Moreover, we find the counterintuitive result that the total population sizes *grow* with increased disorder when compared to their neutral fixed point values (Fig. 4.3).

We also vary the mutation probability p , while holding $\sigma = 0.1$ fixed (Fig. 4.3). We see, unsurprisingly, for low mutation rates, that either predator or prey extinction is highly likely. However, at intermediate p , we see that sometimes the system gets locked in a low diversity state for long times. These happen by chance and will go extinct on an anomalously long timescale (they are metastable). As we increase p , amplitude fluctuations stabilize until we reach $p = 1$, when the population level dynamics fluctuate around a fixed point. While the $p = 1$ limit may not be natural (where all births result in mutations), it is a useful one to keep in mind when trying to understand the mechanism of the stabilizing feedback.

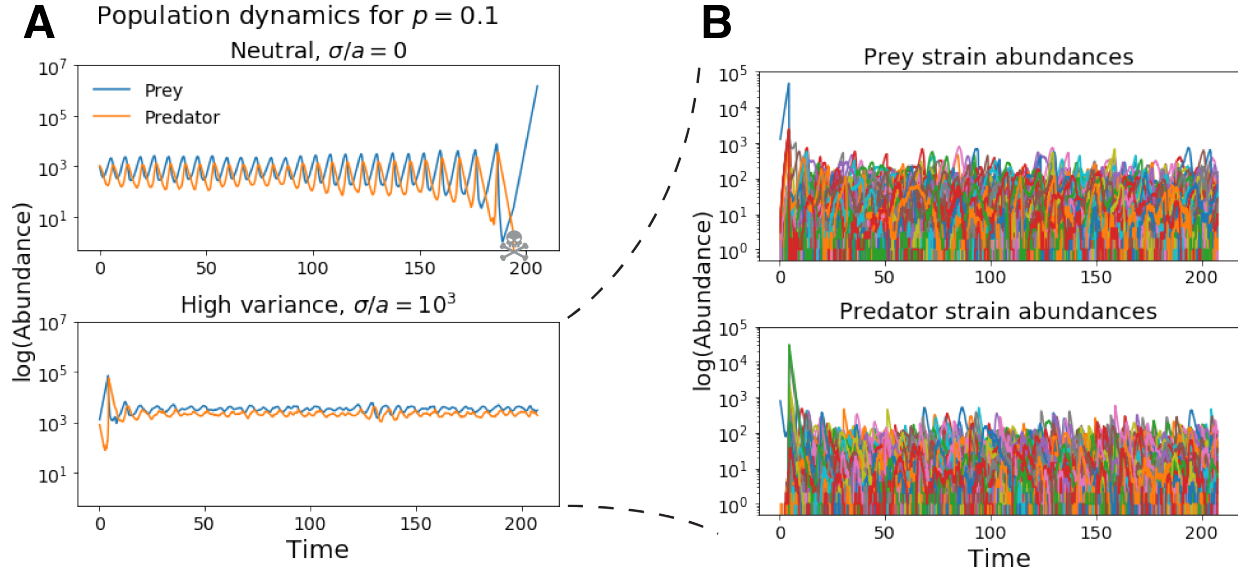


Figure 4.2: A) Stochastic simulations of the total predator and prey population size for different interaction spread, σ (plotted on a logarithmic scale). The neutral eco-evolutionary process ($\sigma_a = \sigma_s = 0$, $a = s = 10^{-3}$) leads to rapid extinction of either the predator or both populations. For highly variable interactions ($\sigma_a = \sigma_s = 1$, $a = s = 10^{-3}$), we see that the dynamics are stabilized. Both simulations have mutation probability $p = 0.1$. Extinction never occurred in the high variance case in over 100 runs with 2 million interactions. B) The strain dynamics underlying the stable population dynamics can be chaotic with constant extinction and recolonization by mutation.

When we look at the stochastic dynamics at the strain level, we see that strong fluctuations can undergird the apparent population stability in stochastic simulations. This is a consequence of the long time strain abundances being sufficiently close to zero that temporary extinction is highly probable. Despite the fact that strains rapidly come up and die off (and are repopulated), they can remain (on average) at finite abundances. However, if we decrease the average interaction strength (increase the average abundance), we see that strains wander around relatively fixed abundances (i.e. they approach the deterministic dynamics). Intuitively, this stabilization effect is due effective intraspecific interactions, which has been identified as a potentially stabilizing mechanism in other contexts. Because evolution is also a relevant process, this intraspecific diversity is continually renewed so that stability can be maintained even in the face of individual strains going extinct.

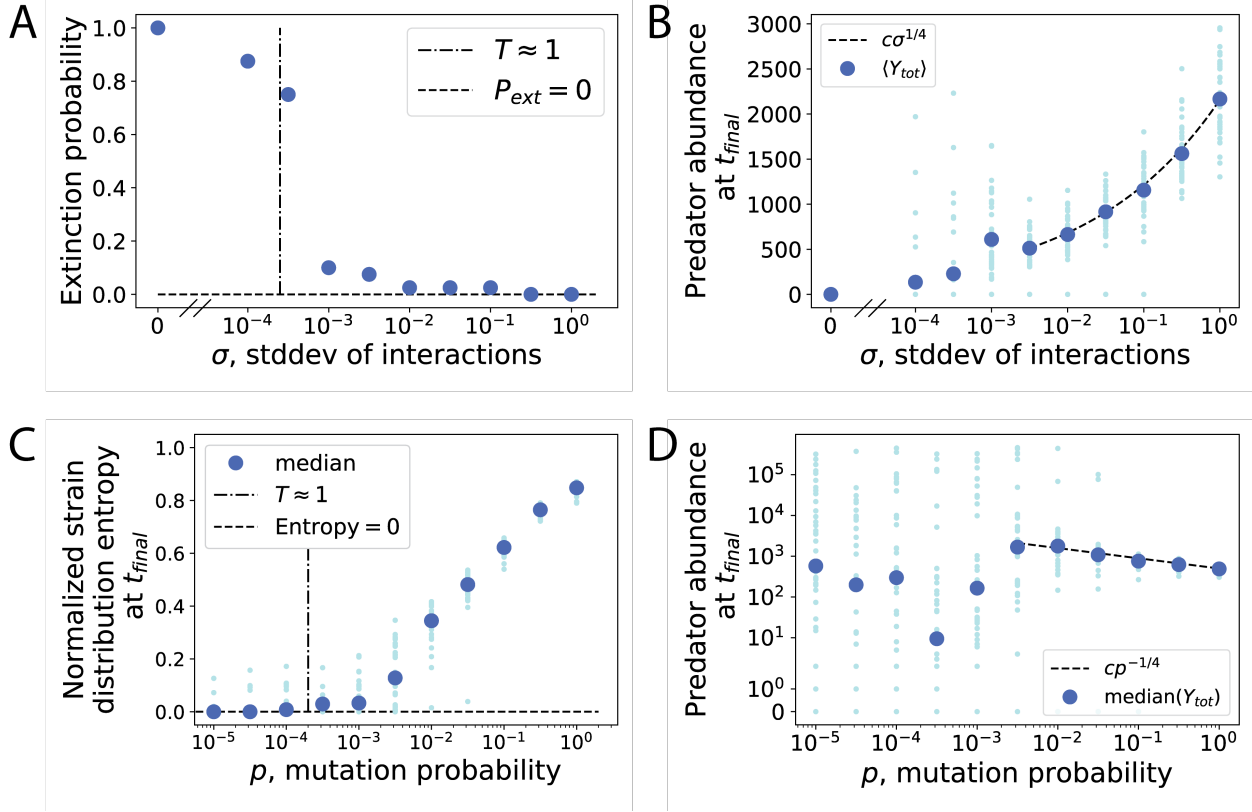


Figure 4.3: The transition between extinction and stability and the accompanying increase in population size. A) The extinction probability sharply drops when the effective ‘temperature’ is close to one. Near the transition, there are long-lived metastable states which lead to a ‘smearing’ effect. We expect many of the simulations to the left of the transition will go extinct given sufficiently long simulation time. Also not that the transition point is systematically underestimated because we drop subleading terms in our estimate of the mutation rate (see Supplement). B) The predator population size at the end of simulation time grows with the variance in the interactions. Large dark blue dots are averages over simulations. Small light blue dots are individual simulations. All averages are over 50 simulations. For both A and B $p = 0.1$. The prey population size grows with a similar exponent when conditioned on predator survival (if the predator goes extinct the prey population grows indefinitely). C) Entropy of the predator strain distribution (normalized by $\log(2^L)$) at the final time point of the simulation, varying the mutation probability, p . D) Predator abundance at the final time point of the simulation, varying the mutation probability, p . Both C and D have $\sigma = 0.1$. Because of the presence of long-lived metastable states we look at the median over realizations instead of the mean.

4.5 Total population sizes are stable in the many strain limit

From simulations, it is apparent that once in the stable phase, the total predator and prey population sizes are approximately constant. Moreover, the averages of these abundances increase with increasing disorder, σ . In order to understand this behavior, we look to the the mean field equations for X_{tot} and Y_{tot} :

$$\begin{aligned}\frac{dX_{tot}}{dt} &= \sum_i X_i \left(b_i - \sum_k (a_{ik} + s_{ik}) Y_k \right) \\ \frac{dY_{tot}}{dt} &= \sum_k Y_k \left(-d_i + \sum_i a_{ik} X_i \right)\end{aligned}$$

The total population size (even in the stable phase) is governed by the same mean field equations as the total population size without evolution (i.e. summing over all types, the evolution terms cancel). However, because there is flux between types due to mutation, this is not the complete picture, as is made clear by simulations. Instead, in the stochastic model we see an inhomogeneous distribution of types at a given time, whereas in the long time limit we expect that mutation (i.e. diffusion in genotype space) and birth-death noise homogenizes the average abundances.

We can formalize this by appealing to dynamical mean field theory [22, 29, 30] to decouple the interactions between predator and prey when the genotype space is large. We find that the mean field dynamics of the total predator and prey population sizes are quite simple. The attractor of these dynamics is a stable fixed point when the mutation probability and the variance of interactions is sufficiently high. This can be understood if we notice that predator mutation switches the sign of the feedback, which in turn stabilizes the deterministic dynamics. Heuristically, when a predator feeds on a highly abundant prey type, it likely produces mutants with sufficiently different feeding preferences so that the focal prey type's growth rate briefly *increases*.

We make an ansatz for the asymptotic behavior of the feedback response in the stable phase, so that the population size grows as a power law of the ratio of the variance of the interactions to the mutation probability. This is consistent with our simulations over a broad range of σ and p :

$$X_{tot} \sim Y_{tot} \sim (\sigma/p)^\eta$$

We argue that this power law behavior is generic rather than a sign of criticality in the Supplement. Moreover, by analyzing the deterministic equations, we are able to estimate the transition point from the extinct phase to stability. We define an ‘eco-evolutionary temperature:’

$$T \equiv \frac{paX_{tot}}{L}Y_{tot}$$

where a high temperature ($T \gg 1$) indicates stability. Using our scaling ansatz for the population sizes, we get consistent values for the transition point (see Figure 4.3).

Taking the limit in which mutation decouples from birth (the μ -model, discussed in the Supplement), we have that the beneficial mutation rate is given by $\mu \approx \frac{paX_{tot}}{L}$ so that the condition for stability is:

$$\mu Y_{tot} \gg 1$$

This condition defines the well-known clonal interference regime from population genetics [6]. In turn, our results can be translated into the language of population genetics, indicating that *clonal interference is crucial for the long term stability of large, diverse predator-prey populations*. Despite the connection to clonal interference, we find that in the stable phase, the behavior of the strain fitness distribution is inconsistent with traveling wave models that are commonly associated with clonal interference in the literature [13, 6, 7] (see Supplement). This is because the predator (prey) strain fitnesses are dependent on the composition of the prey (predator) population, and as such are dynamic quantities that can vary over many orders of magnitude.

4.6 Discussion

We have constructed a model of predator-prey coevolution on a high dimensional genotype space. Using the model, we have demonstrated that, contrary to common intuition, highly disordered interactions serve to stabilize populations when mutations are sufficiently frequent. While the population level dynamics are stable at long times, the strain level abundances can strongly fluctuate on short timescales due to the presence of a so-called Griffiths phase in finite populations. We have discussed the eco-evolutionary ‘phase diagram’ of the model and have provided parameter combinations that might be able to serve as indicators of eco-evolutionary resilience in predator-prey populations. In addition, with limited assumptions, we show that stability corresponds with the well-characterized clonal interference regime from population genetics. However, the population fitness distributions can be qualitatively different from established “traveling wave” theory when the phenotypes are highly diverse.

We expect that our qualitative results should hold for other mutational architectures and for phylogenetically correlated phenotypes, so long as viable strains are sufficiently well-connected (see Supplement). Furthermore, we expect that our qualitative results should hold for other stochastic ecological model choices, including ones that incorporate burst sizes, spatial structure, and other types of interactions. We also expect these results to be relevant to specific model systems including intrahost immune-pathogen dynamics [24] and epidemiological models of individuals’ immune response [37]. We leave these questions for future work.

Finally, we comment on the fact that our theory relies on the asymptotic behavior of our model. In that sense we expect that it should be widely applicable, though it is far from the end of the story. This is primarily a model of microdiversity, but diversity exists at many

scales, sometimes in intermediate (i.e. non-asymptotic) regimes where static phenomena like pattern formation or dynamic phenomena like traveling waves might be the relevant phenomenology in genotype space. Our model and simulations provide an exciting opportunity and base from which to explore these possibilities.

We conclude with a remark. Population modeling in biology often seeks low to intermediate-dimensional representations of dynamics, and in so doing brushes up against complexity and potential instability. However, as we have shown, if we take the opposite perspective and work with explicitly high-dimensional models there can be hope for emergent simplification.

4.7 Materials and Methods

Stochastic simulations

We simulated the model using a version of Gibson and Bruck’s ‘Next Reaction Method’ (NRM)[9], an exact stochastic simulation method with some performance advantages over the more commonly used Gillespie Algorithm (GA)[10]. The GA has an $\mathcal{O}(n)$ time complexity and while potentially faster in situations where the reaction propensities are multiscale (whereby one might order the vector of propensities in a clever way), our system does not exhibit clear multiscale behavior on long timescales. By exploiting data structures like priority queues and a reaction dependency graph, the NRM has a time complexity that scales as $\mathcal{O}(\log n)$ for a simulation with n reaction types. Therefore, it might provide substantial speedups when there are many reactions with sufficiently sparse dependencies, as is the case in our model. Since the NRM is an exact simulation scheme, there are no artefacts due to time discretization or other more ad hoc approximations (such as extinction cutoffs) and individual trajectories are guaranteed to be representative of the underlying stochastic process.

In our simulations, we set the prey birth rates and predator death rates to unity. We do not expect that slight departures from fixed birth and death rates will change our results significantly. We draw the interactions from a log-normal distribution with mean μ and scale σ . This choice guarantees that the total prey and predator populations will fluctuate around $1/\mu$ individuals in the mean field.

We can motivate the choice of distribution physically, based on the picture of interactions mediated by binding due to protein contacts. In this picture, binding events will depend on the interaction energy of a pair of sequences of amino acids. For randomly drawn sequences, this gives an interaction energy which is the sum of many identically distributed site-specific energies, so that the interaction distribution is Gaussian with some mean in the long sequence limit. The reaction rate can then be modeled as an Arrhenius law, so that the rates are log-normally distributed around a mean rate. However, our qualitative results should be independent of our particular choice of distribution so long as it has a well-defined mean and variance.

We varied the width of the interaction rate distribution (σ) and the mutation probability (p) over many orders of magnitude. The maximal genome length we were able to achieve was $L = 10$, due to memory constraints. We ran simulations for 2×10^6 reactions or until either the predator or prey class goes extinct, whichever came first. For simulations where strains are neutral with respect to each other (i.e. exchangeable, or $\sigma = 0$), 2×10^6 reactions is sufficient to observe extinction with high probability.

Deterministic simulations

We also simulated the deterministic dynamics. For these deterministic simulations, we implemented a 4th order Runge-Kutta scheme, with step size dependent on the number of strains, as determined by the genome length, L . For sufficiently many strains, the scheme was unstable unless prohibitively small step sizes were used. This can be understood from the fact that as more and more strains are added, the derivatives become larger and larger so that smaller step sizes are needed to maintain numerical stability. Due to these issues, we were only able to simulate the deterministic model up to $L = 8$. Comparisons to the stochastic model are discussed in the Appendix B.

Dynamical mean-field theory calculations

We employ dynamical mean field theory, a tool from the statistical physics of disordered systems, to analyze the dynamics of the predator and prey population sizes in the ergodic eco-evolutionary chaos phase. This technique has successfully been used for analyzing ecological scenarios without coupled evolutionary dynamics [31, 29]. Our calculations appear in Appendices C and D.

4.8 Supporting Information (SI)

Appendix A: Model choices

Predator-prey vs. generalized Lotka-Volterra models

It is reasonable to model an ecological system as a generalized Lotka-Volterra model (GLVM), with complex interlocking foodwebs, but it is less clear how this type of landscape would fit in to an eco-evolutionary system of the sort that we present, which models fine-scale within-species diversity. Within a random GLVM architecture, by allowing for mutations between strains, it would be possible for a prey individual to give birth to its own predator. This type of process is unlikely on the relatively short timescales represented by our model. Studying a model in which the sign of a mutant's interactions can differ from its parent's is certainly interesting theoretically, but care would need to be made to model this in a way that is biologically relevant. It would certainly be interesting to understand the eco-evolutionary process for more explicitly modeled interactions, which is conceptually possible.

The form of the interactions

We can motivate the interactions in our model through the lens of receptor-ligand binding kinetics. This can be cast as a string matching problem whereby the 'right' pairs of strings have a low energy (and a high interaction rate) and the 'wrong' pairs of strings have a large binding energy (and low interaction rates). A model of the form we present would be able to simulate an evolution experiment given deep binding assay data with K_d measurements for pairs of predator and prey types. These sorts of experiments are increasingly common with mutagenesis of single proteins. There are also some examples of mutagenesis of pairs of interacting proteins.

Appendix B: Heuristic derivation of the phase boundary

We analyze the boundary between the unstable and stable phases by appealing to heuristic arguments. We start with clonal predator and prey populations and work in the small $p \ll 1$ regime so that mutations are rare. Without mutation, it has been shown [27, 28] that the time to extinction of a predator prey model scales with the fixed point population sizes:

$$T_{ext} \sim N_s^{3/2} N_l^{-1/2}$$

where N_s is the smaller (i.e. predator) population size and N_l is the larger (i.e. prey) population size. For simplicity (and consistent with our simulations), we assume that these population sizes are approximately equal so that:

$$T_{ext} \sim N$$

We can then compare this timescale to the timescale on which a strongly beneficial mutation arrives, T_{mut} .

We call the probability that a mutation is strongly beneficial $q(\mu, \sigma, L)$, which depends on the mean of the interaction distribution, the width of the interaction distribution, and the genotype length L . Since N is set by μ , in general q will have some implicit N dependence, but we make the ansatz that this dependence is weak. In general q might be a complicated function, but we can conclude q should increase with both L and σ . In the limit $\sigma = 0$, there are no beneficial mutations so that $q = 0$. Similarly for $L = 0$. In the $\sigma \rightarrow \infty$ and $L \rightarrow \infty$ limits, q will approach 1, but potentially slowly.

We call the turnover rate of the population λ , which is the number of births per individual over the period we are concerned with, and is included so that we have the correct units. If we approximate the arrival of mutations as a homogeneous Poisson process, we have:

$$\int_0^{T_{mut}} N p \lambda q dt \sim 1$$

We can justify this assumption by averaging over the period of oscillation of the dynamics.

When these timescales are comparable, we have a condition on critical boundary between extinction and stability:

$$T \equiv p q \lambda N^2 \sim 1$$

Importantly, we see that the critical mutation probability decreases with the population size so that for larger populations more ‘biologically relevant’ p_c would be expected. We return to this result in a later section and put this expression in terms of model parameters.

Appendix C: Qualitative agreement between stochastic and deterministic simulations

We simulated the deterministic dynamics with random coefficients so that extinction is excluded. We fix $L = 8$, $p = 1$, $a = s = 10^{-3}$ and vary σ . We find that the deterministic simulations qualitatively agree with the stochastic simulations. First we find apparent damped oscillations to a fixed point for large σ . As we increase σ , the population size of the predator and prey at the fixed point increases as a power of σ . We find that the population size grows approximately as a power law $X_{tot} \sim Y_{tot} \sim \sigma^\eta$, with $\eta \approx 1/2$. This power differs from the one we find in the stochastic simulations ($\eta = 1/4$). We expect that this is primarily due to demographic fluctuations, which induce a so-called ‘Griffiths phase’ in the stochastic simulations. The Griffiths phase is also the root cause of the apparent chaos in the stochastic simulations. We discuss this in a later section.

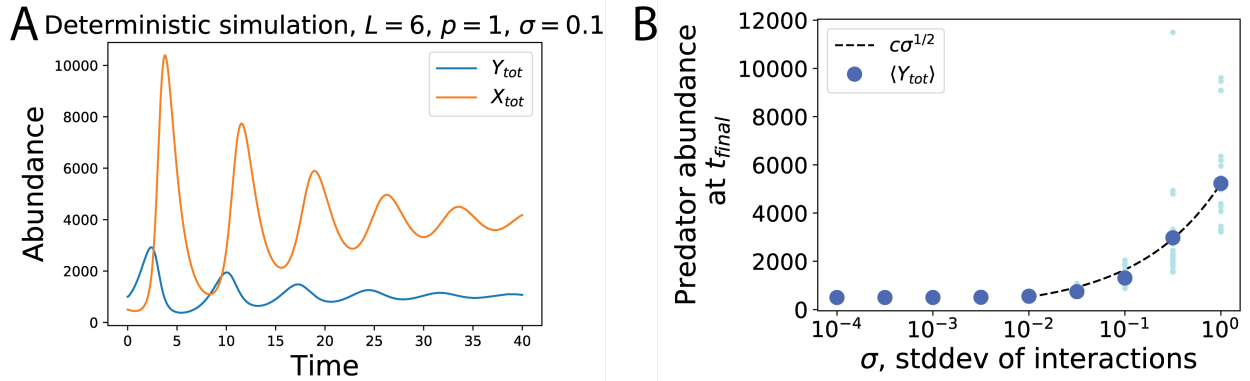


Figure 4.4: A.) Damped oscillations for a random deterministic simulation. Notice that the total population size is larger than that which would be expected from the neutral case, $Y_{tot} = 500$, $X_{tot} = 1000$. B.) for fixed $L = 8$ and $p = 1$, the final abundance increases with the variance of the interactions. The characteristic exponent is different from the stochastic case.

Appendix D: Dynamical mean field theory (DMFT) and analysis of the stable phase

DMFT equations for the total population sizes

Since there is agreement between stochastic and deterministic simulations, we work with the deterministic case to simplify our analysis and understand the mechanism of stabilization and the origin of power law growth in population size. We work in the case where $b_i = d_k = 1$. Since we are concerning ourselves with the mean-field limit, we have that there are no extinctions in this setting. Formally, we work in the thermodynamic limit so that $2^L \rightarrow \infty$, but we leave in these factors for clarity about scaling and the relative size of different parameters. We have that the mean field equations for the total abundances are given by:

$$\begin{aligned}\frac{dX_{tot}}{dt} &= \sum_i X_i \left(1 - (a + s)Y_{tot} - \sum_{k=1}^{2^L} (\delta a_{ik} + \delta s_{ik})Y_k + h_i(t) \right) \\ \frac{dY_{tot}}{dt} &= \sum_k Y_k \left(-1 + aX_{tot} + \sum_{i=1}^{2^L} \delta a_{ik}X_i + h_k(t) \right)\end{aligned}$$

where we have split the interactions into mean and fluctuating components:

$$\begin{aligned}a_{ik} &= a + \delta a_{ik} \\ s_{ik} &= s + \delta s_{ik}\end{aligned}$$

We have also added response fields h_i and h_k for convenience, which we will later set to 0. We can use the lens of dynamical mean-field theory to understand the properties of these quantities. The essence of dynamical mean field theory is that the sums over interactions:

$$\sum_k (\delta a_{ik} + \delta s_{ik})Y_k \quad \text{and} \quad \sum_i \delta a_{ik}X_i$$

can be averaged over and treated self-consistently in the many-strain limit. In the language of statistical physics, the interactions are decoupled and replaced with an effective noise and ‘self-interaction.’ The dictionary for the mapping (which has been derived in other many other contexts [8, 31, 29, 22]) is that:

$$\begin{aligned}\sum_i \delta a_{ik}X_i &\rightarrow \sigma_a 2^{L/2} \eta_i(t) - \sigma_a^2 2^L \int_0^t R_X(t, t') Y_k(t') dt' \\ \sum_k (\delta a_{ik} + \delta s_{ik})Y_k &\rightarrow (\sigma_a + \sigma_s) 2^{L/2} \eta_i(t) + \sigma_a^2 2^L \int_0^t R_Y(t, t') X_i(t') dt'\end{aligned}$$

where we use the natural scaling suggested in [29] and carefully keep track of 2^L prefactors. The first terms are zero-mean noises with correlation:

$$\begin{aligned}\overline{\eta_i(t)\eta_i(t')} &= 2^{-L} \sum_k Y_k(t)Y_k(t') \\ \overline{\eta_k(t)\eta_k(t')} &= 2^{-L} \sum_i X_i(t)X_i(t')\end{aligned}$$

which is demanded by self-consistency. This term determines the invasion properties of the different strains (i.e. it contributes to their fitness while rare).

The second term is the linear response when X_i gets large. Intuitively this accounts for the fact that when X_i large, it will perturb the Y_k (which will result in births, either of Y_k or its mutational neighbors), which will in turn feed back on X_i (thus altering its growth rate). Similarly Y_k at high abundance will deplete their preferred prey and start to decline. The intuition is that, without mutation, feedback effects result in effective time-dependent carrying capacities for the predator and prey populations. Although, as we will see in the next section, sufficiently strong mutation must change this intuition. The response functions are given (once again, due to self-consistency) by:

$$\begin{aligned}R_Y(t, t') &= \left. \frac{\delta Y_k(t)}{\delta h_k(t')} \right|_{h_k=0} = 2^{-L} \sum_k \left. \frac{\delta Y_k(t)}{\delta h_k(t')} \right|_{h_k=0} \\ R_X(t, t') &= \left. \frac{\delta X_i(t)}{\delta h_i(t')} \right|_{h_i=0} = 2^{-L} \sum_i \left. \frac{\delta X_i(t)}{\delta h_i(t')} \right|_{h_i=0}\end{aligned}$$

where these derivatives are in the functional sense.

Summing over all strains in the stable phase (where all strains might be seeded by mutation, so none go extinct indefinitely), we have the DMFT equations for the total population sizes:

$$\begin{aligned}\frac{dX_{tot}}{dt} &= \sum_i X_i \left(1 - (a + s)Y_{tot} - (\sigma_a + \sigma_s)2^{L/2}\eta_i(t) \right. \\ &\quad \left. - \sigma_a^2 2^L \int_0^t R_Y(t, t')X_i(t')dt' + h_i(t) \right) \\ \frac{dY_{tot}}{dt} &= \sum_k Y_k \left(-1 + aX_{tot} + \sigma_a 2^{L/2}\eta_k(t) \right. \\ &\quad \left. - \sigma_a^2 2^L \int_0^t R_X(t, t')Y_k(t')dt' + h_k(t) \right)\end{aligned}$$

We can rewrite these equations, setting the response fields to zero now that they are no longer needed:

$$\begin{aligned}
\frac{dX_{tot}}{dt} &= X_{tot} \left(1 - (a + s)Y_{tot} - (\sigma_a + \sigma_s)2^{L/2} \langle \eta(t) \rangle_{X(t)} \right. \\
&\quad \left. - \sigma_a^2 2^L \left\langle \int_0^t R_Y(t, t') X(t') dt' \right\rangle_{X(t)} \right) \\
\frac{dY_{tot}}{dt} &= Y_{tot} \left(-1 + aX_{tot} + \sigma_a 2^{L/2} \langle \eta(t) \rangle_{Y(t)} \right. \\
&\quad \left. - \sigma_a^2 2^L \left\langle \int_0^t R_X(t, t') Y(t') dt' \right\rangle_{Y(t)} \right)
\end{aligned}$$

The angle brackets denote the (grand canonical) average:

$$\langle \mathcal{O}(t) \rangle_{Y(t)} = \frac{1}{Y_{tot}(t)} \sum_k \mathcal{O}_k(t) Y_k(t)$$

which is (formally) different from the overline (microcanonical) average:

$$\overline{\mathcal{O}} = 2^{-L} \sum_k \mathcal{O}_k$$

We finally assume we are in a statistically steady state so that correlation and response are only a function of time differences:

$$\overline{\eta_i(t) \eta_i(t')} \equiv C_i(t - t'), \quad R_I(t, t') \equiv R_I(t - t')$$

We must make one more key observation in order to make progress, namely that in an (ergodic) steady state, the equivalence of ensembles holds and so the above averages are in fact equivalent in a statistical description of the dynamics (i.e. in the thermodynamic limit or when we average over realizations of disorder). It should also hold in the long time limit of a sufficiently large system ($2^L \gg 1$)¹. In the thermodynamic limit ($2^L \rightarrow \infty$), this is asymptotically exact. Furthermore, we expect that demographic stochasticity accelerates the approach to this limit, since its net effect is to homogenize the long time strain abundance distributions. Incidentally, the canonical ensemble (where Y_{tot} and X_{tot} are fixed to begin with, as in [29, 8, 25]) should provide an equivalent statistical description in this limit. Starting from the canonical perspective is a fruitful avenue to get insight into the individual strain dynamics. In this picture, we surmise that the population size will determine whether the short time dynamics are chaotic or not.

¹It should also be true that for sufficiently large system size, a particular realization of the disorder is similar to the average. The condition for this is that we should have $L \gg 1$ in the p model to sufficiently homogenize the mutation terms. In the μ model this is not an issue since mutation is homogeneous from the model definition. However, this is irrelevant in the present work since we are averaging over disorder, although understanding the variability of particular realizations would be an interesting topic for future work.

An equivalent formulation of this property is that in the large 2^L limit, when averaged over realizations over disorder, the long time averages of X_i and Y_k are homogeneous in i and k . This follows from the fact that the strain labels are arbitrary and exchangeable when we draw random interaction rates. However, in the stochastic simulations, this can only hold in the stable phase when the parameter combination $T \equiv pq\lambda N^2$ is sufficiently large and are technically only truly realized over long times (since temporary extinction is possible and in fact frequent). This is an important, but subtle, point since although the total population sizes follow an almost deterministic trajectory, the individual strains can still fluctuate wildly with constant extinction and re-colonization by mutation. In this sense T acts like a temperature, facilitating diffusion across the interaction landscape, which is why we use the suggestive naming (i.e. to stretch our analogies a bit further, the stable phase could be called an eco-evolutionary gas/paramagnet). In a later section we will more concretely define T in terms of the model parameters.

With this in hand, averaging over disorder and replacing angle bracket averages with overline (microcanonical) averages, we have the following closed equations for the total population sizes:

$$\begin{aligned}\frac{dX_{tot}}{dt} &= X_{tot} \left(1 - (a + s)Y_{tot} - \sigma_a^2 \int_0^t R_Y(t - t') X_{tot}(t') dt' \right) \\ \frac{dY_{tot}}{dt} &= Y_{tot} \left(-1 + aX_{tot} - \sigma_a^2 \int_0^t R_X(t - t') Y_{tot}(t') dt' \right)\end{aligned}$$

We have used that the η_i and η_k have zero mean and are anyway irrelevant to the total population size in the thermodynamic limit (they end up with prefactors of $2^{-L/2}$). To belabor the point: although we have interchanged the types of averages, we should not confuse this statement with the statement that the strain abundances $X_i(t)$ and $Y_k(t)$ are homogeneous across i and k at a given time or for a particular realization of the disorder (which they are certainly not), but rather that they are homogeneous when averaged over disorder and averaged over long times. On average, the population as a whole does not ‘see’ the strain variations.

Fixed point of the DMFT equations

From this we obtain the fixed point equations:

$$\begin{aligned}0 &= 1 - (a + s)Y_{tot} - \sigma_a^2 \chi_Y X_{tot} \\ 0 &= -1 + aX_{tot} - \sigma_a^2 \chi_X Y_{tot}\end{aligned}$$

where we have defined the susceptibilities:

$$\chi_X \equiv \int_0^\infty R_X(\tau) d\tau, \quad \chi_Y \equiv \int_0^\infty R_Y(\tau) d\tau$$

We can solve these equations to get the fixed point:

$$\begin{aligned} Y_{tot} &= \frac{a - \sigma_a^2 \chi_Y}{a(a + s) + \sigma_a^4 \chi_X \chi_Y} \\ X_{tot} &= \frac{(a + s) + \sigma_a^2 \chi_X}{a(a + s) + \sigma_a^4 \chi_X \chi_Y} \end{aligned}$$

From these expressions we can observe a few interesting things. First, in the $\sigma_a = 0$ case, we recover the standard fixed point for non-disordered interactions, as expected. However, at first glance, these expressions seem to indicate that the predator population size decreases as σ_a increases.² But this is in direct contradiction to our observations from stochastic and deterministic simulations!

The key to understanding this discrepancy is that in the stable regime the susceptibilities must be such that χ_Y is negative and χ_X is positive when defined as we have above.³ This can be understood intuitively to be a result of mutation, which changes the nature of the feedback. Mutation has not entered into the total population size dynamics up to this point although it crucially enters into the individual strain dynamics (which are needed to precisely determine the susceptibility in the first place). We can get a handle on this intuition in the strong mutation regime ($p \sim 1$). In this regime, some large fraction of predator births (which only occur due to predation interactions) result in random mutation. So when predator Y_k 's favorite prey X_i is at high abundance, much of Y_k 's consumption results in growth of its mutational neighbors. These mutational neighbors in turn are unlikely to prefer the same prey (if the interactions are sufficiently uncorrelated). So at the end of the feedback loop X_i gets a net boost in growth rate since it's preferred predator has been net depleted. Prey mutations do not feed back on the response χ_X in this way since they do not depend on the interactions.⁴

²This is actually a version of the classic May result when there is no mutation. When σ_a gets too large, the predators go extinct. The prey do not go extinct in our model because in the absence of predators they grow exponentially.

³Note that the change of sign of carrying capacity due (separately) to quenched disorder, diffusion and demographic noise has been reported in the context of a generic model of directed percolation [20]. Since all three are at play in the full stochastic model and strong diffusion (mutation) occurs even at the deterministic level, we expect that a similar effect is at play here. Care would need to be taken in checking this and a full renormalization group analysis of the strain level equations would need to be carried out to verify this.

⁴N.B. In similar models with *mutation rates*, μ , which exhibit regular diffusion in genotype space (versus *mutation probabilities*, p , which exhibit nonlinear diffusion, see below strain level DMFT equations), there should also be a stable ergodic phase but only if μ is sufficiently large. More specifically the important ingredient is that the predator-specific μ is sufficiently large (i.e. $L\mu^{-1}$ should be comparable to the timescale of the interactions) that it makes χ_Y negative. In phage-bacteria systems this may in fact be the case since phage and viruses generally have significantly higher mutation rates than microbes. In the end we argue that the p -model is more true to the underlying biology than the μ -model. However, the μ -model might be more readily compared to existing data, since what is typically measured is the rate rather than the probability. A more direct and detailed comparison of the two models is an interesting avenue for future work but is beyond the scope of the present study.

Once we see this, we can also begin to understand why the population size increases with σ_a from self-consistency requirements. For the fixed point expressions to be valid, they are required to be positive. We see for fixed large but finite 2^L , as we increase σ_a , the susceptibilities χ_Y and χ_X must become sufficiently small to compensate so that the following constraint must be satisfied:

$$|\chi_X \chi_Y| < \frac{a(a+s)}{\sigma_a^4}$$

However, the χ_i should not become arbitrarily small since at least some of the interactions are quite large in the intermediate regime $2^{2L} \gg \sigma_a^4 \gg a(a+s)$ where linear response is still a valid approximation. Therefore we make the scaling ansatz for the asymptotic expansion of the susceptibilities:

$$|\chi_X \chi_Y| \sim \frac{a(a+s)}{\sigma_a^4} - c \left(\frac{p}{\sigma_a} \right)^\eta$$

where c is a constant and the quantity in brackets is a dimensionless expansion parameter. This scaling form can be thought of as a bound on how large the magnitude of the susceptibilities can be and remain self-consistent.

So as we increase σ_a while keeping 2^L and a fixed, we expect the population size to diverge:

$$X_{tot} \sim Y_{tot} \sim \sigma_a^\eta$$

This is borne out in our stochastic simulations with a value for the exponent $\eta \approx 1/4$ and in the above deterministic simulations $\eta \approx 1/2$. Furthermore, we expect the stationary state population size to decline with p :

$$X_{tot} \sim Y_{tot} \sim p^{-\eta}$$

which also agrees with stochastic simulations. For deterministic simulations at long times, any $p > 0$ will result in a fixed point.

Moreover we see that this expansion agrees with simulations even for relatively weak disorder ($a/\sigma^2 < 1$). We make the argument that this power law is generic (i.e. non-universal) in a later section. An important ingredient for better understanding the precise behavior of the susceptibilities (and their scaling) is a detailed description of the strain-level dynamics which requires a more careful and technical analysis than that presented here. This analysis will require the full machinery of reaction-diffusion processes [4].

Finally, we note that while the mean of the ‘lossy’ interactions, s , affects the location of the fixed point, the variance σ_s is irrelevant. These interactions do not form a closed loop between predator and prey, so there is no feedback, which is where the variance comes in at the population level. However, the variability of these interactions will contribute to the stochastic dynamics and may shift the location of the transition to the stable phase. In addition the s_{ik} (and σ_s) will certainly be relevant in the short time strain level dynamics.

Stability of the fixed point

We can perform a linear stability analysis of the equations:

$$\begin{aligned}\frac{dX_{tot}}{dt} &= X_{tot} \left(1 - (a + s)Y_{tot} - \sigma_a^2 \int_0^t R_Y(t - t') X_{tot}(t') dt' \right) \\ \frac{dY_{tot}}{dt} &= Y_{tot} \left(-1 + aX_{tot} - \sigma_a^2 \int_0^t R_X(t - t') Y_{tot}(t') dt' \right)\end{aligned}$$

We focus on $\omega = 0$ mode which is maximally unstable, and which simplifies our analysis since this allows us to set the responses to their $\tau = \infty$ values, χ_i [25]. A simple calculation leads us to the stability eigenvalues:

$$\lambda_{\pm} = \frac{-1 \pm \sqrt{1 - 4 \frac{[(a+s) + \sigma_a^2 \chi_X][a + \sigma_a^2 |\chi_Y|]}{a(a+s) - \sigma_a^4 \chi_X |\chi_Y|}}}{2}$$

where we have the discriminant:

$$\frac{[(a+s) + \sigma_a^2 \chi_X][a + \sigma_a^2 |\chi_Y|]}{a(a+s) - \sigma_a^4 \chi_X |\chi_Y|} > 0$$

since, importantly, the denominator must be greater than 0 for the sake of self-consistency. Therefore the total population size of a collection of predator-prey pairs is a stable focus although the dynamics of a single predator-prey pair is a neutrally stable limit cycle. In the limit of strong disorder, the complex component of the stability eigenvalue grows so that the frequency of the oscillatory component increases. Interestingly, even for weak disorder we have that the mean field population size fixed point is a focus with eigenvalues:

$$\lambda_{\pm} \approx \frac{-1 \pm 3i}{2}$$

whereas the non-zero fixed point of the equivalent neutral predator prey model has pure imaginary stability eigenvalues:

$$\lambda_{\pm}^{neutral} = \pm i$$

so that disorder is a singular perturbation at the deterministic level. Despite the deterministic stability, there are potentially large transient initial fluctuations and also metastable states which are observed in simulations.

The stability results are consistent with simulations of the deterministic dynamics, which show that the total population size is a stable node for a broad range of σ_a . However, in the stochastic simulations, we see that oscillatory fluctuations persist due to a form of stochastic resonance in which the stochasticity (birth-death noise) has frequency components that are comparable to the frequency of the damped oscillations.

Comment on the τ -dependence of the response

In order for our assumptions about the response and its scaling for to hold in the previous sections, it must be sufficiently sharply peaked and importantly, integrable. A function with an exponential cutoff will give us the right properties:

$$R_I(\tau) = f_I(\tau, p, \xi_I) e^{-\tau/\xi_I}$$

where ξ_I are the correlation times:

$$\begin{aligned}\xi_X &= \frac{a+s}{\sigma_a^2} \\ \xi_Y &= \frac{a}{\sigma_a^2}\end{aligned}$$

These are natural timescales that we can construct from the parameters in our model. We also have the functions f_I which should be such that:

$$\begin{aligned}f_X(0, 1, 0) &= 1 \\ f_Y(0, 1, 0) &= -1\end{aligned}$$

and:

$$\begin{aligned}f_X(0, 0, 0) &= 1 \\ f_Y(0, 0, 0) &= 1\end{aligned}$$

Qualitatively, we expect that f_Y is a step-like function in p , where ξ sets the location of the step.

As the variance increases, the correlation time decreases so that the response is more and more sharply peaked (and in the limit it becomes a delta function). In the limit, we get the right leading order scaling for the susceptibilities. This also justifies our stability analysis in the $\sigma_a^2 \gg a$ limit. Importantly, we can in principle check this by measuring the correlation function of Y_{tot} and X_{tot} which will decay exponentially with these predicted timescales, although this is made difficult by the observed resonance in the stochastic model. We will come back to the time-dependence of the response when we discuss the stochastic simulations and the emergence of generic power law behavior in Griffiths phases.

Appendix E: DMFT for strain level dynamics and estimating the phase boundary redux

We have the strain level DMFT equations:

$$\begin{aligned}
\frac{dX_i}{dt} &= X_i \left(1 - (a + s)Y_{tot} - (\sigma_a + \sigma_s)\eta_i(t) - \sigma_a^2 \int_0^t R_Y(t, t') X_i(t') dt' \right) \\
&\quad + \frac{p}{L} \sum_{j \in \mathcal{S}_i} (X_j - X_i) \\
\frac{dY_k}{dt} &= Y_k \left(-1 + aX_{tot} + \sigma_a \eta_k(t) - \sigma_a^2 \int_0^t R_X(t, t') Y_k(t') dt' \right) \\
&\quad + \frac{p}{L} a X_{tot} \sum_{l \in \mathcal{S}_k} [Y_l(t) - Y_k(t)] \\
&\quad + \frac{p}{L} \sum_{l \in \mathcal{S}_k} \left[\sigma_a 2^{L/2} (\eta_l(t) Y_l(t) - \eta_k(t) Y_k(t)) \right. \\
&\quad \left. - \sigma_a^2 2^L \int_0^t R_X(t, t') (Y_l(t') Y_l(t) - Y_k(t') Y_k(t)) dt' \right]
\end{aligned}$$

where the additional terms are the disorder averaged mutation operators. The first of these terms is much larger than the others since it is proportional to the total prey population size. We have included the L dependence although technically in DMFT L has already been sent to infinity (and the sums have some limiting form as infinite dimensional diffusion operators). Unfortunately, we cannot strictly use the equivalence of ensembles as we did for the total population size if we are interested in the strain dynamics since the diffusion sums are only over L terms (2^L is technically our large parameter). In addition, the complicated form of these reaction-diffusion equations (with nonlinear diffusion with memory for the predators) resists obvious solution.

However, we can use these dynamics to more precisely estimate the phase boundary between stable eco-evolutionary chaos and extinction. Note that due to the inherent stochasticity of the population dynamics (i.e. strains fluctuate to extinction and are repopulated), there is more smoothing of the distribution of abundances in the fully stochastic mean-field equations than in the deterministic one, so while the dynamics of a particular realization are by no means homogeneous, we can use homogeneous equations to estimate timescales.

There are two timescales that are important to compare, as in our heuristic description. First is the extinction timescale, t_{ext} and the second is the mutation or diffusion timescale, t_D . We expect that the transition between the extinct state and eco-evolutionary chaos is first order because the extinction timescale differs as we approach the transition from both directions in, for instance, σ_a [12]. This is due to the fact that approaching the transition point from the left, depending on the initial conditions, we move along a neutrally stable limit cycle or divergent chaos (both of which rapidly wander or decay to extinction), whereas

the right side of the transition is asymptotically stable at the population level and takes exponentially more time to go extinct due to stochastic effects.⁵ This hypothesis will be difficult to check directly because in the model, noise, diffusion and disorder are all strong (and so in finite systems, the transition is ‘smeared out’). We discuss aspects of this smearing and the Griffiths phase phenomenon in a later section.

If we approach from the right (from the ergodic phase), the population size sets the extinction timescale. We observe that the predator population is more likely to go extinct first (since it is the smaller population size) so that this is the relevant population to focus on. It can be established that t_{ext} scales as Y_{tot} for a stable fixed point subject to demographic fluctuations [26]. We have that the diffusion timescale is:

$$t_D = \frac{l_D^2}{\langle D_Y \rangle}$$

The averaged diffusion constant can be approximated by $\langle D_Y \rangle \approx \frac{paX_{tot}}{L}$ in the thermodynamic limit. The distance l_D can be understood as how far in genotype space a focal type Y_k needs to mutate to reach a beneficial genotype. In the thermodynamic limit this distance is just $l_D = 1$. So the condition for these timescales being comparable is:

$$\frac{paX_{tot}}{L} Y_{tot} \sim 1$$

This gives a result that is consistent with our heuristic derivation above. Using the scaling forms of the population sizes, we have that $X_{tot} = c_1 \sigma_a^{1/4}$ and $Y_{tot} = c_2 \sigma_a^{1/4}$ so that the critical sigma for the simulations in Fig. 3 is:

$$\sigma_c = \left(\frac{1}{pac_1 c_2} \right)^2 \approx 2.5 \times 10^{-4},$$

which is consistent. We also note that the boundary slightly undershoots the apparent transition point in the simulations, which also makes sense since we neglect the (negative) response term in our estimate of the diffusion constant.

For the simulations with fixed $\sigma = 0.1$ and varying p , we notice anomalously long lived metastable states where small sets of predator and prey exist at finite abundance for long times. This is due to initial conditions and ‘lucky’ sets of predator and prey strains that can

⁵N.B. We can make an alternate argument for why the transition should be 1st order. Grassberger and Janssen conjectured that any active-inactive state transition should fall under the directed percolation (DP) universality class [14, 16]. Our transition is an active-inactive state transition so it should fall under DP. In p , we expect that the transition is second order. In σ , we expect that the transition is first order. Above the critical dimension, there can exist a 1st order transition [14] (when the carrying capacity is the control parameter, as is effectively the case in our model). Our model ($d = 2^L$) is well above $d_c = 4$ so this is still consistent with a first order transition. It is interesting to note that the single locus dynamics ($L = 1$) are below the critical dimension and the two-locus dynamics ($L = 2$) are at the critical dimension. So the transition might only be 2nd order in these cases.

evade extinction due to the random nature of the interactions. Understanding the nature (specifically the connectivity and scale of the interactions) of these metastable states is an interesting avenue for future work and might be extremely relevant in real predator prey systems. Nevertheless, we can still estimate the location of the transition to get:

$$p_c \approx 4 \times 10^{-4}$$

which is also consistent.

These estimated phase boundaries are plotted in Figure 3 in the main text.

The effect of correlations on the transition

Correlations in interactions space enter into the diffusion length described in the previous section. In the limit of no correlations (which we primarily discuss in this work), this length is 1 in the thermodynamic limit (there is always a beneficial mutation one step away). In fact, for any finite correlation length, in the thermodynamic limit $l_D = 1$. In the case of perfectly correlated interactions, this length is infinite – there are no beneficial mutations to be found no matter how far one mutates. This corresponds to the neutral case in our model, in which there is rapid extinction.

However, in real systems (where we may not be sufficiently far into the thermodynamic limit or where correlations might be long-ranged), this length is probably some finite value in between – i.e. there are certainly correlations in interaction space, but there are still large jumps that are accessible by few mutations [19, 5]. This can be ameliorated specifically in viral systems, because even if there are correlations in interaction space, they can compensate by increasing their burst size, or by having a sufficiently wide burst size distribution. This will contribute to the diffusion constant and bring the diffusion time down. An intuitive way to think about this case is that a single birth event can sample interaction space deeply enough to overcome barriers set by correlations.

Appendix : Simulations on correlated landscapes

We simulated a simple version of the model with phylogenetic correlations by explicitly modeling the binding energies between predator and prey. We considered two cases.

The first case we consider is one in which the interaction energy is additive:

$$E_{ik}^a = \sum_{\alpha=1}^L \epsilon_{\alpha}(i_{\alpha}, k_{\alpha})$$

We draw the $\epsilon_{\alpha}(i_{\alpha}, k_{\alpha})$ from normal distributions so that the mean and variance of the rates, $a_{ik} = e^{E_{ik}^a}$, are fixed to be a and σ^2 , respectively (on average). The s_{ik} interactions were defined similarly. As σ grows, closely related neighbors become more and more uncorrelated (a single step mutation can be large) so that we see the same transition and same qualitative behavior as the uncorrelated model.

We note that in this additive random model, fewer random variables are drawn in defining the interaction landscape ($4L$ vs 2^{2L}) so that there is a higher variance among realizations. This is a general property of defining finite correlated landscapes. Furthermore, if we look at the strain level dynamics, we see that they are quite different. For a finite sample, strains are no longer “almost exchangeable,” and there are clear “winners” which persist at high abundance for long times (i.e. it makes sense to consider fitness in this model). However, this does not change the population level dynamics since enough high fitness types (typically those predator-prey pairs with small interactions and high population size) exist at a time so that there is still a stabilizing effect due to intraspecific competition.

The second case we consider is an energy with ‘pairwise’ correlations:

$$E_{ik}^a = \sum_{\alpha, \beta=1}^L \epsilon_{\alpha, \beta}(i_{\alpha}, i_{\beta}, k_{\alpha}, k_{\beta})$$

We draw the $\epsilon_{\alpha, \beta}(i_{\alpha}, i_{\beta}, k_{\alpha}, k_{\beta})$ from normal distributions so that the mean and variance of the rates, $a_{ik} = e^{E_{ik}^a}$, are fixed to be (on average) a and σ^2 , respectively. The s_{ik} interactions were defined similarly. When comparing this to the additive model, we see that there are still some “winners,” which stay at high abundance for long times. However, the gap between them and the low frequency types has closed a fair amount.

We note that in the additive and pairwise random models, fewer random variables are drawn in defining the interaction landscape ($4L$ vs $16L^2$ vs 2^{2L}) so that there is a higher variance among realizations of the random interactions. This is a general property of finite correlated landscapes. Furthermore, if we look at the strain level dynamics, we see that they are quite different. For a finite sample, strains are no longer “almost exchangeable,” and there are clear “winners” which persist at high abundance for long times (i.e. it could make sense to define a scalar fitness in this model). However, this does not change the population level dynamics since enough high fitness types (typically those predator-prey pairs with small interactions and high population size) exist at a time so that there is still a stabilizing effect due to intraspecific competition.

We could increase the number of random draws by using more complicated correlation structure. But for our purposes, since we see the same phenomenology (e.g. stabilization) in the additive and pairwise models, we expect it to be present for less-correlated (i.e. higher order) models. These less-correlated models might include, for instance, NK models with $N = L$ and $1 < K < L - 1$. Our additive, pairwise and uncorrelated models are the $K = 0$, $K = 1$ and $K = L - 1$ cases, respectively. Other models that might be interesting to consider include the “Rought Mt Fuji” model [23] and the Gaussian model of Agarwala and Fisher [2]. However, for the latter, a new simulation scheme would need to be implemented since ours depends on declaring the interaction landscape ahead of time.

We note that there is interesting transient behavior encoded in the model, especially when initialized with a single strain. We see that populations can take a long time to reach the steady state and are much more likely to go extinct by chance when initialized this way. This is similar to what was seen with the fixed σ , p -varying simulation results in the main text. Furthermore, it is interesting that the final (stable) population sizes reached in the correlated models are typically much larger than those reached in the uncorrelated model. This is likely due to the fact that individual types can have much smaller interaction rates in the correlated model than those drawn from an uncorrelated distribution. We leave exploring the details of the additive and other correlated models (especially their asymptotic behavior on infinitely large genotype spaces) as an interesting avenue for future work.




Figure 4.5: Population dynamics on additive landscape for $a = 10^{-3}$, $\sigma = 1$, $p = 0.1$. We see similar behavior as the uncorrelated case, where at long times the population seems to go to a fixed point much higher than the fixed point associated with the average interaction strength. However, since it is difficult to control the mean and variance of a finite sample of such a highly correlated landscape, there is a relatively high variance of the populations' final abundances. Populations were initialized with some fraction of strains at finite abundance with total abundance near their deterministic fixed point values. Occasionally, there is extinction due to initial transients. However, this only happened once out of 20 simulations.




Figure 4.6: Strain dynamics on additive landscape for $a = 10^{-3}$, $\sigma = 1$, $p = 0.1$. Note that in this case several strains go extinct but several persist for long simulation times with relatively small abundance fluctuations.

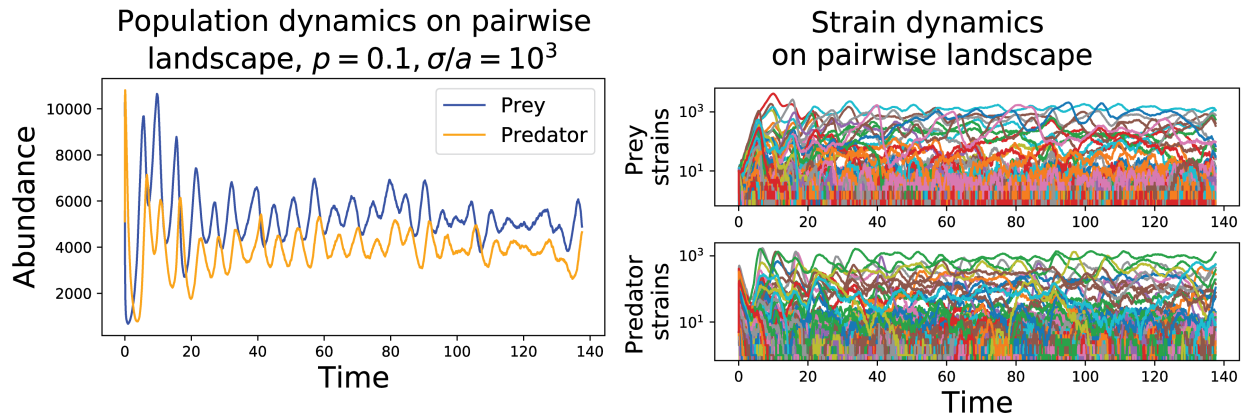


Figure 4.7: Population and strain dynamics on pairwise landscape for $a = 10^{-3}$, $\sigma = 1$, $p = 0.1$. Note that in this case there is stabilization but much fewer strains are extinct at the end of the simulation. This can be understood by the fact that interactions are more ‘high-dimensional’ in the pairwise case.

Appendix F: Apparent chaos as a Griffiths phase

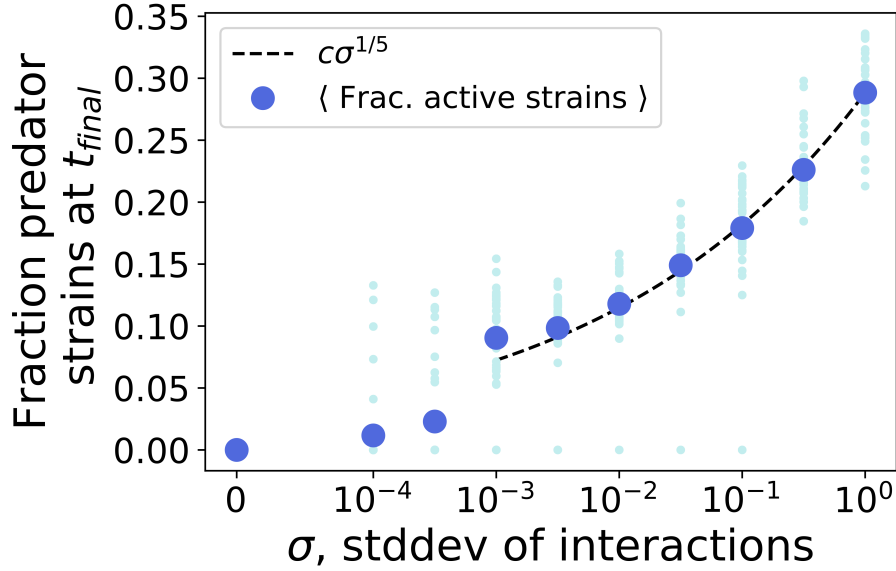


Figure 4.8: Fraction of active ($Y_k > 0$) predator strains in stochastic simulations at the final timepoint. As is the case in Griffiths phases in generic nonequilibrium and quantum systems, only a fraction of sites are active even though in the deterministic dynamics all sites are active. The fitted power law exponent is different than that obtained for the total population size.

The above argument about the transition is analogous to the ferromagnetic (unstable phase) to paramagnetic (stable phase) transition in models of magnetism. A well-known effect in magnetism is that with disorder, there is a continuous region above the transition point known as a Griffiths phase. In magnetism, a Griffiths phase is one in which a system does not fully fall into a paramagnetic or ferromagnetic phase. Rather, it is dominated by macroscopic ‘rare regions’ that relax slowly. This effect has been demonstrated to extend to classical nonequilibrium and quantum models.

We can understand the Griffiths phase phenomenon intuitively in the case of our model. For fixed population sizes X_{tot} and Y_{tot} and for fixed number of genotypes 2^L , when $2^L \gtrsim X_{tot} \sim Y_{tot}$ there are more possible genotypes than can exist at once at finite abundance, given stochastic effects and variance in a given finite realization of the disordered interactions. This means only a subset will exist at finite abundance at a given time, although in the long time limit the average abundances should approach the ones obtained from the deterministic model. In a real predator-prey system, we expect genotype space to be fairly large (much

larger than any reasonable population size) so that only a subset of all possible genotypes may exist at a given time.

The presence of a Griffiths phase does not change our arguments for the location of the transition. However, there should be a second transition where the Griffiths phase gives way to the deterministic expectation. This can happen two ways – by increasing the population size (i.e. having $a^{-1} \gg 2^L$ so all genotypes have average abundance greater than zero) or by increasing σ . In practice, population sizes are much smaller than genotype space. The only reason we choose such relatively large a is because of computational constraints. We simulated a population with $a = 10^{-5}$, so a population of $\sim 10^5 > 2^L \approx 10^3$ individuals. We see that the strains stabilize and look less like the Griffiths type phase. Furthermore, no strain is extinct for very long. With better optimized simulations, and more computing power, simulating populations up to 10^7 should be feasible using our method.




Figure 4.9: Population and strain dynamics on pairwise landscape for $a = 10^{-5}$, $\sigma = 10^{-4}$, $p = 0.1$. Note the remarkable correspondence between the population dynamics of the *different realizations* of the disorder.

The other way we might escape a Griffiths phase is by increasing σ . We note that as σ gets larger (for fixed L), the validity of linear response becomes more suspect since perturbations are no longer necessarily small. Furthermore, once σ gets sufficiently large, some rate parameters can become so large to be physically unrealistic models of predator and prey interactions. So for all intents and purposes the Griffiths phase may well extend over a sufficiently large region of the phase diagram that the deterministic dynamics are not

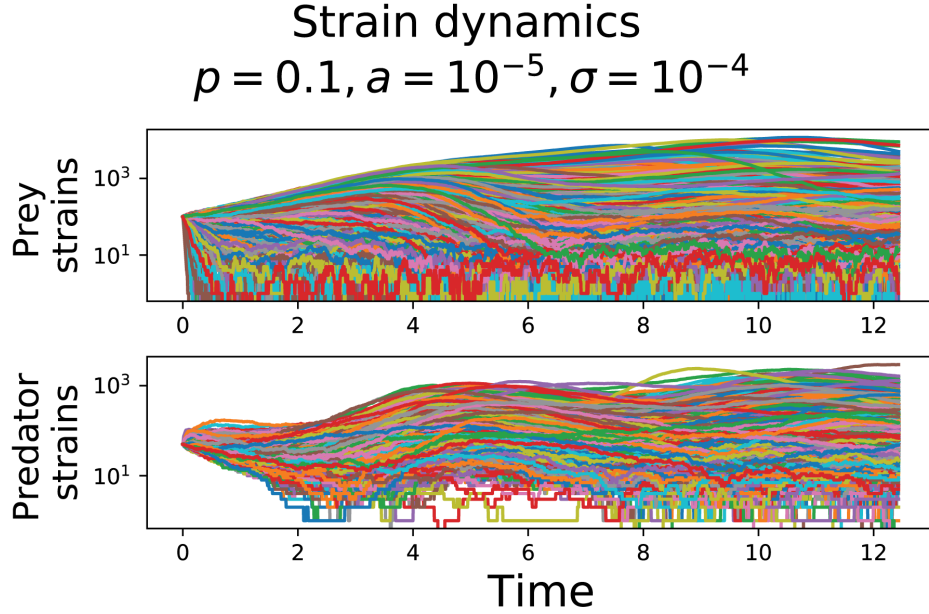


Figure 4.10: Population and strain dynamics on pairwise landscape for $a = 10^{-5}$, $\sigma = 10^{-4}$, $p = 0.1$. Note that while some strains still float near extinction, most of them are pulled away from zero. However in real systems, since $1/a \ll 2^L$, this is unlikely to be a realistic scenario.

representative of any real (i.e. stochastic) system over short times. However, even if this is the case, the deterministic dynamics will still be representative of long time averages.

Finally we comment that the Griffiths phase phenomenon is the likely cause of the discrepancy between the exponent in the stochastic and deterministic dynamics. Generic power laws are common in Griffiths phases due to exponentially large active domains lasting for exponentially distributed times. The convolution of these two factors can lead to a generic exponent. We expect that in the stochastic model, strains are present for exponentially distributed times so that the response must be averaged over these in addition to being averaged over disorder, leading to a different power than in the deterministic model.

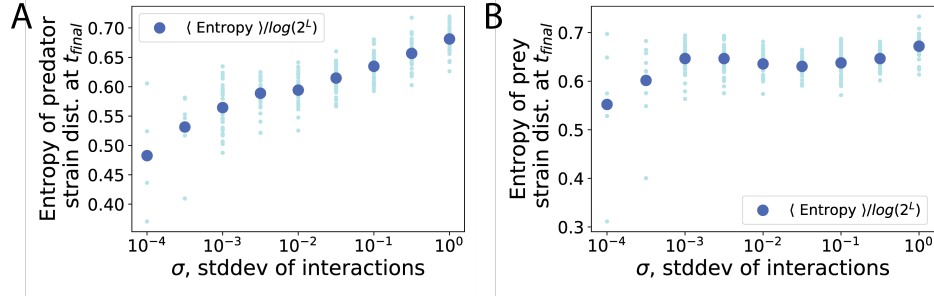


Figure 4.11: Entropy of the (A) predator and (B) prey populations divided by $\log 2^L$ as a function of σ . This approaches unity as σ increases, as the ‘rare regions’ start to span the system and the stochastic dynamics approach the deterministic dynamics. The prey population distribution approaches 1 more slowly since diffusion in prey genotype space is much weaker.

Appendix G: Steady population dynamics belie complex strain dynamics

Previous studies have shown that related eco-evolutionary (or epi-evolutionary) models can be mapped onto so-called traveling wave models of population genetics in certain parameter regimes [37]. Using our simulations, we can also access the strain dynamics. So it is natural then to ask whether such a mapping exists for the predator-prey model under consideration.

An essential ingredient of traveling wave models of evolution is that there be a large influx of mutations of fixed effect size so that multiple mutations compete with each other or even accumulate within lineages. Qualitatively speaking (in the case of beneficial mutations – these models have also been analyzed in the context of Muller’s ratchet/deleterious mutations), a beneficial mutant grows relative to the rest of the population as low fitness clones are purged and new beneficial mutants are seeded. These processes balance to generate an approximately Gaussian wave profile which travels at fixed speed in fitness space [6, 13].

In order to compare our model with traveling wave models, we define the prey and predator fitness:

$$g_i^{prey}(t) \equiv b_i - \sum_k (a_{ik} + s_{ik}) Y_k$$

$$g_k^{pred}(t) \equiv -d_i - \sum_i a_{ik} X_i$$

so that the prey fitness is bounded from above:

$$g_i^{prey}(t) \leq b_i$$

and the predator fitness is bounded from below:

$$g_k^{pred}(t) \geq -d_i$$

In the absence of predators, the prey population grows exponentially and in the absence of prey, the predator population decays exponentially. We also neglect the mutational term, which we are assuming to be relatively small (consistent with the case in Fig 4.12).

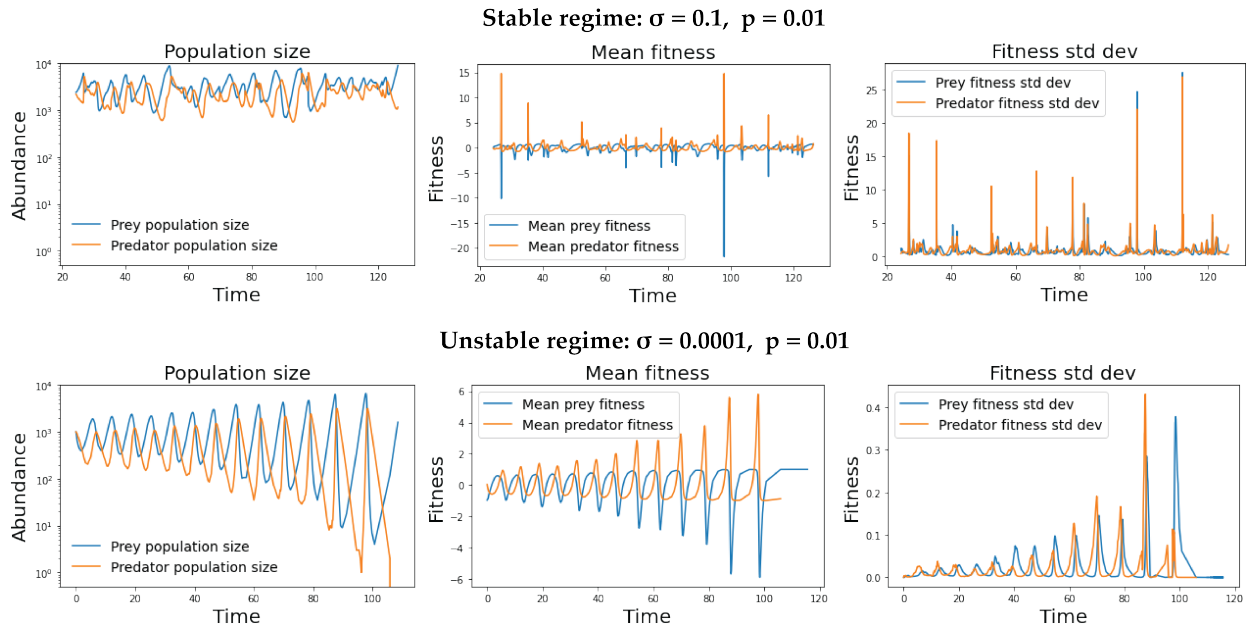


Figure 4.12: Abundance, mean fitness, and standard deviation of the fitness distribution in the stable regime and in the unstable regime. In the unstable regime, the dynamics of the fitness distribution is smooth, whereas in the stable regime it is intermittent. This is markedly different from traveling wave models of adaptation, where a stable fitness distribution forms.

We plotted the mean and standard deviation of the fitness distribution over time (Fig. 4.12). In the unstable weak selection regime, the mean and standard deviation smoothly oscillate, while in the strong selection regime the mean and standard deviation show spiky, almost intermittent behavior, with spikes corresponding to the invasion of highly fit types. However these spikes are sufficiently short that they do not distort the total population size by much.

From this, we can readily see that the population fitness distribution does not have a limiting shape which stands in stark contrast to the traveling wave picture. In our model, we see that the most fit genotype will typically have very high fitness, potentially orders of magnitude larger than the mean fitness. This is due to the fact that the fitness of a prey (predator) type depends on the predator (prey) population linearly. If a prey (predator)

has an anomalously small (large) interaction with a common predator (prey) type, it will have a relatively large fitness. However, this advantage is only transient as this highly fit type quickly gets pulled into the bulk of the fitness distribution and typically goes extinct as the environment changes. Because of this ‘leapfrog’ character, it is difficult to project these high-dimensional dynamics down onto the low-dimensional space of traveling wave models as Yan, et al did for a related model [37].

Instead, we can describe these dynamics in the language of eco-evolutionary feedback. We have established that without sufficient phenotypic variation, either the predator population or both populations will go extinct. However, when there is sufficient variation accessible by mutation, a prey which is anomalously free from predation will be born and will rise exponentially at its growth rate b_i . As it rises, its strongest predator will also start to rise at the point where $a_{ik}X_i \sim d_k$. The two focal types will then peak and start to decline, according to the coupled equations:

$$\begin{aligned}\dot{X}_i &\approx X_i(b_i - (a_{ik} + s_{ik})Y_k) \\ \dot{Y}_k &\approx Y_k(-d_k + a_{ik}X_i)\end{aligned}$$

Eventually X_i reaches low enough abundance (and likely goes extinct because there are sufficiently many weakly interacting predators) that Y_k declines at rate d_k until it is also extinct. This whole process happens on a timescale such that as the populations are declining, a new highly fit prey type is seeded. This results in the observed ‘breathing’ fitness distribution that spreads across genotype space in potentially complex ways. A further investigation of the population genetic characteristics of our model is an interesting avenue for future study.

Appendix H: Relation between the p -model and the μ -model

In the main text, we describe a model in which mutations are coupled to births so that they naturally occur on the same timescale as births. However, traditional population genetic analysis has assumed that mutation occurs sufficiently rarely that it occurs *independently of the underlying birth-death process*. The key operational difference is that this results in different classes of diffusion in mutational space.

The constant rate mutational operator is given by:

$$\frac{\mu}{L} \sum_{j \in \mathcal{S}_j} (X_j - X_i)$$

whereas the mutation probability operator is significantly more complicated and depends on the form of the interactions. For an interaction with a rate given by $F_i(\vec{X})X_i$ which results in a birth (and thus a potential mutation), the appropriate diffusion operator is given by:

$$\frac{p}{L} \sum_{j \in \mathcal{S}_j} (F_j(\vec{X})X_j - F_i(\vec{X})X_i)$$

In a situation where there is sufficient homogeneity, we can make the identification:

$$\mu = \frac{p\langle F_i \rangle}{L}$$

In terms of our model, this is simply:

$$\mu = \frac{paX_{tot}}{L} Y_{tot}$$

which we report in the main text.

Appendix I: Relation to spatio-temporal chaos in models with immigration

Mutation-selection dynamics can be thought of as a type of immigration-selection process, though the two processes differ in some key respects. The mutational picture can be mapped onto a high dimensional spatial landscape, where genotypes correspond to islands and mutation acts as migration. However, interactions happen across islands, instead of being localized to a single site. Furthermore, mutation is inherently different from migration in the sense that mutation can be coupled to the interactions (as we have described in the previous appendix). In situations where mutation is frequent and the population size fluctuates, mutation cannot be simply be decoupled from the population dynamics. When mutation is sufficiently rare (that it can be approximated as regular diffusion), the mapping to spatial models is consistent. In fact, we expect a correspondence between the Griffiths phase behavior in our model and the stable spatiotemporal chaos phase in [29, 30]. However, for many large natural populations (microbes, etc) this mapping would be consistent since the product $\mu N \gg 1$. This is the clonal interference regime (as mentioned in the main text) and would lead to the stable phase described here.

4.9 Bibliography

- [1] G. Ackland and I. Gallagher. Stabilization of large generalized lotka-volterra foodwebs by evolutionary feedback. *Physical review letters*, 93(15):158701, 2004.
- [2] A. Agarwala and D. S. Fisher. Adaptive walks on high-dimensional fitness landscapes and seascapes with distance-dependent statistics. *Theoretical population biology*, 130: 13–49, 2019.
- [3] E. Bairey, E. D. Kelsic, and R. Kishony. High-order species interactions shape ecosystem diversity. *Nature communications*, 7(1):1–7, 2016.
- [4] J. Cardy. Reaction-diffusion processes. *Lecture notes*, 100.
- [5] B. C. Cunningham and J. A. Wells. High-resolution epitope mapping of hgh-receptor interactions by alanine-scanning mutagenesis. *Science*, 244(4908):1081–1085, 1989.
- [6] M. M. Desai and D. S. Fisher. Beneficial mutation–selection balance and the effect of linkage on positive selection. *Genetics*, 176(3):1759–1798, 2007.
- [7] D. S. Fisher. Asexual evolution waves: fluctuations and universality. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(01):P01011, 2013.
- [8] T. Galla. Two-population replicator dynamics and number of nash equilibria in matrix games. *EPL (Europhysics Letters)*, 78(2):20005, 2007.
- [9] M. A. Gibson and J. Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *The journal of physical chemistry A*, 104(9): 1876–1889, 2000.
- [10] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry*, 81(25):2340–2361, 1977.
- [11] N. S. Goel, S. C. Maitra, and E. W. Montroll. On the volterra and other nonlinear models of interacting populations. *Reviews of modern physics*, 43(2):231, 1971.
- [12] N. Goldenfeld. *Lectures on phase transitions and the renormalization group*. CRC Press, 2018.
- [13] B. H. Good, I. M. Rouzine, D. J. Balick, O. Hallatschek, and M. M. Desai. Distribution of fixed beneficial mutations and the rate of adaptation in asexual populations. *Proceedings of the National Academy of Sciences*, 109(13):4950–4955, 2012.
- [14] P. Grassberger. On phase transitions in schlögl’s second model. In *Nonlinear Phenomena in Chemical Dynamics*, pages 262–262. Springer, 1981.

- [15] G. E. Hutchinson. The paradox of the plankton. *The American Naturalist*, 95(882):137–145, 1961.
- [16] H.-K. Janssen. On the nonequilibrium phase transition in reaction-diffusion systems with an absorbing stationary state. *Zeitschrift für Physik B Condensed Matter*, 42(2):151–154, 1981.
- [17] N. Kashtan, S. E. Roggensack, S. Rodrigue, J. W. Thompson, S. J. Biller, A. Coe, H. Ding, P. Marttinen, R. R. Malmstrom, R. Stocker, et al. Single-cell genomics reveals hundreds of coexisting subpopulations in wild prochlorococcus. *Science*, 344(6182):416–420, 2014.
- [18] M. Kondoh. Foraging adaptation and the relationship between food-web complexity and stability. *Science*, 299(5611):1388–1391, 2003.
- [19] T.-L. V. Lite, R. A. Grant, I. Nokedal, M. L. Littlehale, M. S. Guo, and M. T. Laub. Uncovering the basis of protein-protein interaction specificity with a combinatorially complete library. *Elife*, 9:e60924, 2020.
- [20] P. V. Martin, J. A. Bonachela, S. A. Levin, and M. A. Muñoz. Eluding catastrophic shifts. *Proceedings of the National Academy of Sciences*, 112(15):E1828–E1836, 2015.
- [21] R. M. May. Will a large complex system be stable? *Nature*, 238(5364):413–414, 1972.
- [22] M. Mézard, G. Parisi, and M. A. Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987.
- [23] J. Neidhart, I. G. Szendro, and J. Krug. Adaptation in tunably rugged fitness landscapes: the rough mount fuji model. *Genetics*, 198(2):699–721, 2014.
- [24] A. Nourmohammad, J. Otwinowski, and J. B. Plotkin. Host-pathogen coevolution and the emergence of broadly neutralizing antibodies in chronic infections. *PLoS genetics*, 12(7):e1006171, 2016.
- [25] M. Opper and S. Diederich. Phase transition and $1/f$ noise in a game dynamical model. *Physical review letters*, 69(10):1616, 1992.
- [26] O. Ovaskainen and B. Meerson. Stochastic models of population extinction. *Trends in ecology & evolution*, 25(11):643–652, 2010.
- [27] M. Parker and A. Kamenev. Extinction in the lotka-volterra model. *Physical Review E*, 80(2):021129, 2009.
- [28] M. Parker and A. Kamenev. Mean extinction time in predator-prey model. *Journal of Statistical Physics*, 141(2):201–216, 2010.

- [29] M. T. Pearce, A. Agarwala, and D. S. Fisher. Stabilization of extensive fine-scale diversity by ecologically driven spatiotemporal chaos. *Proceedings of the National Academy of Sciences*, 2020.
- [30] F. Roy, M. Barbier, G. Biroli, and G. Bunin. Can endogenous fluctuations persist in high-diversity ecosystems? *arXiv preprint arXiv:1908.03348*, 2019.
- [31] F. Roy, G. Biroli, G. Bunin, and C. Cammarota. Numerical implementation of dynamical mean field theory for disordered systems: Application to the lotka–volterra model of ecosystems. *Journal of Physics A: Mathematical and Theoretical*, 52(48):484001, 2019.
- [32] H. Schenk, H. Schulenburg, and A. Traulsen. How long do red queen dynamics survive under genetic drift? a comparative analysis of evolutionary and eco-evolutionary models. *BMC evolutionary biology*, 20(1):1–14, 2020.
- [33] N. Shores, M. Hegreness, and R. Kishony. Evolution exacerbates the paradox of the plankton. *Proceedings of the National Academy of Sciences*, 105(34):12365–12369, 2008.
- [34] S. Smale. On the differential equations of species in competition. *Journal of Mathematical Biology*, 3(1):5–7, 1976.
- [35] T. F. Thingstad. Elements of a theory for the mechanisms controlling abundance, diversity, and biogeochemical role of lytic bacterial viruses in aquatic systems. *Limnology and Oceanography*, 45(6):1320–1328, 2000.
- [36] C. Xue and N. Goldenfeld. Coevolution maintains diversity in the stochastic “kill the winner” model. *Physical review letters*, 119(26):268101, 2017.
- [37] L. Yan, R. A. Neher, and B. I. Shraiman. Phylodynamic theory of persistence, extinction and speciation of rapidly adapting pathogens. *Elife*, 8:e44205, 2019.

Chapter 5

Conclusions

In this dissertation, we have examined two classes of eco-evolutionary models and have simulated and analyzed their long term dynamics. In both cases, we found important simplification and distinct phases by looking at high-dimensional limits. Exploiting such limits has been a key thread in the development of physics over the past centuries. However, there is much more work to do on the models presented here before these models might be elevated to the level of sophistication and predictive power that statistical mechanics and thermodynamics possess. Importantly, despite ongoing progress, there is still a fairly wide experimental gap that must be bridged before a ‘unified theory’ of eco-evolutionary dynamics might be reasonably proposed. With the development of sequencing and gene editing technologies this sort of direct and quantitatively predictive role for ‘top-down’ theory [3] in evolving biological systems seems tantalizingly close.

However, it is reasonable to ask whether such a theory is possible or even desirable. While many biological systems have key large parameters (population size, genome size, etc) and so can be thought to reside in simplified ‘asymptotic’ regimes, much of the interesting and beautiful phenomena in biology reside in the particular details, or in the *meso-scale*. To spin off Phil Anderson’s famous assertion [1], whereas in condensed matter physics ‘more is different,’ in biology one might also say that *different is more*. A few examples to demonstrate what I mean. Consider the accidental discovery of antibiotics on a contaminated plate [2]. Or the revolutionary biotechnologies discovered to be already existing within the proteome of bacteria [4]. We need only watch a nature documentary to feel a sense of awe at the natural world. Any sufficiently coarse-grained model of population evolution will inherently miss the heterogeneity and difference that makes biology so rewarding to study, both personally and scientifically. In studying biology, in some sense we *hope* to be surprised.

As such, while we develop physics style models of biological processes, it is crucial to step back and consider the big picture, and what our goals might be. Biological models quite obviously play a different role than physics models, where something like the Standard Model precisely constrains what one might expect to observe at the Large Hadron Collider. This is in some sense strictly impossible in a biological system – wholly new biological discoveries can be made every day if one simply looks closely enough! Instead modeling can serve a different

role, demonstrating or enumerating the seemingly infinite spectrum of possibilities. And this is no less important. By enumerating possibilities, biological models can serve as illumination (cf Dobzhansky) and inspiration – both in the concrete sense of laying the groundwork for future scientific inquiries (e.g. future experiments) but also in a more abstract humanistic sense.

It is fitting then to close by discussing the quote that opens this dissertation. In Ursula K LeGuin’s science fiction classic *The Word for World is Forest* [5], Raj Lyubov is a type of extraterrestrial anthropologist tasked with studying the humanoid inhabitants of the planet Athshe. He is struck by the integrated lifestyle that the Athsheans possess with their diverse, forested world (which stands in contrast with the barren Earth in this story). In the course of the novel, the planet Athshe, though primitive by the humans’ standards, is nonetheless able to resist subjugation by technologically advanced human invaders through interconnection and interdependence. In a sense, the Athsheans are made strong and made whole by diversity. In this dissertation, we considered, in various forms, the crucial and potentially strongly beneficial effects that diversity generation might have on the dynamics of model ecosystems. On our Earth, as we develop our natural lands, homogenize the genetic stock of our crops and continually confront the prejudices within our societies, it seems like an opportune time to consider the importance of diversity.

5.1 Bibliography

- [1] P. W. Anderson. More is different. *Science*, 177(4047):393–396, 1972.
- [2] A. Fleming. On the antibacterial action of cultures of a penicillium, with special reference to their use in the isolation of b. influenzae. *British journal of experimental pathology*, 10(3):226, 1929.
- [3] B. H. Good and O. Hallatschek. Effective models and the search for quantitative principles in microbial evolution. *Current opinion in microbiology*, 45:203–212, 2018.
- [4] M. Jinek, K. Chylinski, I. Fonfara, M. Hauer, J. A. Doudna, and E. Charpentier. A programmable dual-rna-guided dna endonuclease in adaptive bacterial immunity. *science*, 337(6096):816–821, 2012.
- [5] U. K. LeGuin. *The Word for World is Forest*. Berkley Books, 1976.