Dissertation

submitted to the
Combined Faculty of Natural Sciences and Mathematics
of Heidelberg University, Germany
for the degree of
Doctor of Natural Sciences

Put forward by

Yunus Sevinchan

born in Münster, Germany
Oral examination: 3 November 2021

Evolution Mechanics

and

Perspectives on Food Web Ecology

Referees: Prof. Dr. Kurt Roth

Prof. Dr. Björn Malte Schäfer

Abstract

The question how hierarchically modularised structures arise from simpler ones is of central importance when desiring to understand our world.
To that end, I present the Evolution Mechanics framework which aims to find a concise description of the mechanisms by which evolutionary systems unfold into hierarchically organised modules. While inspired by the evolution of biological life, Evolution Mechanics is abstracted from it and takes a more general perspective, providing a consistent language to address the fundamental processes giving rise to the complexity we observe all around and within us.
In a second part, I study the evolution and behaviour of ecological interaction networks. Using an evolutionary food web model, I investigate the structures that arise within it, its response to local and global perturbations, and its capacity to be resilient against these perturbations. These studies not only illustrate aspects of Evolution Mechanics, but stress the importance of taking into account evolutionary processes when aiming to understand these systems.

Zusammenfassung

Die Frage, wie aus ursprünglich einfachen Strukturen welche mit hierarchischer Organisation und Modularisierung entstehen können, hat zentrale Bedeutung für ein besseres Verständis unserer Welt.
Um dieser Frage nachzugehen, stelle ich in dieser Dissertation die *Evolutionsmechanik* vor: Eine Theorie, die es zum Ziel hat eine treffende Beschreibung der Mechanismen zu finden mit derer evolutionäre Systeme sich zu hierarchisch organisierten entfalten. Evolutionsmechanik ist inspiriert von der Evolution biologischen Lebens, aber wählt eine abstrahierte und generalisierte Beschreibungsweise. Die dadurch ermöglichte Sprache kann zur Beschreibung der Prozesse genutzt werden, die zur Entstehung der Komplexität in unserer Welt führen.
Im zweiten Teil meiner Dissertation beschäftige ich mich mit der Evolution ökologischer Interaktionsnetzwerke. Mit der Hilfe eines Modells für evolutionäre Nahrungsnetze untersuche ich die entstehenden Strukturen, die Reaktion des Netzwerkes auf Perturbationen, und ihre Resilienz gegenüber diesen Perturbationen. Diese Untersuchungen sind nicht nur Illustration für zentrale Aspekte der Evolutionsmechanik, sondern heben auch die Wichtigkeit evolutionärer Prozesse bei der Untersuchung dieser Systeme hervor.

Table of Contents

Bibliography

Publications

Appendix

Acknowledgements

Prologue

Our world is of magnificent diversity and complexity. Understanding what it is made up of – small and large – and how it behaves has captivated generations of philosophers and scientists and will probably continue to do so as long as conscious life exists.

A remarkable aspects of the current state of this world is that it must have emerged from a simple primordial situation. Over the course of a few billion years, with some periods more eventful than others, this situation evolved into what we see today: Layers upon layers of hierarchically organised structures in the shape of simple cells, complex cells, multicellular organisms, animal societies, and human culture. Some of these structures have emerged a long time ago and shaped the face of planet Earth ever since; others are more recent and constantly appear, adapt, or vanish. While a variety of processes take place at each of these levels, this *Unfolding of Life* is not without patterns, thus offering a chance to describe the fundamental mechanisms giving rise to the complex structures we observe.

The emergence of humankind is another result of the Unfolding of Life and, in the long history of the planet, a rather recent one. Nonetheless, this event considerably transformed the Earth system, most notably since the industrial revolution and the following rapid technological development of the last century, continuing today. By now, the effects of human civilisation pertain to practically all complex life on Earth and have caused not only large-scale biodiversity loss in the planet's ecosystems but also threaten humankind's own habitat. Undoubtedly, humankind faces stark challenges in the years to come; this may serve as yet another motivation for better understanding the intricacies of this world and our role in shaping it.

How these interconnected hierarchical structures came about and how they behave is the larger context of this thesis and I hope to contribute to a better understanding of these aspects of our world. Specifically, the following two questions describe the bigger picture of the two parts of my thesis: What are the fundamental processes that describe the Unfolding of Life? Which processes shape the emergence of ecosystems and how do they react to changes in their environment?

In the first part of this thesis, I motivate and present a novel theoretical framework called *Evolution Mechanics*. The primary aim of this framework is to find a concise yet comprehensive description of the mechanisms by which evolutionary systems unfold into hierarchically organised modules. While inspired by the evolution of biological life, Evolution Mechanics is abstracted from it and takes a more general perspective, highlighting the patterns between different manifestations of evolutionary systems. On the chosen level of abstraction, this framework provides a consistent language to address the fundamental processes giving rise to the complexity we observe all around and within us. Notably, many well-understood observations and theories exist in the study of evolutionary systems. Evolution Mechanics is an attempt to locate, interconnect, and extend these, thus contributing a fresh perspective on the evolutionary unfolding of our world.

In the second part, my aim is to study the structure and behaviour of evolving ecological systems and how they react to changes in their environment. The interest in this topic is twofold: For one, the ecosystems we observe today are the result of a long evolutionary development and thus constitute a fitting example by which to illustrate the ideas of Evolution Mechanics. Second, given the strong anthropogenic effects on ecosystems, better understanding their behaviour in response to these effects is more pressing than ever; explicitly taking into account the evolutionary mechanisms that generated the interactions within ecosystems may play an important role in that. In the bigger picture, central questions are: How do ecosystems evolve into the complex interaction networks we observe them to be? How do these systems react to perturbations and which circumstances or mechanisms may make them resilient to these changes? To contribute to the understanding of these questions, I investigated a numerical model of food web evolution, the results of which I present in the second part of this thesis.

While operating on different conceptual levels, the two parts of this thesis are synergistic in that Evolution Mechanics provides a language and a framework to study evolutionary food web ecology and food web ecology illustrates and informs the abstractions made in Evolution Mechanics. Furthermore, the parts share the idea that in order to understand a system's structure and its behaviour, the evolutionary processes that give rise to it need to be taken into account.

Part I

Evolution Mechanics

I.1

Introduction & Motivation

The quest to find fundamental processes that describe the universe is perhaps the most central aspect of physics. On that quest, scientists have ventured down towards the smallest of scales to comprehend the elementary constituents of our world, and out towards the largest of scales to understand the past, present, and future of the universe at large – with many questions yet to answer. The realm of complex systems opens up another dimension along which to study the intricacies of the universe.

As such, the question posed in the prologue is deeply rooted in physics: What are the fundamental processes that describe the Unfolding of Life?

The first step towards answering this question is to isolate the key phenomena that such a theory would need to describe. Abstracting away from the particularities of individual organisms, the central observation from the Unfolding of Life is that it brought forward life forms and structures that are highly modularised and hierarchically organised. Consider the example of animals, which are comprised of several modularised organs, which themselves are made up of conglomerates of cells, which in turn consist of various organelles, further traversing down a hierarchy. *Lenton and* A. *Watson* [2011] denote this nestedness with the analogy of 'Russian Dolls', with smaller dolls being associated with the smallest modularisation levels and subsequent dolls engulfing the previous structures.

The role of modularisation and hierarchical nesting in allowing for complex structures to form is aptly represented by the watchmaker analogy of Herbert A. *Simon* [1962], where two watchmakers assemble a complex object with and without the use of subassemblies. Given a certain rate of destructive disruptions in the assembly process, an object with intermediate stable structures – modules – can be assembled much faster and more reliably than an object which has to be assembled all-at-once[1]. It follows that, in general, the more complex the structure of an object is, the more unlikely it is that it may have emerged *without* modularised substructures. For instance, it would be practically impossible for a multicellular organism to arise from the successive assembly of individual atoms.

---

[1] Be aware that, elsewhere, the active phrasing of the watchmaker parable is used as an analogy for intelligent design of complex forms – however, *Simon* stresses that random processes (assembly and disruptions) suffice to understand why emergence of modularised objects are more likely than objects without intermediate stable states.

Remarkably, even the relatively ancient structures of biological life like simple cells are already too complex to have emerged spontaneously and all-at-once out of the primordial situation of the universe and planet Earth. Subsequently, the key phenomenon that is to be described by a theory of the Unfolding of Life is the emergence of hierarchically organised structures out of simple ones.

In the realm of biological life it is the process of evolution that describes this emergence. As famously stated by Theodosius *Dobzhansky*, ‘nothing makes sense in biology except in the light of evolution’ [1964]. This statement holds true today and despite the many additional facets that were found to play a role in the emergence of biological systems. It is the powerful generality of the core ideas of evolution – differential reproduction and aggregation of variation – that highlights how a wide range of other (pre- and post-biotic) systems can be considered evolutionary. In combination with the wide-ranging effect that repeated iteration of a process can have on a complex system, a conceivable path towards a description of the fundamental processes underlying the Unfolding of Life appears.

Building on the above considerations, *Evolution Mechanics* (EM) is a framework that aims to describe the Unfolding of Life and beyond. To that end, Evolution Mechanics formulates a set of fundamental mechanisms which, collectively, give rise to the repeated modularisation and hierarchy formation and, as a result: the complex structures we observe in this world.

While descriptions of evolutionary systems are plentiful even beyond biology, Evolution Mechanics is novel in the level of abstraction it chooses. It puts a focus on the evolutionary unfolding at large, highlighting the similarities between the processes occurring on different levels. As noted before, the motivation behind EM is not in replacing existing theories of the emergence of complex structures through evolution, but in locating and connecting them using a consistent conceptual language. Overall, the novel perspective provided by Evolution Mechanics may contribute to a better description and understanding of the processes that play a role in the Unfolding of Life – both in the past and in the future.

The Evolution Mechanics framework itself emerged collaboratively and through many thorough discussions with Kurt Roth, Benjamin Herdeanu, Harald Mack and other members of the TS-CCEES research group over the course of the last five years, including the time of my M.Sc. thesis in the group. What I present in this thesis is built on these discussions.

In the following, I will first give a more detailed overview of the Unfolding of Life. These observations act as the foundation for the Evolution Mechanics framework, which I present in Chapter I.3. Finally, in Chapter I.4, I will reflect on the framework itself, place it in the context of existing theoretical works, and discuss open questions as well as future research perspectives.

I.2

Unfolding of Life

The *Unfolding of Life* is the observation that the complex and hierarchically organised structures we see today arose out of simpler ones. It can be seen as a collection of phenomena throughout Earth history. These phenomena are what *Evolution Mechanics* (EM) aims to consistently describe.

To speak of the Unfolding of Life, requires a definition of life. At this point, a pragmatic definition suffices, namely that of biological life: First cells, capable of reproducing and accumulating variation; any system that forms out of these will also be regarded as alive[2]. Note that the term *Evolutionary Unfolding* is used mostly synonymously.

The following overview of the evolutionary history of Earth is loosely inspired by work of Timothy M. *Lenton and* Andrew *Watson* [2011], where a more detailed account of these topics can be found. I will start out with the beginning of the Universe and the accretion of Earth and then continue towards the Origin of Life. A large part of the Unfolding of Life is that of biological life, from first cells to complex cells and multicellular organisms, later resulting in eusocial organisms. This brief journey through Earth history is then concluded with the emergence of humankind and their effect on the Earth system.

Given the breadth of these phenomena, it is not possible to go into great depths here, neither is it the aim. Instead, I want to highlight events in Earth history that are particularly insightful in the context of Evolution Mechanics; in some areas, this requires more detail than in others. Regardless of the chosen level of detail here, the presented topics offer fascinating journeys down the rabbit hole…

---

[2] I will come back to the question of how life may be defined in Section I.3.2.2.

I.2.1 The beginnings

As far as we know, the universe came into existence in the *Big Bang* roughly 13.8 Gy ago. The initial hot and dense universe rapidly cooled and expanded, with first baryons and antibaryons condensing after some 10^{-6} s. A mass annihilation of particles and their antiparticles followed, after which the universe was dominated by photons. Only a slight imbalance of matter and antimatter is the reason for the observable universe to contain any matter at all – a yet to be fully understood process coined *baryogenesis*. In the *primordial nucleosynthesis*, the left-over baryons then combined into stable nuclei of H, D, He, and some Li.

At about 380 ky, the universe had sufficiently cooled for the first neutral atoms to emerge, mostly hydrogen. With the formerly free electrons now bound to nuclei, the universe became transparent. This transition can be observed today as the *cosmic microwave background*: the radiation emitted from the newly formed atoms when they relaxed into their ground state.

The heterogeneities in how matter was distributed in the early universe then led to self-organised structure formation, with slightly denser regions becoming increasingly dense, eventually forming gas clouds, stars, and galaxies. Through *stellar nucleosynthesis*, stable isotopes of higher masses emerged; other heavy elements were (and still are) produced through supernovae.

At some 4.6 Gya, the Solar System and planet Earth condensed [*Judson* 2017]. Initially in a molten state, denser elements began sinking to the core of the planet and a crust of lighter elements began to solidify. Oldest evidence for a solid crust at that time are zircon rocks which were radiometrically dated to about 4.4 Gya [*Lenton and* A. *Watson* 2011]. A subsequent bombardment with water-rich meteorites from outer parts of the solar system brought a large part of the water, nitrogen, or other volatile substances to the early Earth.

Several properties of early Earth are of importance when considering the habitability of Earth and hence the potential for the Unfolding of Life: Firstly, Earth accrued in a distance to the Sun that allowed for liquid water (under reasonable atmospheric pressures). Second, meteorite bombardment brought enough water and other volatile molecules to Earth and plate tectonics removed a large part of the reductive potential from these outer layers. Finally, the molten core that creates Earth's magnetic field supplied a shield against strong cosmic radiation.

From the above conditions and the early history of Earth, *Lenton and* A. *Watson* put the earliest start of habitable conditions on Earth to the solidification of the crust ca. 4.4 Gya and the latest start to the end of the Late Heavy Bombardment 3.85 Gya.

This initial stage of the Unfolding of Life is dominated by physical aggregation of matter and its self-organisation into structures like planets. Notably, it is only until the end of primordial nucleosynthesis that the universe was in a state which realised all its possibilities: creating all nuclei that were energetically accessible. Henceforth, only an increasingly small subset of possibilities is realised. Unlike the relaxation processes that shaped the early universe and planet Earth, the biological life that emerged later on is comprised of structures that could not emerge spontaneously within the life time of the Universe.

I.2.2 Origin of Life

When, where, and how life on Earth emerged will presumably never be fully known. However, a growing body of direct and indirect observations gives hints about the time and conditions in which the first living cells were formed and there are a number of hypotheses matching these observations.

I.2.2.1 First evidence for cellular life

With today's knowledge, *prokaryotes* – simple single-celled organisms without an enclosed cell nucleus – emerged some time between 4.1 Gya and 3.8 Gya.

The earlier bound marks the oldest indirect evidence: a graphite inclusion in a zircon which is rich in ^{12}C isotope, an isotope that is characteristic of biogenic carbon fixation [*Bell* et al. 2015].

The later date, 3.8 Gya, is the earliest direct evidence for life on Earth [*Dodd* et al. 2017; *Schopf* 2006]. It is in the form of microfossils retained in so-called *banded iron formations*, which are believed to have their origin as a precipitate from *hydrothermal vents* (seafloor fissures ejecting geothermically heated water and minerals). Furthermore, banded iron formations are seen as indicators of an early form of anoxygenic photosynthesis, such biological activity being one explanation for oxidised iron in the predominantly reduced, low-oxygen environment of the early oceans [*Dodd* et al. 2017; *Lenton and* A. *Watson* 2011].

Apart from the study of fossils, analysis of genome data may yield insights into the evolution of early life. From the genetic differences between two species and with an estimate of mutation rates, *molecular clock analysis* allows to reconstruct when the species diverged, thus finding their common ancestor. By repeating this for many species and calibrating the results against the fossil record, a tree of life with estimates for absolute time values may be reconstructed. One particular point of divergence is of interest: the *last universal common ancestor* (LUCA), which denotes the species that is ancestral to all life on Earth today; the root in the tree of life.

In one such study using molecular clock analysis, *Betts* et al. [2018] found the LUCA to likely predate 3.9 Gya. This matches the above estimates for the emergence of life; however, there still may have been life before LUCA (and unrelated to it), a regime that is not accessible via genetic analysis.

As another example of a genetic study, *Weiss and Sousa* et al. [2016] investigated the protein-coding genes from ca. 2000 prokaryote genomes with the aim of reconstructing the protein families that were present in LUCA, thus allowing to draw conclusions about its physiology and habitat. They found that LUCA must have been anaerobic, thermophilic, H_2-dependent and living in a 'geochemically active environment rich in H_2, CO_2 and iron', an environment that is similar to that of hydrothermal vents, hence supporting the hydrothermal Origin of Life hypothesis[3].

I.2.2.2 How did the first cells emerge?

After having looked at when and where these early prokaryotes may have come into existence, the question of *how* remains.

May the first cells have spontaneously assembled from the molecules present in the early oceans? This hypothesis is generally discarded: Already these early life forms were far too complex for them to assemble by chance in the environment of early Earth.[4] At the very least, the first prokaryotes needed to be able to self-replicate; a process that works roughly like this in known prokaryotes: (i) uncoiling the DNA to prepare it for replication, (ii) create a copy of the DNA, (iii) grow the membrane and initiate segregation, and (iv) regrow the cell membrane and cytoplasmic content. In addition, the operation of the cell requires it to synthesise proteins using the information stored in the DNA. This process is carried out in the cell's ribosomes, which itself is a complex and optimised machinery.

Unless the first prokaryotes relied on a vastly different mechanism, which is unlikely, this illustrates that already rather sophisticated structures were necessary for the basic operation and replication of a cell: the biomolecular macromolecules (DNA, various forms of RNA) and protein-based mechanisms for cell division and operation (replication, transcription, and translation) – collectively, this is referred to as *biomolecular machinery* throughout this thesis.

Subsequently, the question of the origin of prokaryotes moves one level down: How did the biomolecular machinery evolve? In other words: how did the geochemistry of Earth's early oceans allow a transition to *bio*chemistry?

Martin and Russell [2003] proposed that this transition might have occurred in the highly compartmented structure of hydrothermal vents [*Russell and Hall* 1997]. In their model, the vents did not only provide an environment with strong gradients in temperature, redox potential, and pH; their compartmented structure allowed sufficient confinement of reactions to reduce loss of reactants into the ocean and the compartment walls acted as inorganic catalysts and persistent electron sources. Later studies of the properties of hydrothermal vents [*Martin and Baross* et al.

---

[3] Another possibility is that life originated before and/or existed elsewhere on Earth, but all record of it was destroyed in the Late Heavy Bombardment, with only those life forms in hydrothermal vents surviving. This is just one example of the difficulties in the search for the Origin of Life.

[4] This is not only improbable on the time scale of the early Earth, but on that of the life time of the universe. Importantly, the same statement applies already to many of their presumed biomolecular precursors.

2008] and the biochemistry of LUCA [*Weiss and Preiner* et al. 2018; *Weiss and Sousa* et al. 2016] further support this hypothesis.

The envisioned route from the initially available simple biomolecules to the DNA-RNA-based biomolecular machinery is through multiple potentially overlapping stages, each with successively more complex structures:

1. prebiotic auto-catalytic chemical reaction networks, producing more complex biomolecules [*Brueckner and Martin* 2020; *Xavier* et al. 2020]

2. an era where RNA-based replicators and enzymes dominated, the so-called *RNA world* [*Higgs and Lehman* 2014; *Takeuchi and Hogeweg* 2012]

3. the *RNP world*, where ribonucleoproteins (RNPs, a complex of RNA and simple RNA-binding proteins) took on successively larger roles

4. the *biomolecular machinery* we know today: protein biosynthesis through *transcription* of DNA-based information into mRNA template molecules and *translation* into amino acid sequences

While an in-depth review of how these stages gave rise to the biomolecular machinery would be out of scope for this thesis, I will address two aspects in more detail below: First, how a simple form of evolution can occur in auto-catalytic chemical reaction networks (see Section I.2.2.3), which is believed to have been an important stepping stone for the biomolecular machinery to arise. Second, the basic ideas behind horizontal gene transfer (see Section I.2.2.4) and its role in the evolution of the universal genetic code.

Today, two domains of prokaryotes are distinguished: bacteria and archaea. While sharing many properties, archaea differ substantially from bacteria in their membrane composition and in their ribosomal RNA [*Woese and Fox* 1977]. Furthermore, archaea are the only organisms known to produce methane via anaerobic respiration and are able to survive in more extreme environments than bacteria.

Martin and Russell [2003] also propose a potential path for the emergence of the two prokaryote domains, which became essential in the origin of eukaryotes (see Section I.2.4). Given their similarity in structure and shared biomolecular machinery, the authors propose that they both evolved in the environment of hydrothermal vents: The confined structure allowed the shared evolution of most of the biomolecular machinery, but also allowed subsequent independent inventions of their respective membrane synthesis and cell wall chemistry. Both proto-cell structures were initially only able to survive in the environment of the vents. Eventually, so the theory, they became free-living cells by replicating the cell wall chemistry present in the pores of the hydrothermal vents and substituting the wall with a lipid membrane [*Martin and Russell* 2003].

I.2.2.3 Pre-biotic evolution

In [*Vasas* et al. 2012], the authors use a linear polymer model to demonstrate how ACS can perform an early form of evolution if they are embedded in a compartmentalised environment. They isolate so-called *viable cores* as the evolutionary units: auto-catalytic loops where each species[5] can be produced directly from a food set (a set of abundantly available species).

These cores are associated with a genotype, as they can reproduce all other parts of the core. Furthermore, viable cores produce a periphery – other species that are not part of the loop but which will always be created once the core exists; this periphery is associated with the phenotype of a viable core.

Vasas et al. [2012] find that the viable cores can act as attractors within an ACS and that there can be multiple attractors. Importantly, these attractors can be selectable, such that a change in species composition (a novel species arising through a rare reaction) or the stochastic loss of a reaction or species can lead to the system flipping to another attractor. In order for the reaction networks to accumulate adaptations, there is one further requirement: The reaction network needs to be embedded in a compartmentalised environment with small rates of species exchange between the compartments. Such a setting allows competition between cores within one compartment, but also between compartments, for instance if a core species finds its way to a different environment and ignites the production of the remainder of the core species there. In effect, adaptations that occurred in one compartment can be selected for or against and may accumulate – a simple form of evolution.

The prospect of simple evolutionary units arising in auto-catalytic reaction networks is especially interesting in the context of *Origin of Life*, as it offers a mechanism by which the biochemical machinery may have evolved: Instead of assembling spontaneously (which would be highly unlikely), it may have evolved as an accumulation of adaptations, selected for by the general mechanisms of simple reaction networks.

While *Vasas* et al. [2012] studied this in the context of chemical reaction networks, they made only fairly general assumptions which translate well to other kinds of reaction networks. In the wider context, their observations illustrate how even simple systems are capable of behaving like (very simple, often structurally limited) evolutionary units, if a few key conditions are fulfilled; see [*Nghe* et al. 2015] for a more thorough study of these conditions.

I.2.2.4 Horizontal Gene Transfer

A notable process taking place in the emergence of the biomolecular machinery all the way to the emergence of eukaryotes is that of *Horizontal Gene Transfer* (HGT).

---

[5] The term species refers to an abstracted concept at this point, not the biological species; see Appendix A.1 for definitions.

It denotes the transmission of genetic information not from parent to offspring ('vertically'), but between organisms. This process plays an important role in unicellular organisms, for instance in the spread of antibiotic resistance between closely-related bacteria [*Frost* et al. 2005; *Koonin and Wolf* 2012; *Poole* 2009].

Furthermore, HGT may have been crucial in the evolution of the genetic code itself, i.e.: of the evolution of the codon table, and the translation and transcription mechanisms. The remarkable aspect of the *standard genetic code* (SGC) is that it is universal among all life on Earth (apart from very rare or isolated exceptions), are highly robust, and show a relatively high level of fault-tolerance [*Koonin and Novozhilov* 2017]. It is certainly not a random code, frozen in place, but it evolved to an optimal form. Specifically in the context of the origin of life – which brought forward the vertical transfer of heritable information via the biomolecular machinery and the SGC – the question arises how the SGC evolved in the first place.

In an influential paper, *Vetsigian* et al. [2006] studied the evolution of the SGC and demonstrated that including the process of HGT into this early stage of evolution produced a more optimal, universal, and robust code. They argue that, before the existence of an SGC, the code was comparably ambiguous and it was the robustness to ambiguity that gave an evolutionary advantage. At that stage, HGT allowed sharing of information, which were only beneficial if they could be properly transcoded. By unifying the innovation-sharing protocol, a collective advantage arose in that innovations could be more easily shared, thus leading to the optimisation of the code.

Based on their observations, they propose to distinguish the origin of life into three phases, depending on the nature of the evolutionary dynamics in each phase: (i) A phase with high tolerance to ambiguous code; followed by (ii) a phase of rapid innovation sharing, reducing the tolerance to ambiguity as a result of the growing complexity; leading to (iii) a phase where 'refinement superseded innovation' and which thus became dominated by vertical information transfer. This last transition is typically referred to as the *Darwinian threshold* [*Woese* 2002], where the previously community-driven innovation-sharing phase (sometimes referred to as *Lamarckian phase* [*Vetsigian* et al. 2006]) is superseded by vertical descent as described by Darwin.

While the term HGT (in the strict sense) only refers to the biomolecular mechanisms of gene transfer, an expanded notion of the term includes any form of non-vertical transfer of heritable information and is not exclusive to DNA and the biomolecular mechanisms. In this thesis, HGT is sometimes used outside of the context of the SGC; in such cases, it denotes the general process of innovation-sharing.

I.2.2.5 Summary

The key takeaways for the Unfolding of Life from the emergence of prokaryotes pertain to the mechanisms that are presumed to have played major roles: modularisation, confinement, auto-catalysis, and innovation-sharing. Furthermore, as demonstrated with the examples of viable cores in auto-catalytic reaction networks, even simple systems may be interpreted as being subject to evolutionary dynamics.

I.2.3 Great Oxidation Event

The *Great Oxidation Event* (GOE, also called Great Oxygenation Event) denotes a period roughly 2.4-2.1 Gya in which the composition of Earth's atmosphere changed from weakly reducing – mostly N_2 and CO_2 with practically no O_2 – to an oxidising one that included O_2. Despite qualified as 'Great', the increase in atmospheric O_2 by the GOE stayed far below today's levels [*Holland* 2006; *Lyons* et al. 2014]. Yet, it is considered a crucial milestone in Earth history, because it indicates the emergence of organisms that used oxygenic photosynthesis for their metabolism. By this process, they not only transformed the chemical composition of early Earth but also lay the foundations for more complex life forms depending on the considerably more efficient oxygen-based metabolic pathways [*Lenton and* A. *Watson* 2011].

Photosynthesis refers to a set of light-dependent ('light') and -independent ('dark') chemical reactions that oxidise an electron donor and thereby fixate carbon from CO_2. If the electron donor is water, this process is called *oxygenic photosynthesis* and oxygen is a by-product of the overall reaction:

$$\mathrm{CO_2} + \mathrm{H_2O} + 8\nu \longrightarrow [\mathrm{CH_2O}] + \mathrm{O_2} \tag{I.2.1}$$

The carbohydrate is the result of the light-independent reactions and is typically stored in form of a sugar like $C_6H_{12}O_6$. Its synthesis is driven by the reduction of NADPH and ATP, widely used biomolecules that are the products of the light-dependent reactions.

Oxygenic photosynthesis is believed to have evolved in ancestors of today's cyanobacteria and predate the beginning of the GOE [*Betts* et al. 2018; *Lyons* et al. 2014]. While the timing of the GOE starting roughly 2.4 Gya is largely undisputed, the emergence of oxygenic photosynthesis is matter of heated debate, with their latest emergence thought at 2.5 Gya but presumptive evidence going back hundreds of My further [*Lenton and* A. *Watson* 2011, ch. 10]. Specifically, if oxygenic photosynthesis evolved at an earlier time, the delayed rise of atmospheric O_2 requires explanation. It is currently thought that a number of factors may have played a role, like nutrient shortages or various carbon burial and buffering processes [*Lenton and* A. *Watson* 2011; *Lyons* et al. 2014; *Olejarz* et al. 2021].

The oxygenation of the oceans and atmosphere by early cyanobacteria had a profound impact on system Earth: The rise in atmospheric oxygen led to the oxidation of atmospheric CH_4 to CO_2, which is a much weaker greenhouse gas than methane. In combination with the lower solar radiation at that time, this appears to have caused a ‘Snowball Earth’ that lasted for roughly 300 My.

Another result of the oxygenation was a mass extinction of obligately anaerobic life forms for which oxygen was toxic. However, higher oxygen levels also opened new potentials for evolution: For one, the changed geochemistry made more nutrients become available. Second, it allowed for life forms to emerge which relied on an aerobic respiration – a substantial ingredient for the emergence of complex life due to the higher energy efficiency. In addition, the formation of the ozone layer provided protection against strong UV radiation, hence facilitating the evolution of terrestrial life forms [*Lenton and* A. *Watson* 2011, ch. 10].

In the larger context, the GOE can be seen as an example of a key innovation sparking environmental changes on the planetary scale. While triggered by an innovation, the delay between the emergence of oxygenic photosynthesis and the rise of atmospheric oxygen illustrate the importance of considering the state of the whole system and the feedback processes between different parts of the system.

I.2.4 Eukaryotes

The domain of *eukaryotes* not only makes up the majority of the biomass on planet Earth [*Bar-On* et al. 2018], it also comprises a vast diversity of complex organisms and practically all multicellular life forms, including the kingdoms of animals, plants, and fungi.

The following description of the properties and origin of eukaryotic cells is only a glimpse into the world of the eukaryotes, their composition, behaviour, and the living conditions they appear in. For the Unfolding of Life, their origin is of particular importance, which is thought to have happened through repeated endosymbiosis of different prokaryote lineages; a remarkable example of how the cooperation of several different species can lead to their integration into a vastly more complex and capable organism.

I.2.4.1 Properties

Eukaryotic cells are much more complex structures than prokaryotic cells: For one, they are much bigger, typically with a volume many thousand times larger than that of prokaryotes. The eponym of the eukaryotes is its *cell nucleus* – in contrast to the free-floating DNA of prokaryotes, the genetic material of eukaryotic cells is confined to a membrane-enclosed region that isolates it from the rest of the cell. Outside the nucleus, eukaryotes contain a wide range of complex cell organelles like *mitochondria*, which perform aerobic respiration to produce the ATP needed elsewhere in the cell[6]. Plant or algae cells further contain *plastids*: organelles with their own DNA which perform additional actions like photosynthesis (in *chloroplasts*) or pigment synthesis (in *chromoplasts*). Also, eukaryotes may live as unicellular organisms or form a multicellular organism; in fact, almost all multicellular organisms are eukaryotes – more on this in Section I.2.5.

Replication of eukaryotic cells may happen either asexually through *mitosis*, or sexually through *meiosis*. In sexual reproduction, genes from both parent cells' genomes are randomly recombined; by this process, considerably more variation is introduced, allowing for faster evolutionary adaptation.

Eukaryotes are considerably more energy-efficient in their operation than prokaryotes while using a much smaller fraction of their cell volume for metabolism [*Martin* 2017]. This is mainly due to the existence of the mitochondria, which allow eukaryotes to perform respiration not at the cell membrane but in a strongly localised and efficient fashion. Furthermore, the eukaryotes' larger size reduces redundancies in protein synthesis. Assuming that 1000 prokaryotic cells are of roughly the same volume as one eukaryotic cell, *Martin* makes the following comparison:

> Using the same energy per volume, the 1000 prokaryotes have to make 1000 cell walls, 1000 genome copies, 1000 cell division machineries and everything else required to make 1000 energy converting self-replicating systems. The eukaryote has solved the energy problem with 10 % of the cell volume, and does not need to synthesise 1000 cell walls, 1000 host genomes, 1000 plasma membranes or 1000 lipopolysaccharide layers. [. . .] The eukaryote has three orders of magnitude more energy per gene than the prokaryote. *Martin* 2017

By freeing energy and cell volume, *Martin* argues, the eukaryotic cell's energy budget allows for *tinkering*: synthesising non-essential proteins which *may* end up leading to innovations, without being a direct disadvantage if the synthesis does not provide an evolutionary benefit. This is unlike in prokaryotes, which have a much more restrained energy budget.

I.2.4.2 Origin of eukaryotes

Symbiogenesis The origin of eukaryotes is today thought to be the result of symbiosis and merging of different kinds of prokaryotes. First ideas of *symbiosis* – the living-together of unlike organisms – being the crucial mechanism in the origin of eukaryotes go back to the early 20th century [*Mereschkowsky* 1910]. In the 1960s, Lynn Margulis (then: Lynn *Sagan* [1967]), further substantiated and advanced this

---

[6] As not uncommon in biology, there are exceptions to this, albeit rare: Some eukaryotes without mitochondria exist, but they are presumed to either have lost the mitochondrion at some point in their evolutionary history or they have organelles that can be seen as reduced forms of mitochondria [*Martin and Russell* 2003].

theory[7]. While details are still debated [*Doolittle* 2020; *Martin* 2017; *Martin and Garg* et al. 2015], the *symbiogenesis* theory is the most widely accepted theory explaining the emergence and structure of eukaryotes.

Hints for this theory arose from the study of the eukaryote substructure, specifically triggered by properties of the chloroplasts and mitochondria: Both these organelles have their own genome and membrane and thus show strong resemblence to prokaryotes; the chloroplasts in particular show a high similarity to photosynthesising cyanobacteria. This led *Martin and Russell* [2003] to refer to eukaryotes as 'cells with something inside'. However, the associated structures in eukaryotes lack the ability to replicate on their own and are strongly dependent on the rest of the cell for their functioning.

Symbiogenesis proposes that these organelles were formerly free-living prokaryotes which went into symbiosis with a large host prokaryote. The symbiotic relationship might have been such that the host profited from the waste products of the free-living cell; to facilitate the interaction, the host may then have evolved the ability to provide it with nutrients. At one stage, the symbiont became encapsulated by the host prokaryote, thus becoming an *endosymbiont*. Eventually, the endosymbiont lost functionality that was of no benefit in the environment of the host cell, e.g. because the host cell started to provide it; also, part of its genome may have become tranferred to the host's genome in a process called *horizontal gene transfer* (HGT) [*Keeling and Palmer* 2008]. This allowed a drastic reduction in its genome, with the endosymbiont no longer being able to survive outside the host cell and the host cell controlling its function and reproduction – thus becoming a fully integrated part of the cell: an *organelle*.

For animal and fungi eukaryotes this process is presumed to have happened once, leading to the assimilation of a proto-mitochondrion; in many theories, the endomembrane structure formed as a part of this process as well. For plant cells, a second endosymbiosis is proposed to have assimilated proto-cyanobacteria, leading to the formation of the chloroplasts [*Martin and Garg* et al. 2015].

How exactly and how many of these merging processes occurred is unknown – *Martin and Russell* [2003] call it 'one of biology's messiest problems' and there are a wide range of hypotheses surrounding it; for a good overview see [*Martin and Garg* et al. 2015].

Further topics of debate are about the order these processes took place in, which properties the endosymbionts had, and with which domain of prokaryotes[8] they are to be associated with, the *archaea* or the *bacteria* (or neither). Genetic studies show that eukaryotic cells are overall more closely related to archaea, with their mitochondria and plastids being more closely related to bacteria [*Gribaldo and Brochier-Armanet* 2019; T. A. *Williams* et al. 2019]. Hence it is now said that

---

[7] Her work is not lacking in exceptionally well-drawn sketches of early eukaryotes, for example [*Sagan* 1967, p. 248].

[8] See Section I.2.2 for a brief description of their differences.

eukaryotes emerged from 'well within the domain of archaea' [*Betts* et al. 2018], successively acquiring proto-bacteria that became their organelles.

Fossil record In the fossil record, candidates of early eukaryotes are observed at about 1.7-1.6 Gya [*Isson* et al. 2018]. Less debated evidence of near-modern eukaryotes is found for 1.4-1.2 Gya, where fungi and algae are presumed to already have diverged [*Lenton and* A. *Watson* 2011]. The molecular clock study by *Betts* et al. [2018] places the emergence of eukaryotes to a time no earlier than about 1.8 Gya; however, whether this general methodology warrants such statements is put into doubt by *Lenton and* A. *Watson*.

Regardless of the time of their origin, eukaryotes seem to have played only a minor role in the Earth's oceans of 2.1 Gya. They are presumed to only have become ecologically dominant some time around 0.8 Gya [*Butterfield* 2015; *Isson* et al. 2018].

Tree of life The relatedness of species is often represented in a *phylogenetic tree*, a tree diagram extracted from the phenotypical or genetic characteristics of species. It is typically in the form of a binary tree, leaves being associated with observable species and the root with their common ancestor. If the diagram aims to represent all of life on Earth, it is called *tree of life*[9].

While this approach works reasonably well for higher life forms, which can be categorised more easily, it reaches its limits when aiming to represent the emergence of eukaryotes from archaea and bacteria [*Lenton and* A. *Watson* 2011]. Specifically, the non-vertical transmission of genetic information undermines the assumption of a binary tree [*Martin and Russell* 2003], and the process of endosymbiosis is difficult to represent in such a framework. Hence, the development of eukaryogenesis theories was accompanied by various new representations of the tree of life, which I want to briefly mention here.

Historically, it was *Mereschkowsky* [1910] who first realised a symbiotic nature of the emergence of eukaryotes and reflected this in his visualisation, which shows a tree with two symbiotic stages that denote information transfer between different branches of the tree.

In the context of the more discretised classification envisioned by *Woese and Kandler* et al. [1990, Fig. 1], purely binary trees became more wide-spread. While the debate about the number of basic domains appears to become settled in recent years [T. A. *Williams* et al. 2019], the representations as binary trees remain relevant for visualisation of phylogenetic data. For instance, in figures like [T. A. *Williams* et al. 2019, Fig. 2], distances in the tree inversely relate to genetic sim-

---

[9] An alternative metaphor to the tree was already proposed by Darwin: the coral. A coral is a more fitting metaphor in several aspects: there is no branching in their root, only their uppermost parts are alive, branches may merge, and their diameter may vary [*Podani* 2017]. As the tree metaphor is prevalent, it is also used throughout this thesis.

ilarity; this figure specifically supports the classifications of life into two domains and eukaryotes emerging from within the archaea.

In contrast to these binary trees, [*Martin* 2017, Fig. 3] shows a tree that does not aim to qualify whether the origin of eukaryotes was from primarily one of the prokaryote domains, but emphasises the symbiotic nature of the emergence of eukaryotes: The root of the tree is the LUCA of bacteria and archaea, drawn as an unstructured blob rather than a single point. From there, both domains fan out into many separate branches but continue to be subject to HGT. In the center of the tree, some of these branches come together and form a proto-eukaryote which subsequently picks up additional characteristics through further symbiotic branches. A feature of this visualisation by *Martin* is that individual lines from the LUCA are carried through, such that it is clear that the resulting eukaryotes are the result of multiple symbiotic processes – an idea also used by *Mereschkowsky.* This follows suggestions by *Lenton and* A. *Watson* to regard the tree of life as a 'tangled bush of life' [2011, Fig. 3.1] rather than a binary tree.

Overall, these visualisations exemplify how difficult categorisations become when trying to account for the endosymbiotic nature of the origin of eukaryotes. This is to some extent reminiscent of the situation described by *Goldenfeld* et al. [2017] of a 'collective network phase of life' – while the authors refer to the time before the LUCA and the origin of life itself, the origin of eukaryotes appears to also require a form of non-vertical inheritance channels.

I.2.4.3 Summary

In summary, the origin of eukaryotes appears to be a unique event [*Lane* 2011] in Earth history, intertwined with the evolutionary history of both archaea and bacteria but with fundamentally advanced capabilities in terms of efficiency and coordination. The observations made in the context of their emergence emphasise the interconnected and coevolutionary nature of this phase of life on Earth. Thus, as *Lenton and* A. *Watson* stress, no single date can be assigned to the origin of eukaryotes; rather than emerging suddenly, fully-formed, eukaryotes appear to have gained additional characteristics over a long period of time.

In the larger context of the Unfolding of Life, these eukaryotes formed the basis for substantially more complex (multicellular) life in the shape of plants, fungi, and animals. Also, it is the only known transition where individuals from different species combined into a higher-level collective.

I.2.5 Multicellular organisms

Multicellular organisms are organisms that are made up of multiple individual cells. The (independent) emergence of multicellular organisms was observed more than 25 times throughout Earth history and for both prokaryotic and eukaryotic cells. Earliest evidence for multicellular prokaryotes goes back at least to 3 Gya;

multicellular eukaryotes are presumed to have emerged before 1 Gya, with further diversification 700-600 Mya [*Grosberg and Strathmann* 2007]. The resulting structures are diverse, ranging from simple aggregates to highly complex organisms like fungi, algae, plants, and animals which can comprise billions or trillions of cells [*Pennisi* 2018]. While there are vastly more unicellular organisms living on Earth[10], multicellular life introduced a new organisation structure – and a fascinating variety of life forms.

Multicellularity can be evolutionarily advantageous alone by the larger size of the collective compared to the individual, allowing to hinder predation; in turn, it becomes advantageous for predators to also grow in size and evolve mechanisms to feed on larger prey. Apart from the size of the collective, individual cells in a multicellular organisms performing different functions can be beneficial. This divison of labour allows for cells to specialise on certain tasks like different metabolic processes.

The transition from unicellular to multicellular organisms appears to have been relatively easy, at least when compared to the transition from prokaryotes to eukaryotes. Yet, for plants and animals, it only occurred once [*Pennisi* 2018]. *Butterfield* [2015] argues, that in order for complex multicellular organisms to be stable enough, gene regulatory mechanisms needed to evolve in the unicellular organisms first. Only by adapting these regulation mechanisms on the scope of the multicellular organisms, they would gain an advantage over their single-celled ancestors.

One hurdle to multicellularity is that of genetic conflict between the cells in the collective. As *Grosberg and Strathmann* [2007] put it, 'the advantages of multicellularity depend upon cooperation among cells, but coop- eration invites cheating.' Subsequently, for collectives to have a benefit over solitary cells, the collective needs mechanisms to reduce genetic conflict or some form of policing mechanisms: For the latter, they may perform active detection of defector cells and exclude them from the collective. For the former, they can reproduce through a single-cell reproduction path, such that only a single genome is passed on to the next generation. As a result, there is reduced selection pressure within the multicellular organism because mutations will not be passed on, thus reducing evolutionary benefits from cheating. In addition, if a defector cell would become the basis of an offspring organism, it would become a defector within a collective of defectors – and would thus probably not be viable [*Grosberg and Strathmann* 2007, Fig. 2].

In general, the requirements for multicellular life to be possible are seen in cell adhesion (that allows to form the collective), cell-cell coordination (for instance, to coordinate division of labour or transport processes), and programmed cell death (to hold defective or malfunctioning cells at bay). All these functions may already have evolved in unicellular organisms; for the transition to multicellularity, adopting and potentially adjusting these mechanisms may have sufficed for an evolution-

---

[10] In pure numbers, this is certainly true. In terms of biomass it is hard to estimate; those given in [*Bar-On* et al. 2018] are very uncertain for bacteria and archaea, but allow for this case.

ary advantageous multicellular organism [*Grosberg and Strathmann* 2007; *Pennisi* 2018].

Notably, the transition to multicellularity has also been replicated multiple times in laboratory conditions and over surprisingly few generations [*Ratcliff* et al. 2012].

I.2.5.1 Complex multicellular organisms

A closer look at complex multicellular organisms – plants, fungi, animals – is worthwhile, given their remarkable ability to coordinate a huge number of cells.

A crucial ingredient for complex multicellular life is seen in the capabilities of the eukaryotes themselves, specifically the mitochondria. In addition to their high energy efficiency described above [*Martin* 2017], the genomic capacity of eukaryotes is five orders of magnitude larger than those of prokaryotes, allowing to encode vastly more heritable information [*Lane* 2011; *Lane and Martin* 2010].

A second aspect is that of communication and coordination between cells in the collective. To that end, *gene regulatory networks* (GRNs) affect the expression rates of certain genes in complex cells [*Staps* et al. 2019]. This is achieved, for instance, by synthesising transcription factors which attach to the promoter regions of a gene, thereby inhibiting or completely suppressing the mRNA transcription. On a more abstract level, GRNs can be understood as reaction networks which detect some input signal (from the cell or the wider environment), perform a computation, and create an output signal which regulates gene expression. GRNs may also contain positive or negative feedback loops.

In multicellular organisms, these regulatory circuits affect the behaviour of cells and can cause them to take a specific shape or perform a certain task in the collective. The GRNs are the local machinery that gives a cell some form of plasticity over its life cycle. On the level of the whole organism, communication between cells can then occur via chemical signalling through hormones; simply speaking, hormones bind to cell receptors and be interpreted as one of many inputs to a cell's GRN. This allows for a complex organism to grow according to a body plan, with initially identical cells specialising into a form the fulfills a certain purpose on the level of the multicellular organism.

Finally, another important property of complex multicellular life is that it is often host to a whole range of other life forms. This collective of a host organism and many symbiotic organisms[11] is referred to as a *holobiont*, with the combined genome of all these life forms comprising the *hologenome* of the holobiont [*Theis* et al. 2016]. The idea to describe organisms as holobionts is pivotal in that it recognises the multi-level nature of complex life and its embedding into an environment. However, it complicates the notion of the relevant evolutionary unit and may lead to several misconceptions, which are aptly summarised in [*Moran and Sloan* 2015, Box 1].

---

[11] Note that this may include parasites, which can strongly drive evolution of the symbiotic relationship.

As an example, humans and other mammals are host to a vast number of microbes, living on their skin, on mucous membranes, or in the gastrointestinal tract. They provide the host with different functionality (like nutrient uptake or protection), but are in constant flux and competition with each other – an ecosystem in itself [*Foster* et al. 2017]. In plants, the situation is similar, but with fungi playing a larger role [*Vandenkoornhuyse* et al. 2015].

Holobionts have different kinds of host-microbiome interactions and control mechanisms. For instance, in mammals, the host actively controls its microbiome ecosystem to keep it in check and facilitate an ecosystem that is beneficial for the host; *Foster* et al. [2017] characterise this as an 'ecosystem on a leash'.

I.2.5.2 Cooperation

Cooperation denotes a form of interaction between individuals that brings an overall evolutionary advantage. The advantage may either pertain directly to the individual, or to the group of cooperating individuals as a whole.

In the context of the transition to multicellularity, cooperation plays an important role. Here, I want to briefly review which approaches there are to describe the emergence of cooperation in a system; these considerations are not restricted to multicellularity in particular, but are also applicable in a wider context.

A central question is how cooperation may emerge and be maintained, especially in situations where individuals have an immediate disadvantage by cooperating compared to not cooperating ('defecting') – the famous Prisoner's dilemma situation. These kind of questions are studied in the field of *Evolutionary Game Theory* (EGT). As was shown by *Axelrod and Hamilton* [1981], such situations can be resolved if there are multiple interactions between individuals, allowing reciprocity, or mechanisms like mutual policing that suppress defection. In essence, a lot of research pertains to the question how delayed benefits can be conveyed within a group such that individual disadvantages ('costs') can be balanced for the greater good of the whole group.

In the evolutionary context, attainable benefits are in the form of evolutionary success. Thus, a particular aspect of this question relates to the level at which selection acts, because this affects which part of the group is evolutionarily successful. Proponents of *kin selection theory* or *inclusive fitness theory* [*Hamilton* 1964] argue that as soon as individuals in a group share genetic material, costs for an individual are balanced by its indirect propagation of genetic material through its kin: benefits are conveyed through genetic similarity. This was (and still is) a matter of heated debate [*Birch* 2017; *Birch and Okasha* 2014; *Nowak and McAvoy* et al. 2017]. A contrasting – and more general – approach is *multi-level selection theory* [*Nowak and Tarnita* et al. 2010; D. S. *Wilson and* E. O. *Wilson* 2007], arguing that selection may act on multiple levels, specifically including that of the group. D. S. *Wilson and* E. O. *Wilson* [2007] famously conclude: 'Selfishness beats altruism within groups. Altruistic groups beat selfish groups. Everything else is commentary.'

While a lot more can be said about the mechanisms that may give rise to cooperation, this brief overview suffices for this thesis. For a more in-depth study of this topic, refer to the work by Benjamin *Herdeanu* [2021].

I.2.5.3 Fraternal and egalitarian group formation

Having described the emergence of both eukaryotic cells and multicellular organisms, an important observation can be made: Group formation can happen in many ways, depending on the function the individual entities fulfil in forming the group. One frequently made classification [*Bourke* 2009; *Queller* 1997; *Queller* 2000] distinguishes groups that formed with functionally identical constituents which later diversified, from groups that formed with functionally diversified entities: In essence, the distinction is made using the order in which functional diversification takes place in relation to the change in the level of selection from the individuals to the group.

For *egalitarian* group formation, diversification occured *prior* to group formation, such that the group is made up of diversified entities coming together with each part being equally relevant for the success of the group (hence 'egalitarian'). If diversification occured *after* the group was formed, this is *fraternal* group formation: identical entities (hence 'fraternal') coming together and later diversifying within the group.

In [R. A. *Watson and Mills* et al. 2015, Fig. 1], this is referred to as 'two dimensions of change' on the transition from identical solitary particles to a heterogeneous collective. Eukaryotic cells would be classified as an egalitarian transition, multicellular organisms as a fraternal transition.

I.2.5.4 Summary

To summarise, given the many times it occurred, cells aggregating into multicellular collectives appears to have been a simpler evolutionary transition compared to the previous ones. Still, it was a highly consequential one, giving rise to all complex life.

The key takeaways from the transition to multicellular life are, that (i) individuals can aggregate and coordinate into a new organism that acts as a whole; (ii) the individual cells in these fraternal collectives can specialise to take up different roles within the group; (iii) single-cell reproduction paths or policing mechanisms are evolutionary strategies to avoid genetic conflict; (iv) multicellular life requires coordination mechanisms; and (v) complex multicellular life can become holobiontic.

I.2.6 Radiation & extinction events

Throughout Earth history, certain time periods show rapid changes in environmental conditions and biodiversity. The rates at which new species arose and became extinct are not uniform, but are a result of a multitude of feedbacks within the Earth system and individual ecosystems. Of course, some kind of 'background rate' of speciations and extinctions can be assumed, influenced by the intrinsic mutation rates. Environmental changes or key innovations may cause time periods with a high speciation rate (radiation events) or high extinction rate (extinction events).

One example for a so-called radiation event was the *Cambrian explosion*, a time period around 541 Mya during which a large diversity of species emerged, including practically all modern animal phyla, all within the span of a few million years [*Knoll and Nowak* 2017; *Lenton and* A. *Watson* 2011]. As just one of many examples, animals with bilaterian body plans[12] became significantly more wide-spread as a result of this radiation event [*Budd* 2003].

How 'explosive' this event actually was is put into question in [*Lenton and* A. *Watson* 2011, ch. 14], who argue that many of the environmental, ecological, and evolutionary changes may have occurred in the preceding hundreds of millions of years already. During the Cambrian explosion, they may then have came together and caused a larger diversification rate – but not as explosive as often portrayed.

For instance, the first fossils of bilaterian animals were dated to around 558 Mya, suggesting that the innovations associated with the Cambrian explosion may have evolved earlier already. These already-diversified organisms were then super-charged by the rise in atmospheric and oceanic oxygen as a consequence of the preceding *Neoproterozoic Oxygenation Event* (NOE). The authors trace the roots of the rapid diversification during the Cambrian explosion even further back to the period starting 1.8 Gya and denote the time in between as 'not-so-boring billion' [*Lenton and* A. *Watson* 2011, ch. 13] – this is in contrast to the widespread notion that it was a time of relative stasis and hence not particularly relevant for the Cambrian radiation.

There are many hypotheses and questions regarding the causes for the Cambrian explosion, ranging from the change in environmental conditions to the passing of a genetic complexity threshold (powered by the higher oxygen levels), which allowed to encode substantially more structural features than before. An ecological mechanism that may have played a role is a dynamic coined an *evolutionary arms race* between predator and prey species, leading e.g. to the emergence of hard shells as defensive structures against predators. Dynamics like these can cause a relatively large amount of genetic diversity within short time spans.

---

[12] Bilaterian organisms have a left-right symmetric body plan, additionally allowing to define head, tail, belly, and back.

Similarly, there were times in Earth history which led to a large number of species extinctions within a short time. One of these was already mentioned above: The GOE led to the extinction of most species relying on anaerobic respiration.

Since the time of the Cambrian explosion, five such major mass extinction events were observed. One was the *Permian-Triassic extinction event* (PT extinction) around 252 Mya, which is considered the largest such event since then and caused the extinction of a majority of species in all branches of life. The youngest extinction event on geological time scales is the one that led to the extinction of dinosaurs roughly 66 Mya.

The exact causes for these periods of rapid biodiversity loss is debated. For some events, the extinctions coincide with huge volcanic eruptions or asteroid impacts [*Knoll and Nowak* 2017] which lead to a short-time cooling of the atmosphere and reduced photosynthesis. In addition to these direct causes, there are indirect effects on the Earth system (e.g., ocean acidity or levels of greenhouse gases) which may be more wide-ranging and long-lasting.

A sixth mass extinction is associated with the present time: the *Holocene mass extinction*. It is caused primarily by the effects humankind has on the Earth system, like habitat destruction or over-exhaustion of ecosystems [*Payne* et al. 2020].

In the larger context, these radiation and extinction events are of interest because they shape the diversity of life forms on system Earth and hence: the potential evolutionary path. They can be triggered by key innovations and are strongly coupled to changes in the environment, which may also be caused by the organisms themselves. Through positive feedbacks, these radiation and extinction events may have different severities. Specifically radiation events illustrate the dynamics of evolutionary systems when new ecological niches become available.

I.2.7 Eusocial organisms

Another level of organisation becomes apparent when studying the behaviour of insects like ants, termites, or bees: In many of these species, individuals organise into colonies, which act as a whole and can hence be understood as a *super-organism*[13]. There are some parallels to multicellular organisms in that both are collectives of many similar individuals (fraternal collectives). However, what makes these insect species particularly interesting is that they have additional forms of communication and coordination that allow them to behave as a super-organism, specifically their social structure.

These capabilities are generally referred to as *eusociality* [*Nowak and Tarnita* et al. 2010; E. O. *Wilson* 2012], and are defined by the following properties of a colony: (i) individuals of overlapping generations are living together, (ii) they are collectively taking care of offspring individuals, and (iii) show reproductive division

---

[13] While this may appear similar to the previously described *holobiont*, the two terms denote different ideas, with *super-organism* referring to collectives of similar individuals.

of labour. The latter point includes that some individuals may become sterile and warrants a differentiation into different *castes* within the eusocial society.

While other animals, including humans, may also be regarded as exhibiting some form of eusociality, I will focus on less debated examples from arthropod species. Specifically, the following examples are of ants, which are particularly well studied in this respect – see [*Hölldobler and* E. O. *Wilson* 1990; E. O. *Wilson* 2012; E. O. *Wilson and Hölldobler* 2005] for more details.

Ants are estimated to have evolved around 150 Mya [E. O. *Wilson* 2012] and strongly diversified after the emergence of flowering plants 100 Mya, rising to ecological dominance around 60 Mya. Within this time, a multitude of mutualistic relationships and an intricate web of interactions and dependencies evolved, making ants (and other eusocial insects) become strongly integrated parts of many ecosystems. Notably, ants are estimated to make up roughly a third of insect biomass, while comprising only about 2 % of the number of insect species [E. O. *Wilson and Hölldobler* 2005].

Within a eusocially organised ant colony, there are many different roles, ranging from worker ants and soldiers to brood carers and the queen. Individual ants may be permanently assigned a role (called 'caste'), presumably through a certain type of nutrition during their larvae stage, or they may be temporarily assigned a task.

To coordinate tasks within the colony, ants communicate via pheromones. For instance, pheromone trails can guide ants towards food, signal attacks, or lead to the assignment of an ant to a certain task group like nest-building. They are also able to coordinate collective transport of larger objects or form ant bridges to cross obstacles. In addition, individual ants carry a colony-specific scent and ants from other colonies will typically be attacked; interestingly, these scents differ even between colonies from the same species of ants.

An important determinant of eusociality is the reproductive division of labour. In ants, there are many reproduction strategies, one of which is through a single queen and a few workers leaving a nest and founding a new colony elsewhere. Another is the nuptial flight, where many males and young queens leave the colony, mate, and attempt starting new colonies.

Furthermore, ants engage in agriculture. A prominent example is that of the leafcutter ants, which live in obligatory mutualism with a certain kind of fungus. Queens carry the fungus in their mouth when they found a new colony and henceforth cultivate it in their nest. The ants not only provide the fungus with shredded leafs for cultivation, but also try to keep it free from parasites (with the help of the secretes of a symbiotic bacterium). The fungus is then fed to the ant larvae.

Extending this list by one more remark, ants have been observed to form super-colonies which range over large areas and consist of thousands of individual colonies, including interconnected nests. Ants in these super-colonies still discern between colonies for mating and typically do not mix, but they appear to be non-aggressive towards ants from other colonies [*Steiner* et al. 2007].

From the capabilities of these ants and their ecological importance, it appears that eusociality was a highly successful innovation for these species. It allowed them – and other eusocial animals – to organise into super-organisms. Notably, this is much alike the emergence of multicellular organisms through fraternal group formation: A collective of many similar individuals organises such that individuals take up different roles within the collective, all made possible by new coordination capabilities.

I.2.8 Emergence and rise of humankind

On the time scale of the Unfolding of Life over the last billions of years, the emergence of humans some 300 kya and their subsequent spread across the planet is an incredibly short time span – even more so if considering that most of humankind's growth occurred as a result of the agricultural revolution, starting 12 kya, or the industrial revolution, starting a mere 200 years ago. Today, humankind is a dominant force on planet Earth, with effects ranging into all realms of their biotic and abiotic environment.

The following is a brief overview of the emergence and rise of humankind, based primarily on [*Lenton and* A. *Watson* 2011, ch. 19]. To keep this overview concise, I will sketch only one particular path through this chapter of human history and will do so in a matter-of-fact way, aware that there are many uncertainties regarding the time and context of certain events.

I.2.8.1 History of humankind

Modern humans, *Homo sapiens*, are animals (complex multicellular eukaryotes and, conceivably: holobionts) from the order of primates and the family of hominids. One of their ancestor species, *Homo erectus*, emerged about 2 Mya in Africa and spread throughout Europe and large parts of Asia, giving rise to further speciations. These early humans are thought to have lived in hunter-gatherer societies where their bipedalism and ability to run long distances benefited them in exhausting large prey. Also, they are known to have developed simple stone tools and coordinate the transport of raw materials over long distances, suggesting a proto-language to coordinate such efforts. With the development of more sophisticated stone tools some 400 kya, their (already comparably large) brain size further increased. As another skill, early *Homo* species learned to control and create fire, using it for cooking, keeping warm, or for hunting.

The emergence of *Homo sapiens* itself is dated to roughly 300 kya and to regions in South Africa [*Schlebusch* et al. 2017]. While not particularly different from their ancestors in anatomy, they learned to build advanced tools from bone and antlers and used them for specialised hunting. Apart from hunting, *Homo sapiens* societies show signs of having conducted ritual burials, creating works of art, and trading over long distances in some early form of economy.

All these skills were assisted by – or: coevolved alongside – their advanced capacity for communicating through spoken language, with syntax and universal grammar. With this ability came a new form of information storage, where information was encoded into stories which were told and re-told within the group. This opened an additional inheritance channel that went beyond that of biological inheritance: The group was now able to convey information between generations and even between groups.

Homo sapiens migrated to all continents but Antarctica between 130 kya and 10 kya, leading to the successive extinction of other hominin species in those regions. The arrival of humans in a new region further coincided with declines in many large mammal species – these can be considered to have been precursors to the aforementioned Holocene extinction.

Human evolution made a qualitative jump roughly 12-10 kya when it transformed from hunter-gatherer societies to sedentary communities of farmers in what is called the *agricultural revolution*. This transformation was made possible by a stabilisation of Earth's climate after the last ice age, which marks the beginning of the geological epoch of the Holocene. The domestication of crop – more precisely: the coevolution between farmers and their crop – allowed for groups of humans to settle in one place while generating a food surplus. Furthermore, these societies allowed for or even required division of labour, giving rise to first forms of larger, hierarchical societal structures. With the growing need to coordinate these societies, writing sytems emerged which were used for trade and book-keeping.

However, in these early days of agriculture, the benefit of settling down and taking up farming was only given in particularly fertile regions, like in Mesopotamia. Compared to the highly adaptive foraging communities, farming communities had disadvantages from malnutrition and more frequent disease outbreaks (due to higher population density and lack of sanitation). Furthermore, unsustainable irrigation practices or over-exploitation lead to low yields. Despite the generated food surplus, this may explain why the global population size did not rise noticeably in the first few thousand years of human agriculture. In the following years, however, population growth increased, driven by expanding agricultural areas. At approximately 4 kya, the total global human population is estimated to about 50 million; 500 years ago it was roughly 500 million.

In what is called the *industrial revolution* (or: second agricultural revolution), starting roughly 200 years ago, humans made the transition from solar-powered[14] societies to fossil-fuel-powered societies. While coal was already in use earlier, it became particularly relevant only with the invention of the steam engine in the late 18th century, drastically boosting productivity and agricultural yields. This accelerated the growth in population size, passing one billion some 200 years ago and doubling to two billion about 100 years ago. In a parallel development, all kinds

---

[14] All their energy needs were met by solar energy in the form of crop, animals, or wood.

of technological and scientific advancements occurred throughout the last 500 years – this is sometimes referred to as the *scientific revolution*.

Starting in the 19th century, humans became increasingly confronted with the waste products of their own civilisation, be it in the form of air pollution, acid rain, or polluted water sources. Later, the more globally acting waste products became apparent, particularly through greenhouse gases that changed the composition of the atmosphere in non-negligible ways. Accompanied by further technological advances throughout the 20th century, this time is often characterised as 'the great acceleration', referring in particular to the impact of humankind on the Earth system.

I.2.8.2 The Anthropocene

Today, the Earth is habitat to some 8 billion people and their presence – our presence – is noticeable practically anywhere on the planet. Within a few thousand years, and especially within the last century, human societies grew from living in scattered agricultural settlements to a global, interconnected civilisation, affecting all other life on the planet.

These effects appear as traces of manufactured materials in sediments, including radionuclides from nuclear weapons, as substantially increased concentrations of greenhouse gases, or as increased rates of sea-level rise and species extinctions [*Waters* et al. 2016]. These persistent and pervasive markers of human activity led to the proposition of a new geological epoch: the *Anthropocene*, as a time where anthropogenic effects in the Earth system are clearly discernible from previous epochs.

An idea of the sheer scale of the anthropogenic changes to System Earth was given by *Zalasiewicz* et al. 2017, who estimated the extent of the so-called *technosphere*: the combined 'complex social structures together with the physical infrastructure and technological artefacts supporting energy, information, and material flows that enable the system to work'. They found the mass of the technosphere to be in excess of 10^{16} kg – in more comprehensible terms, this is equivalent to $50\,\mathrm{kg}/m^2$ of land surface; or five magnitudes larger than the combined human biomass. A recent study estimates that human-made mass may now exceed all living biomass on the planet [*Elhacham* et al. 2020].

I.2.8.3 Summary

Compared to all previous evolutionary development on Earth, humankind evolved from a relatively insignificant hominin species to a global force in the blink of an eye.

Presumably the most important aspect in their development was their capacity for universal language, giving rise to culture and mechanisms that allowed coordination within groups, leading to the formation of ever-larger societies. It is this capacity for cultural evolution that allowed the subsequent scientific and technological innovations.

Key observations from the emergence and rise of humans pertain not only to their capabilities compared to other species, but also to the effects that this rapid development had on the Earth system. The innovations and the spread across the globe were so rapid that – unlike in eusocial insects – no considerable coevolution with other parts of the Earth system occurred. Given the wide-ranging effects of humans on all aspects of the Earth system, including their own habitat, the future of humankind is unclear.

I.2.9 Summary

Concluding this glimpse into the Unfolding of Life, I want to very briefly list the most notable observations. The processes re-occurring throughout the evolutionary history of the Earth system are:

- changes in energy sources: from geochemical to light to flesh to fire
- modification of and coevolution with the environment
- group formation in two characteristically different ways: egalitarian & fraternal
- changes in information processing: DNA, coordination processes, language
- modularisation: entities combining into larger structures

All these will play an important role in the formulation of Evolution Mechanics.

I.3

Evolution Mechanics

This chapter presents the *Evolution Mechanics* (EM) framework. As noted initially, the foundations of this work are the thorough discussions with Kurt Roth, Benjamin Herdeanu, Harald Mack, and others over the last years – Evolution Mechanics came about in that environment. Here, I elaborate and extend on this foundation.

The core idea and claim of Evolution Mechanics is that the emergence of complex, hierarchically organised structures – like the Unfolding of Life, but not limited to it – can be described and understood as part of one conceptual framework. The Evolution Mechanics framework is built on three main abstractions:

- the *self-replicator* as the atomic particle of an evolutionary system;
- the *DAT processes*, a set of processes which lead to the emergence of a new self-replicator and hence: a new organisational level; and
- the *Evolutionary Spiral*, which describes the emergence of multiple such organisational levels by the aforementioned processes and locates those constructs on a spiral trajectory.

Taken together, Evolution Mechanics aims to be a guiding framework to better understand the autonomous Unfolding of Life: how complex, hierarchically organised structures may emerge from simple ones.

In the following, I will first specify some premises and define a basic vocabulary. I will then describe the three main constructs of the EM framework in more detail.

I.3.1 Premises & definitions

I.3.1.1 Premises

A premise for the emergence and maintenance of complex structures is that the system of interest is in non-equilibrium: Structure formation entails a reduction in entropy, hence any such system in which structure emerges either needs to be a relaxation from an excited state or have access to an external energy source.

The former scenario is applicable on the scope of the universe; the latter applies to systems like Earth, which is kept out of equilibrium by an external energy flux [*Kleidon* 2010].

I.3.1.2 Definitions

In the following, I define a basic vocabulary for the Evolution Mechanics framework.

Entity An abstracted structure that can in some way be discerned from its environment, e.g. in morphology, interactions, or capabilities. For instance, this may relate to a biological organism, but also to a viable core within an ACS, or a sufficiently strongly interacting group of entities.

Environment Evolution Mechanics postulates many organisational levels, hence the *environment* cannot be defined in a fixed way, but needs to adapt to the currently investigated level of organisation. In essence, it can only be defined in contrast to an *entity*. Specifically, the term does not refer solely to the abiotic surroundings of an entity, but may include other entities and the interactions with them.

Energy & resources Generalisations of the corresponding physical concepts.

Genotype The heritable information that allows to reconstruct an entity. In EM, this is specifically *not* limited to information stored in DNA, but to all forms of information storage.

Phenotype The realised structure of an entity, based on genotypic information but influenced by the state of and interactions with the environment. Selection pressure acts on the phenotype.

Genotype space & phenotype space The high-dimensional, non-euclidean spaces of possible genotype and phenotype states. Depending on the mechanism that translates a genotype into a phenotype, it *may* be possible to formulate a *genotype-phenotype map* to define their relation to each other. Such a map may allow for multiple genotypes to map to the same point in phenotype space. In general, these spaces have no unique definition; but it is helpful to envision their existence.

The genotype of an offspring entity may be different from that of the parent, leading to a different position in genotype space and possibly also in phenotype space. This 'movement' may be caused by *mutations* in the genotype introduced during the reproduction process. Alternatively, sexual reproduction leads to recombination of the parent genotypes, thereby introducing variation.

In general, the position in genotype and phenotype space is not fixed during the lifetime of an entity. The phenotype changes during development of the entity and is influenced by all its interactions. The genotype may also accumulate heritable information; for DNA-based life, this may be in the form of epigenetics.

State space If a distinction between genotype and phenotype is not necessary, the term *state space* is used to refer to the space of possible structures a system may exhibit.

Species A set of entities whose phenotypes are so similar that the competition between them is much weaker than with different entities. This definition includes biological species, but also includes more general concepts like those of quasi-species [*Takeuchi and Hogeweg* 2012]. Due to the difficulties of delineating this term, it is not widely used in EM.

Niche A niche describes the favourable conditions for an entity to persist in some environment. It can be understood as a hypervolume[15]in a high-dimensional niche space that contains all possible configurations of the environment, including other entities in it. Depending on the context, different subspaces of niche space may be highlighted, for instance the spatio-temporal location of the entity in physical space, their energy and resource demands, or their need for specific interactions.

Invention & innovation An *invention* denotes the discovery of some new process, structure, or energy source. An *innovation* is an invention that is particularly useful in that it gives a substantial benefit to the entities or leads to a rapid spread.

In effect, an invention can be thought of as giving access to an unpopulated volume in niche space; an innovation gives access to a particularly large volume in niche space and leads to many entities populating that new volume, making the invention wide-spread.

Autonomy Autonomy refers to an entity's capacity to perform a certain action in a predominantly self-controlled manner, i.e. without other entities needing to

---

[15] More precisely, favourable conditions may be described as an entity's utilisation function defined over the niche space. The *niche* then denotes a volume of niche space where utilisation rates are sufficiently large.

be involved. With entities never being fully independent, the delineation of this quality of an entity is soft.

Autotrophy From Greek *auto-* (self-) and *-troph* (-nutrition). The ability of an entity to robustly extract all energy and building blocks it requires for its formation and maintenance directly from components of the environment. Here, this term is used in the literal sense of the word and is specifically not limited to the context of biological life. Hence, an entity's environment and the resource and energy availability therein defines whether it can be regarded as autotrophic in that environment. For instance, a food-generated ACS would be regarded as autotrophic: the food set is part of its environment and, by definition, the food set provides all the building blocks for an ACS to construct itself.

To make this definition more robust to different environment compositions, we add the requirement that the entity shall require only fairly general components from its environment, which are available in sufficiently high concentrations. Note that the property to be autotrophic does not imply or require limitless growth; the definition is merely meant to allow statements regarding the generality of an entity's energy and resource needs.

I.3.2 Self-replication

At the core of Evolution Mechanics is the concept of *self-replication* and the corresponding atomic particle, the *self-replicator* (SR). The idea of self-replication is in many ways equivalent to that of *evolutionary units* in the Darwinian sense, but it focusses on the essential capabilities these units need to have and thus allows to abstract away from biological contexts.

In the following, I will first reiterate on the necessities for a system to be evolutionary in the classical sense and then show how self-replication fulfils these requirements while also being a more condensed definition.

I.3.2.1 Fundamental processes of evolving systems

Classically, for a population of entities to be considered an *evolving system* would require three fundamental processes taking place: (i) phenotypic variation in the population, (ii) differential survival and reproduction depending on phenotypic differences, typically called *fitness*, and (iii) heritability of fitness. This idea, famously, goes back to Charles *Darwin and* Alfred *Wallace* [1859; 1858], who described the above dynamics as 'descent with modification'.

A more condensed version of these necessities is formulated in terms of 'two kinds of variation': (i) variation in heritable traits, and (ii) variation in reproduction rates. Whichever formulation is chosen: As soon as a system exhibits these mechanisms, it *has* to be regarded an evolutionary system. This consequence is illustrated by simple replicator systems which show selection for fitter traits [*Takeuchi and Hogeweg* 2012].

The entities' capabilities are determined by their underlying structure and their interactions. For simple replicators, these are minimal – for life on Earth, they are highly diverse.

I.3.2.2 The self-replicator

Self-replicating entities fulfil all of the above requirements for being regarded an evolutionary unit:

- Self-replication trivially implies inheriting genotypic information from the replicating entity and with it: its phenotype and fitness; otherwise it could not be called *self*-replication.
- In a world with stochasticity, self-replication can never be perfect. A self-replication process will thus necessarily introduce some form of variation into a system, even if minute.
- As soon as there are phenotypic differences, subsequent self-reproduction and survival rates will be different; the actual rates depend on the environment and the interactions the offspring entity has with it.

These properties allow making *self-replicators* (SR) the atomic particles of the Evolution Mechanics framework. In contrast to the classical definition of evolutionary units, the formulation using self-replication is slightly more concise, and – more importantly – makes it easier to isolate the emergence of new organisational levels: As soon as a group of entities is capable to reproduce *as a whole*, they may be regarded an evolutionary unit. Hence, the focus is on the ability to self-replicate; it is the crucial property distinguishing a higher-level evolutionary unit from a group of strongly-interacting, yet *individually* replicating entities.

As discussed later, there are two ancillary concepts – the autotrophy and autonomy of SRs – which are important qualifiers for self-replicators. Before that, some clarifications regarding the definition of self-replicators as evolutionary units is needed.

On imperfect self-replication

The claim that self-replication can never be perfect and hence always introduces variation needs some further argumentation, specifically when it comes to very simple systems which would seemingly be able of perfect self-replication.

Consider the simplest autocatalytic chemical reaction, which leads to the self-reproduction of the catalyst – how could variation arise in such a setting? Written as $\mathrm{A} + \mathrm{C} \rightleftharpoons 2\mathrm{C} + \mathrm{W}$ (see Appendix A.1 for context), the reaction would indeed never allow for variation, because it is considered to be completely isolated. However, in practice, chemical reactions do not occur in isolation, but in an environment with non-zero thermal energy and a large number of other molecules. Depending on the reaction kinetics, it cannot be precluded that other reactions occur occasionally; and even if they are very unlikely or of short temporal stability, these fluctuations would be deviations from the typical reaction path and thus be classified as variations.

While the above scenario is highly constructed, it illustrates that it is difficult to assign a lower boundary to what counts as variation and what does not. Instead of attempting to set such a threshold value, Evolution Mechanics directly includes variation into the self-replication process with the assumption that every self-replication can and will give rise to variation. The motivation behind this is to put the conceptual focus not on the binary capacity for *any* form of variation, but on a self-replicator's ability to generate new structures – which is what is relevant for the Unfolding of Life.

Crucially, the emergence of new structures through self-replication does not depend solely on variation, but additionally on the space of *possible* structures that may arise from the self-replication. Essentially, the variation rate relates to the speed with which that space is explored; however, the structure of the self-replicator is what determines the extent and topology of structure space[16]. Focussing solely

---

[16] The term *structure space* is used here instead of genotype and phenotype spaces, because the distinction of the two is not relevant for this argument. Note that the structure of an entity also prescribes the interactions it has with its environment.

on the capacity for variation would neglect the important role of the SR's own structure in determining which complexity may evolve from it.

This perspective can be applied to the context of individual autocatalytic reactions: In EM, these systems would indeed qualify as self-replicators and hence also as evolutionary units. However, the structure space that becomes accessible by variations introduced by their self-replication is too small to expect considerably more complex forms. In contrast, the structures that can be attained by imperfect self-replication of an RNA molecule or multi-cellular organism are considerably more numerous, hence reflecting their larger potential to create more complex structures.

Summarising the above, the concept of the self-replicator in Evolution Mechanics unifies in it all necessities for an entity to be an evolutionary unit. While deliberately reducing the threshold for entities to be regarded as evolutionary units, it instead prompts for an assessment of an entity's capacity to form complex structures.

On the definition of life

In this context I want to reflect on the definition of life, also because it exemplifies how Evolution Mechanics can contribute another perspective on topics like this.

A definition of life was extracted from the scientific literature in a meta-analysis by *Trifonov* [2011]: 'life is self-reproduction with variations'. With EM *defining* self-replication to include variation, that statement would be further compressed to: 'life is self-replication' or 'any system capable of self-replication is alive'. The initial response to such a definition may be to question its usefulness and applicability: Would this not lead to systems like ACS with viable cores being classified as 'alive'? And if so, does such a definition fulfil any purpose?

These questions resemble the ones from the previous section of the threshold to be regarded an evolutionary unit – which is not surprising, given that they both arise from the amalgamation of variation and self-replication in EM. In fact, the argumentation is along exactly the same line: With EM focussing on describing the Unfolding of Life, it is more important to assess an entity's capacity to generate complex structures rather than to classify it into living and non-living. The former can inform the role a certain entity may have played in the evolutionary history of a system, while the latter can only do so if the implicit assumptions made during classification are sufficiently clear.

Definitions of life have historically shifted depending on the knowledge about the systems at hand, the context in which it was to be used, and the questions that were studied. Early definitions were of a primarily descriptive nature, assessing the capabilities of biological organisms (like their metabolism or their ability to reproduce). Modern definitions not only are more abstracted and include the concept of evolutionary units, but put the threshold for life to far simpler systems: In astrobiology, for instance, life is referred to as 'a self-sustaining chemical system capable of Darwinian evolution' [*Lammer* et al. 2009], which is essentially equivalent to what

Trifonov extracted, but with the added limitation to chemical systems. Going one step further, in the field of artificial life, the particular medium on which life may arise is not prescribed at all [*Taylor and Dorin* 2020].

Both these fields operate at the transition from non-living to living system, which may explain the shift in definitions towards more abstracted definitions. Despite the increased abstraction, these definitions still include variation as a key ingredient. However, as argued above, the capacity for variation cannot easily be regarded as a binary property of a system; to use it for categorisation, it requires to assume some form of threshold variation rate above which a system is classified as showing variation, which prompts for the question: At which point would a system be regarded to show *enough* variation to be considered alive? Without a naturally-given 'canonical' threshold, any such choice will be arbitrary.

Furthermore, as argued above, the structure of the system itself is of importance for the complexity that may arise. Assuming a chemical system already opens a large structure space and, as shown in the context of pre-biotic evolution (Section I.2.2.3), it may lead to similar precursors as on the early Earth. This is less so for the highly restricted structure space that can be represented in simulation-based artificial life studies.

In effect, binary definitions of life appear to require making explicit but presumably arbitrary choices exactly in those regimes where such a choice may have a large impact on the classification result[17].

In the context of Evolution Mechanics, it is more suitable to regard life as a continuum rather than a binary characteristic (dead or alive) and use it to describe the capacity of a system to form complex, hierarchically organised structures through evolution[18]. This potential cannot be assessed by looking at the self-replicator alone, but the whole system needs to be taken into account: the structure and abilities of the self-replicator, the DAT processes, and the environment. In a sense, the question is not whether a system can be regarded as living or non-living, but it needs to be characterised to which extent it is capable of evolving complex structures.

This perspective aligns well with the aim of EM to describe the evolutionary Unfolding of Life: the process of complex structures forming out of simpler ones.

---

[17] Naturally, definitions of intrinsically fuzzy concepts or characteristics – such as life – will always be pragmatic to some extent; even if they are in some ways arbitrary, they may be useful. To clarify: The partly deconstructive approach taken here is not to suggest that this would not be the case; specifically, it is not to say that existing definitions of life are not useful in their respective fields or that they are the sole characteristic that would be studied. Yet it would be an omission to not include a definition of life from the perspective of Evolution Mechanics and to contrast it to existing ones.

[18] The idea of a continuous description of life is not entirely new; a similar view is taken by *Krakauer* et al., albeit formulated with a focus on individuality.

I.3.2.3 Autotrophy and autonomy of self-replicators

Having established the basic definition of self-replicators in EM and contextualised the role of variation and structure space, two further aspects are of relevance that were only implicitly assumed above: the autotrophy and autonomy of a self-replicator. These properties qualify in how far a SR can be seen as independent in initiating and carrying out the self-reproduction process. If this is the case, the SR itself can control its self-replication through some internal mechanism, limited primarily by the speed of its replication processes.

In the context of a newly-emerging SR, these properties are especially relevant: An autotrophic and autonomous SR is less dependent on the specifics of its immediate environment can initiate its own self-reproduction, hence it can be expected to more rapidly populate an unpopulated part of niche space. In the following, both these properties are implied when speaking of a newly-emerging self-replicator.

Autotrophic SR Autotrophic self-replicators are of special interest because they have robust access to the energy and resources they need for their own sustenance and self-replication, which makes them less dependent on a specific configuration of the environment. In contrast, non-autotrophic self-replicators have to be regarded as being directly or indirectly controlled by the availability of resources and energy in their environment. This uses the abstracted notion of autotrophy (see definitions) which puts a focus on the generality and robustness of an entity's resource needs. In effect, it is not solely dependent on the capabilities of the entity but also on the environment it is embedded in.

For example, consider a specialist archaea cell which is strictly anaerobic – a scenario used in [*Martin and Russell* 2003] to model the origin of eukaryotes[19]. This SR would only be considered autotrophic if it is embedded in an environment with sufficient levels of CO_2 and H_2; otherwise, the availability of these substances controls the reproduction.

Subsequently, upon changes in the environment, organisms that evolve in a way to remain autotrophic will be more successful. Following the path towards the eukaryote envisioned in [*Martin and Russell* 2003], the archaea may evolve to increase its interaction with a bacterial cell, which provides the required substances as waste products of its own metabolism. By incorporating the bacterium, the former archaea becomes dependent on the metabolism of the bacterium; hence, the effective environment of the archaea changes and its autotrophy now has to be evaluated regarding the resource needs of the bacterium. The resulting structure of the archaea with an endosymbiotic bacterium (basically, a proto-eukaryote) would be more independent on CO_2 and H_2 and could instead feed on organic material.

---

[19] Note that the term autotrophy used there is that of the biological literature.

Autonomous SR Autonomous self-replicators are able to initiate and carry out the self-replication process largely independently from other entities and without requiring an external trigger to begin the self-reproduction. To that end, an autonomous SR needs to have all the relevant machinery for creating a copy of itself readily accessible – a 'complete toolkit'.

While autonomous self-replication should not require external control, this does not mean that these SRs may not take cues from their environment regarding the time they initiate self-replication. In fact, the ability to use additional information for the timing of self-replication will most likely be an evolutionary benefit. For autonomy, the focus is on the general ability to perform this process independently. In a way, this is similar to the definition of autotrophy, which also focusses on a general ability of an entity: that of depending only on sufficiently abundant components of their environment.

As with autotrophy, the point beyond which an SR is considered autonomous is soft and depends on the processes that may be externalised to the entity's environment. For instance, an isolated virus is not able to self-replicate; only if it is in an environment that contains the appropriate machinery for its reproduction – a host cell – will it become a self-replicator. The crucial point in the case of the virus is that it takes over the host cell (partially or completely), hence gaining control of cell functions including the reproductive machinery.

Furthermore, despite similarities, the term *autonomy* should not be confused with *agency*, which is used in sociological contexts and includes higher cognitive functions or decision making processes. Autonomy, in turn, may also apply to a viable core of an ACS, for which the machinery for self-reproduction is a direct consequence of their structure.

I.3.3 Diversification – Aggregation – Transformation

With the concept of the self-replicating entity established, we can turn to the emergence of organisational structures like modules and hierarchical levels, both of which are crucial aspects of the Unfolding of Life. In the EM framework, a hierarchical levels are defined by the emergence of a new self-replicator. This new self-replicator can be understood as a new module which comprises other entities (including other SRs) and which is capable of autonomous, autotrophic self-replication.

Note that *hierarchical levels* in EM are exclusively referring to the emergence of a new self-replicator. In contrast, *modules* are defined via their structure and need not have a certain set of capabilities. While recombination or nesting of modules technically can be described as a hierarchy, this use of the term is avoided to reduce ambiguities.

How do new self-replicators come about in EM? We subdivide the processes leading to the emergence of a new hierarchical level into three classes: (i) *diversification*, (ii) *aggregation*, and (iii) *transformation*. This set of processes is referred to as *DAT processes* or DATs and culminates in a transition to a new self-replicator, an *evolutionary transition in individuality* (ETI, here subsumed under transformation).

The DAT processes define the minimal set of processes that need to occur such that a new self-replicator can emerge: through *diversification*, SRs change in their structure and establish new ways of retrieving energy from their environment, potentially changing it in the process; through *aggregation*, diversified SRs begin interacting more strongly, forming loose synergistic relationships or more tightly-coupled encapsulated groups; and finally, as part of the *transformation*, the group internally reorganises up to a point where it is able to replicate as a whole, this last step marking the *transition* to a qualitatively new self-replicator.

The DATs do not necessarily have to occur in exactly the given order for a new self-replicator to emerge. For instance, SRs might have been encapsulated by chance and only begin forming cooperative interactions after their encapsulation. Depending on the conditions, the processes may be more or less pronounced; furthermore, they may overlap in time, sometimes to the end that they occur essentially in parallel, and they may occur at different speeds. This speed is affected by the likelihood that a certain process occurs, also depending on the conditions at the time.

To arrive at a concise conceptualisation, we divided and ordered these processes in a way that best matches the observations from the Unfolding of Life. We argue that it is indeed *all* of these processes that are relevant and need be considered for the transition of one self-replicator to a higher-level self-replicator.

In the following, I will go into more detail about each of the process classes and how they may result in the formation of a new self-replicator.

To improve readability, many statements below are written in active language, despite evolutionary processes typically giving the unit of selection no room for an active role in their own evolution. For instance, statements like 'the SR evolves to access a new niche' should always be understood as 'there is selection pressure that tends to give mutated SRs with access to a new niche an evolutionary benefit over other SRs that do not have such a mutation'. Also note that the perspective taken here is in the form of a projection, going forward in time; this is not to suggest that such a projection would be possible at each point, but is again done to simplify the narration. In effect, we are tracing the path of an entity that evolved and changed as a result of all the DATs, while putting aside the many other entities that did not, including those that branched-off from the selected entity.

I.3.3.1 Diversification

Diversification is the process by which an evolutionary system explores its state space, resulting in qualitatively new structures – these are important precursors for the emergence of a new self-replicator. For instance, in a biological context, diversification would denote the emergence of new species. This exploration of state space pertains not only to the self-replicators but also to their environment, influencing each other and coevolving.

Furthermore, diversification is typically accompanied by – or even: driven by – modifications to niche space, leading to improved access to energy and resources. This may be caused by aforementioned interactions with the environment, which strongly influences niche space, but may also be the result of particular innovations.

However, improvements in energy access need not induce diversification in state space but may actually lead to the opposite. Hence, these domains need to be carefully distinguished in EM, which is why *diversification* refers exclusively to diversification in state space, i.e. a SR's genotype and phenotype space. The changes in niche space are addressed separately when considering the cause or effect of diversification.

Diversification is by no means a monotonic process, nor is it a simple random walk in state space. Instead, it includes a wide range of dynamics: times of near-stasis, slow but continuous adaptation, or very sudden changes[20]. As a background process of diversification, we can assume the base error rate during self-replication. Such a base rate may be caused by the inaccuracy of the self-replication process itself or by external factors; for instance, time periods with stronger cosmic radiation may lead to more mutations in DNA-based life forms. In addition, innovations like that of sexual reproduction can introduce variation via recombination, thus leading to faster base rates.

---

[20] Debates regarding the exact nature of these dynamics are plentiful. Most prominently perhaps: whether *punctuated equilibrium* or *gradualism* best describes evolutionary innovations. While interesting, the distinction is not particularly relevant in EM.

There are processes that counteract diversification, for instance by causing SRs to become extinct and thus removing structural information from the state space of the system. Similarly to diversification, base rates for extinctions can be assumed, which may fluctuate depending on many factors.

In addition, two phases can be distinguished where the characteristics of diversification and extinction depends on the state and topology of niche space: During *radiation*, a large volume of niche space becomes suddenly accessible, leading to SRs rapidly populating those niches and diversifying in the process. In contrast, *extinction* phases denote time periods where niches become suddenly uninhabitable, reducing diversity in the short-term, but shaping the long-term evolution of both niche space and possible structures.

For diversification, different evolutionary driving forces can be discerned, which are more or less dominant in the respective phases. To illustrate these forces, consider an unpopulated volume of niche space that becomes accessible:

First, a SR that has acquired access to that unpopulated volume of niche space (e.g. by some new capability, an *invention*) can quickly populate those niches without feeling particularly strong competition from other SRs (letting the invention become an *innovation*). During such a phase, SRs diversify primarily via *neutral drift*, meaning that selection pressure is either weak or undirected – this would indeed constitute something akin to a random walk in genotype space, leading to diversification in phenotypes.

Second, once this newly-opened volume in niche space becomes more populated, competition between SRs increases, subjecting them to selection pressure that are no longer neutral: The *competitive exclusion principle* (CEP) becomes dominant, by which prolonged coexistence of SRs is only possible if they are sufficiently distinct in niche space. Subsequently, there is a benefit in also occupying the less accessible niches (specialisation), further facilitating phenotypic diversification.

Third, diversification may occur through SRs sharing innovations in a process called *gene flow* or *horizontal gene transfer* (HGT[21]). This process is only indirectly dependent on the occupancy of niche space, but SRs that gain innovations via HGT may have a competitive advantage in a scenario with strong competition between similar SRs.

Finally, the environment plays an important role in this context: The environment is not only the medium on which SRs diversify, but the diversification also causes environmental changes, thus leading to alterations of niche space – in turn forming selection pressures, which then influence diversification. This may either be a self-enforcing feedback, leading to further niches becoming available, or a self-inhibiting one. Due to this coupling and delayed self-modification, the diversification of SRs cannot be isolated from its coevolution with the environment.

---

[21] Note that both terms use the abstracted meaning of *gene* as any form of heritable information, which need not be DNA-based; see Section I.2.2.4 for more information.

In the following, I will first take a closer look at the coevolution with the environment before turning to the aforementioned radiation and extinction phases.

Environment coevolution

The environment plays a crucial role during diversification as it prescribes the limits and possibilities of evolution while, at the same time, being constantly transformed by the SRs and their changing interactions with it. In essence, diversification occurs not only in the genotype and phenotype spaces of the SRs, but also in the structure of the environment. Hence, these two aspects cannot be studied in isolation, but are intrinsically linked as constituents of the same self-modifying system – entities and environment are coevolving.

This may manifest in various ways and to different extents. In the most abstract description, occupying a previously unoccupied niche constitutes a new interaction taking place between entities and their environment, which has to be understood as a change to the environment. For instance, consider a population of birds reaching an island devoid of any other birds; by feeding alone, the presence of these birds will lead to selection pressures that have not been there before.

While qualitatively the same, cases where the environment modifications are large or lead to significant changes in selection pressures are of special interest. These modifications are referred to as *niche construction* (if they pertain mostly to the benefit of the entity that causes them) or *ecosystem engineering* (if they affect the success of other species) [*Erwin* 2008]. The classic example of niche construction are beavers: their dams cause the formation of lakes, transforming the local ecosystem and having a large effect on nutrient cycling and other species. Another example is that of the cyanobacteria of the Great Oxygenation Event (see Section I.2.3), which lead to persistent planetary-scale changes.

In the section on the potential capabilities of self-replicators (Section I.3.2), the capacity for formation of complex structures was ascribed to the structure space of the self-replicator alone. With the considerations made here, this statement needs to be amended to include the structure of the environment and the interaction between the two.

Radiation

In the EM framework, *radiation* denotes a rapid diversification of SRs that is caused by a new volume in niche space becoming available[22]. There are several ways how unpopulated niches may become accessible, for instance: a mutation causing an

---

[22] Note that, in evolutionary biology, this process is referred to as *adaptive radiation*: the rapid diversification of a subset of species, e.g. after a new environmental niche has opened or after a key trait (an invention) was acquired. As EM uses a wider formulation of niche (which may include geographic niches, for instance), further differentiation of different kinds of radiations are not required here.

innovation and hence allowing to populate a niche; mobile entities becoming able to enter a habitat they did not have access to before; new niches developing due to changes in the environment; or a competing SR becoming displaced and thus freeing up a previously populated niche. Naturally, the niches that become available by any of these processes may differ considerably in size, depending on the specific innovation and feedback loops with the environment.

An additional scenario by which a large niche space may become available is that of a newly-emerging self-replicator: Consider an entity which acquired some invention throughout its evolutionary history; further, assume that the invention has the potential to give access to many niches, but that this potential could not be fully utilised due to a limited self-reproduction rate. The first manifestation of a higher-level self-replicator will be an entity that can autonomously and autotrophically reproduce, thereby reducing limitations on the self-reproduction rate. Acquiring the ability to self-replicate without external control and with fairly general resource requirements can be a large benefit over other SRs, thus leading to high reproduction rates and diversification by the above-mentioned processes.

The Cambrian explosion (see Section I.2.6) is an example for a particularly large radiation event. It is unclear, whether a sole trigger or underlying cause can be found for the Cambrian explosion; the above processes certainly all played a role to some extent. However, in the context of EM, the following abstracted description is sufficient: The Cambrian explosion made a very large volume in niche space accessible, self-enforced by environment coevolution, and lead to a high diversification of SRs.

Extinction

As an opposite to *radiation* phases, *extinction* phases denote a time of rapid reduction in the diversity of SRs. Diversity loss is not necessarily due to a reduction of accessible niche space, but is often caused by it (cf. Section I.2.6). Such a niche loss can be caused directly by sudden changes in the environment, or indirectly by other SRs gradually modifying the environment. Again, changes in one part of the environment are often coupled to other parts, allowing for self-amplifying extinction events. Hence, extinction phases can be are conceptually similar to radiation phases.

What is the role of extinction events in the overall diversification of species, i.e. the exploration of state space? Their direct effect certainly is a diversity loss, but one that does not affect all SRs equally: those that are more robust towards the environmental changes and fluctuations in their niches are more likely to persist while those that have specialised on a very specific niche are more prone to become extinct. This can be understood as an additional, temporary selection pressure applied on the combined phenotype space of all SRs. In effect, the losses of information in the combined genotype space are not distributed uniformly but are shaped by this additional pressure.

Now consider the situation after the strong fluctuations in niche space abated and the extinction rate has normalised[23]. The niche space may then either slowly recover to its previous state (with the effect of the fluctuation wearing off) or be in a permanently transformed state. Regardless of which state it is in, it represents a situation where diversity is reduced but occupation of niche space is also reduced; specifically, niches that required strong specialisation will often be unpopulated.

In this situation, SRs evolving to re-populate such niches can form different structures than the SRs that populated the niche prior to their extinction, thus exploring parts of phenotype space that have not previously existed. However, depending on the specific phenotypic features needed to populate such a niche, it may also lead to the remaining species evolving to show analogous structures, a process called *convergent evolution*[24].

One example for a very large extinction event is the P-T extinction where it is estimated that 90 % of species became extinct (see Section I.2.6). An example for the transformative nature of an extinction event is the GOE (see Section I.2.3), where the innovation and subsequent radiation of one family of SRs (oxygenic cyanobacteria) led to the demise of another family of SRs (those for which oxygen was toxic); this is presumed to have played an important role in the origin of eukaryotes (see Section I.2.4).

I.3.3.2 Aggregation

By the process of *aggregation*, a subset of a diversified system of SRs begins to cooperate and eventually form an encapsulated *group* of SRs. Like diversification, aggregation can bring SRs an evolutionary benefit by opening up new niches or improving existing ones.

The EM framework highlights two key mechanisms as the substructure of the aggregation process: *cooperation* between SRs can bring mutual benefits to the interaction partners, thus initiating group formation. With strong-enough interdependencies between them, *encapsulation* strengthens the newly-formed group by facilitating in-group interactions and inhibiting disadvantageous out-group interactions.

As described in Section I.2.5.3, two modes of group formation may be distinguished: those of identical parts (fraternal) or those of differing parts (egalitarian). Depending on which kinds of SRs come together, the aggregation phase may have different characteristics.

---

[23] Separating these periods is of course a simplification; in practice, they will strongly overlap.

[24] Convergent evolution is not limited to the context of extinctions but generally applies to all situations where rather specific structures are beneficial when populating a niche. In the case of extinctions, it can be understood as essentially the same niche, separated in time. In other cases, there may be multiple niches for which the benefits apply (e.g. geographically separated).

Cooperation

Cooperation denotes a form of interaction between a set of SRs that leads to an overall evolutionary advantage. This benefit need not pertain to all involved SRs as long as the emerging construct – the *group* or *collective* – has some overall advantage over competing SRs; this may indirectly lead to benefits for the SRs, despite individual disadvantages. It can thus be seen as a form of self-organisation by which SRs act towards a common benefit.

In EM, cooperation marks the beginning of group formation and hence of the aggregation phase towards a higher-level SR. This aggregation can be observed via systematically repeated interactions between certain kinds of SRs, thus allowing to associate them as a (loose) collective. If their interactions bring an evolutionary advantage in the competition for niches, it is likely to become further pronounced. As such, cooperation introduces a new level on which selection may act: the group.

There are many mechanisms by which the advantages on the group-level are conveyed back to its constituents[25]. In the simplest case, the cooperative interaction may lead to a direct, mutualistic benefit for all involved individuals. In other scenarios, the benefit may be delayed; this makes the dynamics of the interaction more complicated, because the cooperative interaction can be assumed to incur a short-term cost and thus constitutes a disadvantage for one or all interaction partners. However, even if the benefits arise only indirectly[26], cooperation may still constitute an overall evolutionary advantage.

Examples of cooperation in nature are plentiful and can lead to the formation of groups of dissimilar (egalitarian) or similar (fraternal) entities. As an example of the latter, cooperation of individuals in eusocial societies gives rise to the collective having more capabilities than the individual, as seen for eusocial insects in Section I.2.7. Examples of cooperation between different kinds of SRs, there is the symbiosis between fungi and cyanobacteria or algae, giving rise to lichens [*Lenton and* A. *Watson* 2011]; or the endosymbiosis of proto-mitochondria and proto-eukaryotes discussed in Section I.2.4.

These two configurations of fraternal and egalitarian group formation through cooperation differ in the potential benefits and costs of the initial cooperative interaction. An initial advantage can be crucial, as it can lead to a self-enforcing situation, akin to autocatalysis.

In fraternal groups, *Bourke* [2009] sees the large number of interacting entities as a potential for why a cooperative interaction may become advantageous: Even small individual costs may lead to emergent behaviour on the scale of the group. This is counteracted by the intrinsic competition of these entities for the same niche, which brings an advantage to those individuals that do not cooperate, but

---

[25] The terms *individual* and *constituent* are used synonymously to refer to a part of a collective.

[26] The exact mechanism is not of further relevance in the context of EM, but is topic of an extraordinary range of research – keywords are the *prisoner's dilemma* and the field of *evolutionary game theory*, or the debate surrounding *kin selection* and *multi-level selection.*

still benefit from the group. The relation between the group-level benefit and the level of selection is decisive in determining whether cooperation may manifest itself between the SRs.

Egalitarian groups, in contrast, are typically small in number and may be comprised of as few as two entities. However, these entities are structurally and functionally diversified and thus depend on different niches, alleviating the competition between them. Any change in interaction between the SRs that benefits at least one of them (and does not disadvantage the others) is thus beneficial; a fundamentally simpler situation for cooperation to arise than in the fraternal case.

Two adjacent forms of repeated, long-term interaction between SRs should be briefly mentioned, which are distinct from cooperation: First, in the case of *parasitism*, host and parasite could be regarded as forming a group. This is not considered cooperation because the host has a clear disadvantage from the interaction with the parasite. It may however *turn* into a cooperative relationship, for instance through some mutation or change in the environment that makes the parasite-host interaction become commensalistic or mutualistic. From that point on, it would be classified as cooperation.

Second, trophic interactions – predator-prey interactions – are typically not regarded as cooperative. However, despite the obviously negative effects on individual prey, an 'overall benefit' *could* be conceived: the prey providing resources for the predator, the predator preventing overpopulation of the prey. In the widest sense, this could indeed be understood as cooperation, if prey overpopulation would be a considerable threat to the survival of that SR and if there are no other interactions in place that would prevent it. However, trophic interactions have substantially more dominant other dynamics, which warrants neglecting this minute cooperative aspect they may have.

In summary, cooperation denotes an increase in interaction between SRs that leads to an overall evolutionary benefit for the resulting fraternal or egalitarian collective. In the context of the DAT processes, it is associated with the initial aggregation towards a higher-level self-replicator. Notably, this initial aggregation may be a weak form of cooperation; it may even be facultative, simply an additional capability of the SRs that is beneficial but not crucial for their survival. Cooperation becomes obligate only through the following stages of aggregation and transformation.

Encapsulation

Encapsulation describes the evolution of a physical or process-based barrier that encloses the components of a group, thus forming a clearer boundary between the 'inside' and 'outside' of the group. A precursor to encapsulation can be the evolution of a certain kind of cooperative interaction: an improved coordination between the cooperating SRs. Coordinating their life cycles may not only lead to more efficient cooperation, but would be a first step towards a collective action that may give rise to the formation of an outside boundary.

Through this process of encapsulation, a collective of cooperating entities may facilitate their in-group interactions, for instance by increasing local resource concentrations or reducing diffusive losses. Additionally, adverse interactions like free-riding or predation on individual components may be suppressed. As a result, the level of selection may shift further towards the group.

For example, consider the origin of eukaryotes (see Section I.2.4) where encapsulation occurred through endosymbiosis: The existing membrane of the host cell evolved to enclose the symbiont. The membrane is a permeable structure that actively or passively transports nutrients into the intra-cellular space, thus providing the individual constituents with the resources required for their functioning. Simultaneously, it reduces loss of resources between the symbiont and the host and acts as a protection against adverse influences.

Non-physical boundaries can be found in colonies of eusocial insects. Here, encapsulation can be seen as a behavioural capability, specifically that individuals from one colony will typically not transfer into another colony (see Section I.2.7). In addition, colonies often construct nests which (among other purposes) serve as physical protection or as a breeding space.

An intriguing example is that of the *origin* of cells because it illustrates differences between a (spatially confined) environment and the process of encapsulation. The first cells are presumed to have evolved at deep-sea hydrothermal vents, as described in Section I.2.2. The porous structure of these vents is thought to have facilitated the evolution of the biomolecular machinery. At a later stage, evolution of a lipid bi-layer membrane substituted the functionality of the pore structures and allowed for the cells to become mobile.

In EM, the pores would be associated with a confined environment and the evolution of the membrane would be regarded as the encapsulation process. The difference between the two is that the membrane is an evolved structure which is maintained by the group and brings an evolutionary advantage to the group. In contrast, the pores are a previously existing structure which provides favourable conditions for the constituents and requires no maintenance; they are hence closer to the definition of the environment.

However, the distinction between the two is not easily made. Picking up the example of the origin of eukaryotes again, the membrane of the host cell could be seen as a pre-existing structure which is maintained primarily by the host. In a way, what was referred to as encapsulation above could be seen as the symbiont moving to a more favourable, confined environment. If seen from the perspective of the group, however, the constituents evolve to utilise a previously existing structure for the benefit of the whole group. Whichever perspective is chosen, both are essentially appropriate pictures in this case, as they effectively provide the same functionality to the group of cooperating SRs and lead to an evolutionary benefit.

The above examples are all from groups of SRs that went through further processes on their way to a higher-level self-replicator; they are retrospective observations of encapsulation and, accordingly, are interwoven with other evolved functions of the new SR.

The situation is more difficult to conceive if the group is solely a set of cooperating entities and has not (or not yet) evolved to a higher-level SR, where it could coordinate the synthesis of the membrane with the reproduction of its constituents. Without these coordination capabilities, encapsulation would require the constituents to spontaneously aggregate and then synthesise a membrane that encloses them. Furthermore, due to a lack of coordinated reproduction, this process would need to be repeated in each life cycle of the group's constituents. Overall, early versions of encapsulation processes would appear to be extremely expensive and error-prone and thus be detrimental to the success of the group.

To a certain extent, this perception is a result of the conceptualisation made in the context of the DAT processes, aiming to differentiate processes that – in reality – occurred in a coevolving fashion and cannot be isolated. Despite this classification, the DAT processes should be understood as potentially occurring simultaneously; thus, there is no requirement for encapsulation to evolve all-at-once *after* the constituents reached a certain degree of cooperation.

Nevertheless, overcoming initial disadvantages appears to be one of the challenges for encapsulation to take place. As also seen above, there are conditions that may facilitate an encapsulation: For prokaryotic cells, it was presumably assisted by the structure of their environment; for eukaryotes, existing structures were repurposed to include further entities of the group. Furthermore, having evolved capabilities to coordinate group behaviour becomes beneficial for encapsulation: Such mechanisms may provide a direct benefit to the group and hence reduce initial disadvantages; having the general mechanism in place could then reduce evolutionary hurdles of evolving the encapsulation. While assisting the evolution of encapsulation itself, evolving group-level coordination would also allow other collective action – which will become increasingly relevant the closer the constituents of a group interact.

I.3.3.3 Transformation

Transformation is the process by which aggregated SRs become capable of self-replicating as a whole, thereby evolving into a higher-level self-replicator. The point at which the transformation to the new self-replicator is completed is denoted as an *Evolutionary Transition in Individuality* (ETI): the point at which the collective of SRs has acquired the complete toolkit needed for autotrophic, autonomous self-replication. As a result, the evolutionary system contains a qualitatively new form of SR. While this moves the level of selection towards the resulting new SR, the previously existing lower-level SRs continue to play a role – they are the foundation of everything above and continue to evolve under the changed selection pressures. In many cases, the lower levels can be considered to be reasonably stable, such that they may be approximated to play a negligible role for the further evolution of the new SR.

Two processes that precede such a transition can be isolated: Through *specialisation*, the constituents of a group diversify their role within the collective, thereby broadening the groups' abilities and refining in-group interactions. This is of particular importance for fraternal groups, which may lack in functional diversity.

Further, in what we call *unification*, the constituents may successively transfer capabilities from the level of the individual to that of the group. This can make the group more efficient, for instance by constituents shedding abilities that are unnecessary in the context of the group. The unification process also includes the acquisition of the machinery necessary for group-level reproduction. This machinery typically requires an additional channel for storing and transmitting information between constituents of the group. Both these processes result in the interdependencies between members of the group becoming more pronounced. Notably, this may include that the constituents lose the ability to reproduce *without* the collective.

The result of specialisation and unification is the aforementioned *Evolutionary Transition in Individuality* (ETI) and concludes the transformation process: A higher-level SR has emerged, being made up of strongly integrated subunits which act and reproduce as a whole. The transition also concludes the DAT processes, yielding a SR that is a qualitatively new entity in the evolutionary system – at the same time, it is conceptually equivalent to the SR we started out with, but with one further level of organisational structure beneath it. This conceptual equivalence gives rise to the idea of the Evolutionary Spiral (presented in Section I.3.4).

Specialisation

During *specialisation*, the entities and interactions within a collective adapt such that the group has an evolutionary advantage. This is typically achieved by a diversification of the function that each entity has in the group and that it performs for the success of the group, resulting in a group where the constituents have specialised roles.

While this process has similarities to the general diversification process discussed before, there are two important differences: First, the focus is on the capabilities and interactions relevant to the collective, rather than that of the constituents on their own. Second, the term specialisation suggests that a certain direction of diversification is of a general benefit, namely that which gives the whole group an advantage. Making this distinction between general diversification and functional diversification within the collective – here coined *specialisation*[27] – is useful as it

---

[27] Note that the term *specialisation* is sometimes used (in the literature, but also in this thesis) to refer to an entity evolving to specialise on a narrower niche. While the meaning has some overlap (interpreting the narrow niche as the specialised configuration of the group), the term carries more information in the context of the DAT processes. Subsequently, a more original term would be desirable.

highlights that the selection pressure driving this in-group reorganisation is not acting solely on the constituents of the group, but on the newly-formed collective.

The specialisation processes for groups of similar entities (fraternal) or groups of dissimilar entities (egalitarian, cf. Section I.2.5.3) have fundamentally different starting positions:

Egalitarian groups are, by definition, already structurally diversified and thus (most likely) already fulfil different functions in the context of the group. For instance, the eukaryotic cell is thought to have been formed by the ever-intensifying interaction and endosymbiosis of archaea and bacteria, each with their own set of distinct capabilities (see Section I.2.4). Subsequently, it can be said that the collective always had a high degree of functional diversity, because this is what made their cooperation beneficial in the first place. In effect, no further functional diversification may be necessary for collectives formed on the egalitarian path. They may of course still further evolve to improve the efficiency of their interaction (and *specialisation* would still denote any changes in functional diversity after the group was established), but the functions they fulfilled in the group were large even before they became encapsulated.

The situation is more complicated in fraternal collectives (cells in multicellular organisms, individuals in colonies of eusocial insects, see Sections I.2.5 and I.2.7). In contrast to egalitarian groups, fraternal groups consist of structurally similar entities – hence, their initial level of functional diversity depends on the degree to which the constituents may perform diverse roles *despite* structural similarity.

In these groups, *division of labour* is the process by which the entities of a group may fulfil a wider range of functions than the individual entities would be able to: The entities dynamically take up different roles in the functioning of the collective, hence allowing for emergent behaviour. In EM, this term is used in a wider sense than in the context of eusocial organisms: It refers to a functional diversification *during the lifetime* of a specific manifestation of the group, regardless of whether they are of a purely behavioural nature or include structural changes to the constituents of the collective[28]. This definition captures both the scenario of a multicellular organism, where stem cells diversify into different morphologies, and that of a eusocial insect colony, where individual insects differentiate into castes.

The ability to coordinate and carry out such a differentiation has to evolve at some point. If such an invention is only beneficial in the context of the collective, it is unlikely that it would have evolved prior to the existence of the collective: Without an advantage, the capability is not selected for; with a disadvantage, it may even be selected against. Naturally, the evolution of this capability can hardly be put into separate stages; especially because it may already play a role when entities begin cooperating and only subsequently begin to form a group. The

---

[28] As such, it is similar to *phenotypic plasticity* – an entity's phenotype being influenced by specific aspects of the environment – but with 'the environment' referring only to the collective the respective entity is part of, not the entirety of the system.

extent of this ability for dynamic functional diversification can thus be assumed to coevolve with the formation of the group. In EM, the *latest* point at which it would evolve is stipulated to be during the *specialisation* phase.

Summarising the above, *specialisation* denotes a phase of the transformation to a higher-level SR during which the constituents of the group optimise their functional roles in the collective. As a result, the evolutionary success of the group is improved. At the end of specialisation, both egalitarian and fraternal collectives can be assumed to be in a comparable configuration with respect to their roles within the collective, no longer requiring a strong distinction between the two paths.

Unification

During the *unification* phase, the collective's constituents transfer essential capabilities to the group, increasing their dependency on the collective and improving the overall functioning of the collective. One fundamental capability of the group's constituents – their autonomous self-replication – is singled-out at this point, because it is of special importance: Once this capability is transferred to the group and the group becomes able to self-replicate *as a whole*, a *transition in individuality* has occurred; it is thus handled separately, see below.

This aspect being excluded, in which ways are unification and specialisation distinct? While both improve the functioning of the group, they do so in different ways: Through specialisation, the interaction that gave rise to the groups cooperation is optimised in the context of the group. In contrast, unification is an optimisation that pertains to functionality of the constituents that is not relevant for the group, but that was relevant for the survival of the entity outside of the collective.

In principle, any capability of a constituent could be substituted by the group. If the group provides a sufficiently stable substitute for that capability, this may lead to the constituent losing that ability altogether, because maintaining it would be at a disadvantage to the group.

For instance, defensive measures of the subunits may no longer be necessary when in the collective. These could either be transferred to the group or lost altogether, if the group no longer requires these defences. As another example, the waste products of one constituent may include the resources needed by another – which might have been the mutual benefit that brought them into cooperation in the first place. Once robustly in the collective, the entities could specialise their respective waste removal and resource retrieval mechanisms to the specific situation of the collective.

These examples also show that the distinction between the unification process and the transition in individuality can be somewhat arbitrary: If a constituent loses its defensive measures or its ability to efficiently extract resources, it would most probably no longer be viable outside the collective; hence, pragmatically, its individuality is already affected by the described changes. These difficulties are inherent to aiming to describe overlapping quasi-continuous processes, which additionally have fuzzy boundaries – a topic that I will briefly address in the discussion, Chapter I.4.

Evolutionary Transition in Individuality

An *Evolutionary Transition in Individuality* (ETI) describes the process of a collective becoming capable of autonomous self-reproduction. It is the culmination of the *transformation* and creates a self-replicator with a new level of organisation, and hence: a new level of individuality. With self-replication at the centre of Evolution Mechanics, the process by which this capability evolves plays a correspondingly important role, which is why it is handled separated from the more general unification process discussed before.

Throughout the transition, reproductive control is transferred from the individuals to the collective such that the self-reproduction of the collective is also controlled *by* the collective; this is denoted by *autonomous* self-replication. Therefore, the changes that need to take place include the constituents' loss of their (individual) reproductive autonomy. This does not mean that the reproductive machinery of the constituents is lost or that they would necessarily be unable to self-reproduce outside the collective, but only that they – as long as they are part of the collective – are not controlling their own self-reproduction.

To illustrate this, consider a particular collective in an evolutionary state prior to the transition, i.e. a collective that is not yet able to autonomously self-reproduce. In such a case, the individual constituents are not only still in control of their self-reproduction, but they still *need* to self-reproduce. To generate an offspring of the group, all constituents need to first self-reproduce and then again aggregate into a new group, potentially synthesising a physical barrier around them – and this process needs to occur reliably-enough for each reproduction of the collective. Such a reproduction path is most likely to not be robust enough, thus entailing high costs if the reproduction fails to aggregate into a functioning group (group-centred adaptations are typically not as useful outside the group).

The key to making reproduction of the collective more efficient is for the constituents to coordinate their reproduction such that the group is in control. This coordination will be evolutionarily advantageous if the accompanying costs for coordination are lower than those of unreliable aggregation after individual reproduction.

What is referred to as *reproductive control* is exactly that ability of the group to coordinate its own reproduction. Hence, for a transition in individuality, these coordination mechanisms need to evolve and they need to become dominant over those of the constituents. General coordination capabilities may of course already have evolved prior to the transition, as they are also useful to facilitate cooperation or sustain encapsulation. The transition refers to the evolution of specifically those coordination abilities that are needed for robust reproduction as a collective[29].

Examples of ETIs throughout the Unfolding of Life Having described the process of ETI, which examples are there of it? From the observations of the Unfolding of Life given in the previous chapter, successive ETIs resulted in the following kinds of higher-level SRs, with proposed examples included in parentheses:

- prebiotic self-replicators (ACS?, RNA world?)
- simple cells (archaea, bacteria)
- complex cells (eukaryotes, plastids)
- multicellular organisms (plants, fungi, animals, . . .)
- super-organisms (eusocial animals, human societies?)

This list is arguably a very rough representation of the evolutionary history of life on Earth – and may offer many starting points for discussions, especially when it comes to the examples given and concrete associations with species that first fulfilled these criteria. In the context of EM, these associations are not in focus, but the qualitative changes in the organisational structure of the respective groups of self-replicators is, namely: the emergence of an overarching structure that coordinates collective reproduction. I will pick up some discussion points regarding this proposed list in Section I.4.2.

Furthermore, as hinted at by the examples in parentheses, the ETIs manifests not as a linear chain, but as a branching tree, where transitions may occur multiple times. This is easier to discuss in a larger context and without the sole focus on the transition process; I will hence address this in more depth using the concept of the *Evolutionary Spiral* introduced in Section I.3.4.

The abstraction of the transition process that is presented here is a simplification that focusses on the transition from one organisational layer to the next-higher one. However, the SRs can never completely decouple from all lower organisational levels; these still exist and remain relevant: Throughout the aggregation and transformation phases, selection pressure shifts to the group level, which may cause adaptations in those underlying structures. In addition, for a transition to occur, adaptations in lower levels may even be a requirement; this depends on the independence of higher hierarchical levels from lower ones and the restrictions that the capabilities of lower levels put on those above them.

This also puts the reproductive autonomy into context: a higher-level SR will never be fully independent from all of its substructure. For EM and specifically here

---

[29] As such, the transition should be seen as a *process*. However, note that the term is also used to refer to the *point* at which this ability was acquired, as this marks the emergence of the new SR.

in the transition process, only the next-lower level is considered, which is merely an approximation of the multi-level nature of the Unfolding of Life.

Reproduction paths What kind of coordination mechanisms need to evolve to allow for the collective to self-reproduce? The reproduction mechanisms for the SRs listed above – which are already *past* a transition – may indicate which processes might be relevant in their self-reproduction. One thing they have in common is a constricted *reproduction path*: All the information that is required to produce an offspring SR is bundled such that it may be efficiently transferred and subsequently used for construction of the offspring.

For instance, in colonies of eusocial insects, the relevant self-reproduction would be that of the whole colony, which occurs via the queen: In essence, the queen contains all information to initiate a new colony (which can be considered an offspring of the original colony).

Similarly, in many (eukaryotic) multicellular organisms, reproduction happens via sexual reproduction. The constricted reproduction path is hence through the gamete cells, which – again – contain all information to construct a new multicellular organism. As described in Section I.2.5.1, this germ line sequestration additionally alleviates many genetic conflicts within the collective.

Single-celled eukaryotes typically reproduce asexually through mitosis. Mitosis is a highly-coordinated multi-step process in which the cell grows, replicates its chromosomes, extracts the replicated chromosomes from the nucleus, separates them spatially into two sister sets, and finally causes the membrane to contract and create two separate cells. For prokaryotes, the situation is a bit simpler: they reproduce through binary fission, where only the single DNA molecule needs to be replicated; after spatial separation, a cell membrane grows and splits the cell into two parts. In both these cases, information is transferred via DNA and potential auxiliary structures[30] – these channels are arguably less constricted than in the more complex cases sketched above.

To address all SRs listed above, consider a prebiotic self-replicator: What would be the constricted reproduction path in the context of prebiotic chemical reaction networks? As described in Section I.2.2.3, for a viable core to reproduce in a new environment a single molecule from the viable core suffices. If such a core molecule is transferred to a new environment where it did not exist previously, it would catalyse all reactions that produce the remaining core (given an equivalent food set). Hence, the single molecule needed to recreate the core would be the constricted path. Of course, the information about the resulting structure is also encoded in the makeup of the food set and the fact that the exact interactions are a direct consequence of their structure; nevertheless, the seed molecule in a new environment will reproduce the viable core and its periphery and thus represents a self-replication process.

---

[30] Possible additional information channels may be epigenetic or through shared cell organelles.

These examples show that a frequent approach to coordinate reproduction is to tightly bundle information and have processes in place which unpack this information, creating an offspring SR from it. Hence, an ETI requires of the collective to evolve a way to encode information, store it in some form, and transmit it to the offspring – basically, this is the inheritance channel of the newly emerging level of individuality.

How exactly this may manifest depends on the structure of the involved SRs which may themselves already have a certain 'language' established with which they coordinate certain functions. For instance, in the realm of cells, the biomolecular machinery is ubiquitous and provides a communication channel; these structures can either be used directly, or they can be extended or repurposed such that they apply for further coordination between constituents of the collective. As a more specific example, consider the reproduction path of eukaryotes, which is based on DNA but – compared to prokaryotes – additionally includes chromosomal structuring and more advanced regulation mechanisms for gene expression. When turning to eusocial insects, their coordination occurs via a different language: chemical signalling via pheromones.

In summary, the coordination mechanisms needed for an ETI pertain in large parts to an efficient and robust reproduction path. From the examples studied here, it appears that new information channels are necessary to achieve this and that it may be beneficial to repurpose and extend existing communication methods.

I.3.3.4 Summary

With the transition to a new self-replicator, the description of the DAT processes concludes. Above, I motivated and argued how this transition can be the result of three classes of processes: diversification, aggregation, and transformation. Through diversification, an evolutionary system acquires novel structures; through aggregation, entities aggregate to form loose collectives that have an evolutionary advantage; and through transformation, the reorganise into a closely interacting collective which may become able to self-replicate as a whole.

These processes are an abstraction which aims to be general enough to be applicable to any kind of self-replicator. How this leads to an Evolutionary Unfolding with multiple hierarchical levels emerging will be further elaborated using the concept of the *Evolutionary Spiral*.

I.3.4 The Evolutionary Spiral

The *Evolutionary Spiral* brings the aforementioned constructs of Evolution Mechanics – self-replicators and DAT processes – together. It can be seen as a conceptual simplification which allows to think in terms of multiple evolutionary transitions rather than focussing on a single one. Furthermore, it inspires a visualisation that helps to illustrate the processes involved in EM and how they are intertwined.

The main idea of the Evolutionary Spiral is to situate the self-replicator and the DAT processes on a spiral trajectory, going upwards and outwards. The starting point of the spiral is the primordial self-replicator at the centre. Again taking an active and forward-looking perspective, 'walking' along the trajectory is associated with entities being subject to different sets of processes and potentially evolving as a result of these. Depending on the processes that have already taken place in a system, entities can be localised on a point on this trajectory. Subsequently, a full 'turn of the spiral' denotes that an entity was subject to all DAT processes and formed a new self-replicator. After this turn, the higher-level self-replicator acts as the new starting point for the next iteration, again going though conceptually equivalent processes on its way to the next hierarchical level.

In the following, I will introduce a visualisation of the above idea. The visual language used therein coevolved with the Evolution Mechanics framework and aims to provide a concise representation of the main ideas of Evolution Mechanics.

As such, it is rather general and need not be restricted to the specific context of the Unfolding of Life on Earth. For that specific case, the starting point of the trajectory may be associated with the origin of the prokaryotic cell – or any other SR that is deemed to be the primordial SR. In a more general setting, the starting point can be any abstract entity that qualifies as a self-replicator.

I.3.4.1 Visualisation

Figure I.3.1 is a schematic depiction of the idea of the spiral trajectory, taking a bird's eye view on it. The colour gradient denotes that the respective processes can not be delineated strictly from each other but may occur during overlapping time periods. As noted in Section I.3.3, not all processes have to occur strictly in order, either; Figure I.3.1 makes a visual simplification by assuming that they are somewhat in order, with the gradient hinting at the possibility that this may not be the case.

This first schematic is arguably *over*-simplified: It shows only a single, hypothetical lineage of SRs and does not visualise any dependencies between them, nor any substructure. Thus, it does not illustrate more than the idea of a spiral trajectory, where the angle is associated with the DAT processes and the number of turns with the number of ETIs.

A more illustrative depiction of the idea of the Evolutionary Spiral is shown in Figure I.3.2 – this is what we refer to as *the* (visualisation of) the Evolutionary Spiral.

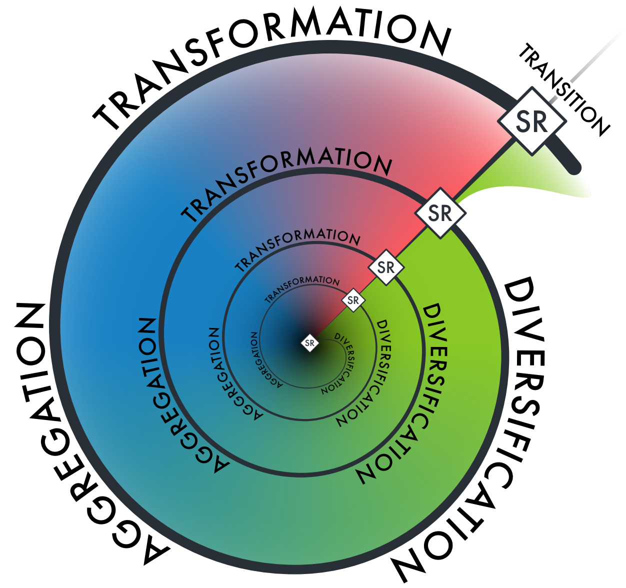

Figure I.3.1 Starting from a primordial self-replicator (SR) in the centre, this schematic visualises how the unfolding of complex structure by repeated modularisation may be associated with a spiral trajectory: Each 'turn of the spiral' represents the evolution of a SR to a higher-level SR via the DAT processes. The colours and the angle of a point on the trajectory refer to a certain 'progress' along that sequence of processes, with the gradient highlighting that they are overlapping.

Here, subsequent iterations of the spiral are collapsed onto the same circle[31], highlighting that they can be regarded as equivalent in terms of the DAT processes needed to evolve to a higher-level SR. The gradients outside the circle hint at the diverging and aggregating nature of the first and last half of one iteration, respectively. While the segments inside the circle denote the three hierarchy-generating processes, symbols outside the circle are added to illustrate them in the respective segment of the spiral, with each symbol denoting a SR that is a descendent of the initial SR and its shape morphing to symbolise a phenotypical change.

Assuming a hypothetical SR, traversing one turn of the spiral would go as follows:

1. The starting point is the SR, which is depicted as a perfect circle. The gradient hints at it having just-so acquired the capabilities to autonomously self-replicate.
2. The self-replication leads to diversification of the initial SR, symbolised by deviations from the perfect circle. The SR and its descendants disperse in

---

[31] We still refer to this as 'the spiral', despite it being visualised as a circle; conceptually, it represents a spiral.

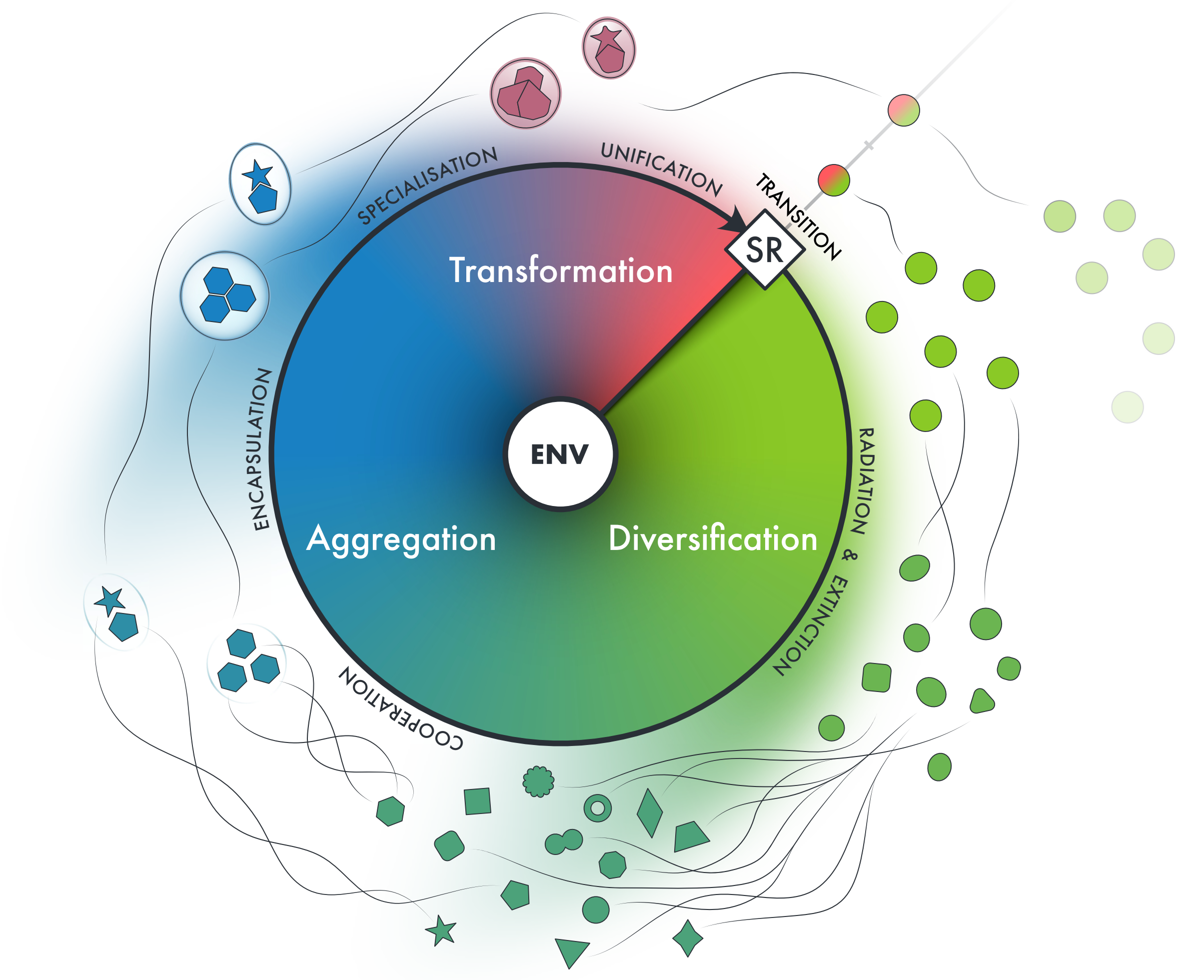

Figure I.3.2 The **Evolutionary Spiral**.
The spiral trajectory is collapsed onto a circle, emphasising the conceptually equivalent situation after each transition to a higher-level self-replicator.
The symbols on the outside, starting from the SR, exemplify particular aspects of the DAT processes: Diversification in structure and function, aggregation through cooperation and encapsulation, and finally transformation to a qualitatively new SR, from which point on the next iteration of the spiral may commence. All of this happens in coevolution with the environment.

niche space, affected by *radiation* and *extinction* phases and the coevolution of the environment. With niches becoming less easily accessible, competition makes high specialisation of SRs advantageous, causing a wide variety of structures.

3. At the beginning of the aggregating branch, *cooperation* between SRs offers additional advantages, leading to more frequent interaction – shown here with them being closer together and loosely encircled. Through cooperation, the level of selection begins to move from the individuals to the collective. This may happen between similar (fraternal) or dissimilar (egalitarian) entities; both options are depicted here and distinguished throughout the rest of the aggregating branch.

4. Through *encapsulation*, cooperating SRs strengthen their in-group interaction and reduce interference by other SRs.

5. Once encapsulated, *specialisation* of the individual parts of the group may lead to further variation of the SRs structure and their function within the collective. Specialisation is especially relevant if the group was formed between similar entities, which are not structurally diversified and may not yet have the capabilities to dynamically diversify within the collective.

6. During *unification*, the collective's constituents transfer capabilities to the group. This culminates in an internal reorganisation that leads to the encapsulated group of SRs becoming capable of autonomous self-replication: the *transition* to a higher-level SR. In this visualisation, the transition pertains to the point at which this was first possible; the process towards that point can be associated with part of the unification process.
   The symbol of the higher-level SR is again depicted as a perfect circle, albeit not the one the iteration initially started out with but on a new hierarchical level, further spiralling outwards…

This walk-through can be seen as a concise summary of the main ideas of Evolution Mechanics.

However, there is a certain potential for misinterpretation in Figure I.3.2, specifically of the added symbols on the outside, which I want to pre-emptively address here. The underlying cause for these misinterpretations is in the abstractions made throughout EM, their translation into a visual language, and the inherent difficulties of visualising an evolutionary process.

First, the symbols should be interpreted as nodes of a phylogeny, with lines representing some form of inheritance relationship. The phylogeny is reduced to a very specific subset: that of two kinds of SRs (fraternal and egalitarian) which went through the ETI to become a higher-level SR, with some peripheral SRs during the diversification phase. Hence, the symbols should be regarded as different kinds of SRs, potentially at different times; they do not refer to one and the same entity.

Second, the line from the fraternal and egalitarian branches towards the higher-level SR is not meant to suggest that the higher-level SR has *both* these branches as ancestors. The higher-level SR will be conceptually equivalent regardless of whether it resulted from a fraternal or an egalitarian group. In effect, walking the spiral trajectory along the symbols requires to *choose* one of the branches; from that point on, these are two separate trajectories.

Third, the entities are not in isolation but interact with their environment and other entities. There are typically many interaction partners, but only few are shown here: those with the same initial SR, emerging during diversification. However, due to the symbols representing a phylogeny, these need not even have existed at the same time, nor in relevant proximity (in niche space).

Fourth, the symbols illustrate the DAT processes as a procedure with strictly separated steps; this is far from reality, where they occur in parallel and in a coevolutionary fashion. While the DAT processes also attempt this isolation, words allow to articulate this more clearly; the visualisation can only hint at it using the gradients but has no apt solution in the case of the symbols.

Despite the above, Figure I.3.2 captures the central ideas of EM and casts them into a concise visualisation. With the wide range of abstractions made throughout EM, it additionally serves as the necessary simplification allowing to think ahead and not get lost in the forest of conceptual ideas.

Tree representation

The spiral representation shown above emphasises the conceptually equivalence of processes between transitions – it does not itself represent a particular evolutionary history of a certain line of SRs, but abstracts away from them. For that reason, trying to combine a phylogenetic tree and the idea of repeated processes is visually difficult: each branching point of the tree could be seen as the start of a new spiral trajectory.

The following is a brief exploration of an attempt to bridge these difficulties and visualise ideas from both Evolution Mechanics and phylogenetics.

Figure I.3.3 shows this attempt: A (purely fictional) evolutionary tree which starts from a primordial self-replicator at the bottom and grows into a tree akin to a phylogenetic tree[32]. Each line denotes the long-term evolutionary development of some SRs and a branching point denotes that the SRs speciated. The horizontal distance along the lines denotes diversification from the nearest branching point. The position on the vertical axis refers – very roughly – to how far a line of SRs have evolved along the trajectory of the evolutionary spiral, with dotted lines referring to transitions to higher-level SRs.

---

[32] Here, the alternative metaphor of the *coral of life* would again be appealing.

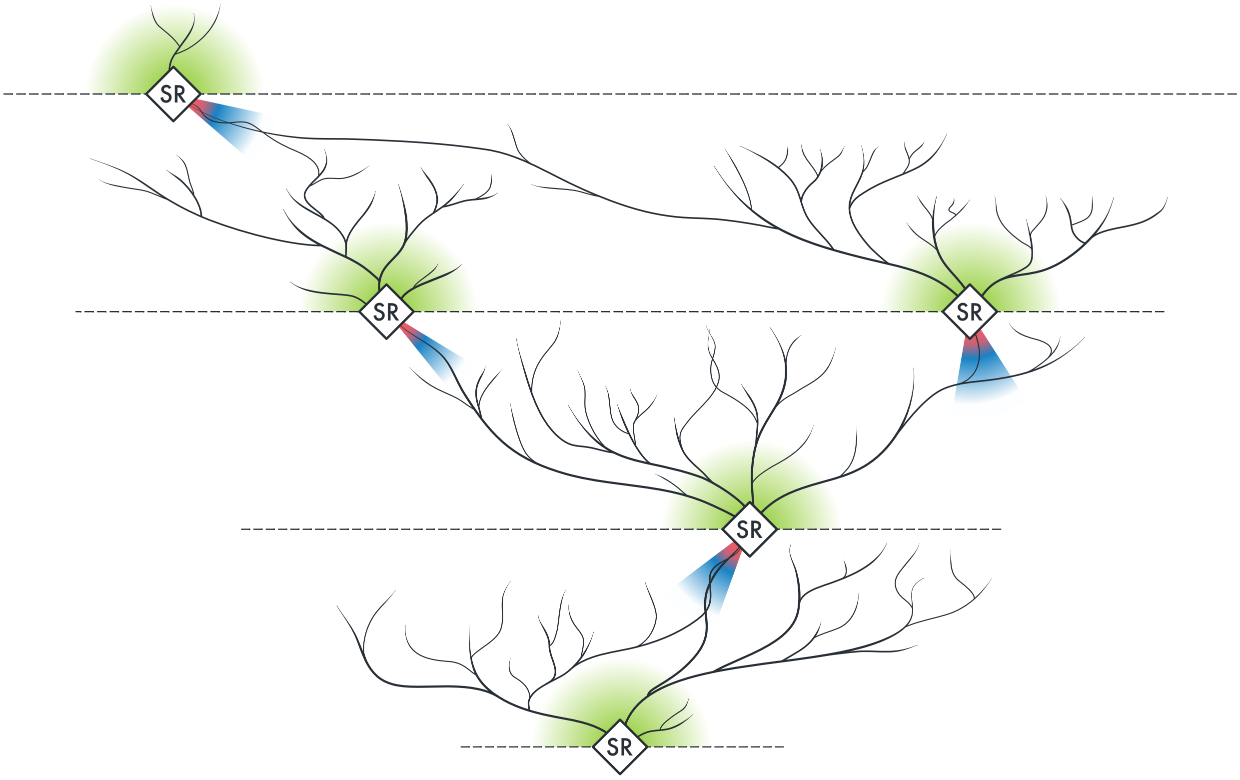

Figure I.3.3 A *fictitious* evolutionary tree with multiple evolutionary transitions (dashed lines). Starting from a primordial self-replicator, different lineages diversify (green gradients), some of which aggregate and transform (blue-red gradients) into a higher-level self-replicator. Lines that intertwine with other lines represent aggregation via egalitarian group formation; single lines denote fraternal group formation.
The horizontal distance *along* each line can be interpreted as the diversification from the respective branching point. The vertical axis can be associated with 'progress' along the DAT processes; it does *not* represent time.

Importantly, the vertical axis does not represent time, as in phylogenetic chronograms. In effect, two points at the same vertical position cannot be assumed to refer to SRs living at the same time, and a line ending does not necessarily mean that the SR became extinct but only that it did not further evolve along the spiral.

For instance, many lines end at similar vertical position before the highest-level transition. This denotes that these lineages have in common that any further evolution is disadvantageous to them, which is why they do not occur. The underlying reason for this evolutionary stasis may be of a structural nature; for instance, their structure may make it unlikely that they can successfully cooperate with other SRs. All this is valid only given the current point in time and configuration of the whole system – it cannot be precluded that these lineages would not evolve further given different circumstances.

Figure I.3.3 illustrates one particular aspect of evolutionary transitions quite nicely: Evolution Mechanics abstracts to levels of individuality, regardless of how exactly the new organisational structure of a SR is manifested. Hence, there can be multiple SRs associated with a level: All have in common that they evolved one additional organisation level that encompasses lower-level SRs.

Finally, a few assorted remarks that should clarify further aspects of Figure I.3.3:

- The vertical distance between two points can not be associated with a 'difficulty' to evolve; some evolutionary changes may occur in a short time, others may take a very long time or never find the right circumstances to be beneficial.
- This schematic does not show any cross-level interactions nor evolutionary transitions that are aggregates of SRs from different levels. However, all lower-level SRs are relevant in the evolution of higher-level SRs; the focus on a single level is only an approximation. This inherent multi-level nature of evolving systems is difficult to represent visually and was not attempted here.
- The environment is not represented here; it is implied that all evolution and interactions are on the backdrop of an environment, which is itself modified by the evolving SRs.
- Lines could in principle move downwards, meaning that losing some previously-acquired innovation was more beneficial than maintaining it.
- Diversification can be seen as a process acting on the whole niche space and all SRs; aggregation and transformation, in contrast, pertain only to a small subset of SRs. To depict this, diversification is hinted at with the 180ř gradients from each new SR, while aggregation uses the pie-shaped gradients that only encompass the relevant lineages that form a new SR.
- This tree does not show HGT, which certainly plays a role at multiple stages of the evolution along the spiral trajectory.

As can be seen from this list of comments and clarifications, Figure I.3.3 is clearly not a figure that speaks for itself. Yet, by providing a different perspective on the Evolutionary Spiral, it contributes to the understanding of core ideas of Evolution Mechanics.

Prospective and retrospective interpretation

The walk-through of the Evolutionary Spiral given above was primarily prospective: Which evolutionary processes need to occur in order for a higher-level SR to emerge?

While this was specked with observations of certain properties of SRs, the overall perspective was towards the evolutionary future of an SR, not its past. As initially mentioned, the language to talk and write about evolutionary processes can quickly become bulky. In such situations, using an active and prospective language is an appealing simplification that assists in conveying the main ideas.

However, evolution is not an active process of the evolving entities; it is a system's characteristic to aggregate certain changes more than others. This all happens in a variable and self-modifying landscape with an inherent stochastic nature, such that patterns only become observable when considering long time periods and many different manifestations of these processes. Therefore, I want to briefly contrast these two perspectives, starting with the differences in a purely retrospective view. I will then reflect on the extent to which projections are possible in the context of the Evolutionary Spiral.

Retrospection Looking back on the evolutionary history of a species of interest, the Evolutionary Spiral allows to trace the emergence of that species back to some primordial SR and additionally associates the corresponding evolutionary changes with certain types of processes. As such, it focusses on a very narrow corridor through a phylogenetic tree, backwards in time, including only those SRs that played a considerable role in their emergence. The effect of other SRs (e.g. through competition) are only indirectly included in the form of descendants from the nearest common ancestor; in Figure I.3.2, these are those other SRs that emerge as a consequence of diversification.

A difficulty in the retrospective view is that there is a considerable bias on available observations and subsequently in the many different branching points of the tree. Furthermore, given the multi-level nature of evolutionary systems, assessing which path is of interest may not be easy.

Projection In what way is the forward-looking view along the spiral trajectory sensible and can it be of any practical use? This would be akin to a projection, asking: Which processes still need to occur until a certain SR transitions to a higher-level SR?

While this view can in principle be taken, the main difficulty is to project the abstractions made in Evolution Mechanics back to the entity of interest. As such, a projection may formulate only the abstract necessities for a higher-level SR to emerge, but will not be able to make statements about the necessary evolutionary steps that this would take.

I.3.5 Summary

The Evolution Mechanics framework presented in this chapter aims to be a general theory for the emergence of complex, hierarchically organised structures out of simple ones.

It is a highly abstracted theory that isolates the key processes playing a role during the Evolutionary Unfolding of a system. Given the high level of abstraction, its contribution to the understanding of evolutionary systems is equally abstract: It can assist in finding common patterns between systems operating at different spatial and temporal scales, be they of a biological nature or abstracted to the realm of technological or cultural systems. Perhaps most importantly, it provides a coherent conceptualisation and language with which to address all varieties of evolving systems.

I.4

Discussion

Evolution Mechanics, as it was presented in the previous chapter, is a very young theoretical framework. Crucially, it has yet to be discussed – and scrutinised – in the wider community; I hope that this thesis (and those of my co-workers Harald Mack and Benjamin Herdeanu) may provide starting points for such a discussion.

In this chapter, I will first reflect on the Evolution Mechanics framework itself, including a contextualisation in the landscape of existing theories. I will then isolate open questions that go beyond the formulation of the framework itself, but may potentially be addressed by it.

I.4.1 Reflections on Evolution Mechanics

Evolution Mechanics is a highly abstracted framework of the Evolutionary Unfolding of systems. As a result, it can be useful in some areas while it may be too generic in others.

I want to reflect specifically on two aspects of Evolution Mechanics: First, in which ways it is a novel framework and whether the observational basis is sufficient to formulate it. Second, and referring directly to its formulation, how it could better address the multi-level nature of evolutionary systems. In addition, I will try to place Evolution Mechanics in the landscape of other theoretical frameworks of evolutionary systems.

I.4.1.1 Why Evolution Mechanics?

There is no scarcity of theories surrounding evolution and the Unfolding of Life (and I will take a more detailed look at them a bit further below). Given the many existing approaches, one may ask: why another theory? What does Evolution Mechanics contribute which is not already covered by other theoretical frameworks?

From what we see, Evolution Mechanics is novel in that it describes the Evolutionary Unfolding of *any* evolutionary system: It is abstracted away from biological systems and can just as well be applied to cultural or technological contexts, given the corresponding abstractions. At the same time, the level of abstraction is what limits the applicability of Evolution Mechanics. For instance, if focussing only on a single layer (say: niche construction in ecosystems, or cooperation in animal societies), existing theories allow to go into more detail regarding the concrete dynamics.

Evolution Mechanics aims for a holistic representation that includes both a conceptualisation of the processes relevant for a single transition as well as how these transitions relate to each other in the larger context of the Evolutionary Unfolding. It puts the Evolutionary Unfolding centre stage: the successive hierarchy formation that gives rise to complex structures. Evolution Mechanics highlights the similarities and patterns in the Evolutionary Unfolding and offers a conceptual language with which the evolution of systems can be formulated.

While many of the incorporated and abstracted processes are already well-studied on their own, we know of no theory that takes this level of abstraction and aims to formulate a coherent, overarching framework of the Evolutionary Unfolding.

I.4.1.2 Is there sufficient evidence for the abstractions made in EM?

In a second step, it should be reflected whether the abstractions made in the formulation of Evolution Mechanics have a sufficiently sound foundation in the observations of the Unfolding of Life on Earth. After all, there is only this single unfolding that we know of and there are many details of these observations which are unclear or unknown. For instance, a bias in which organisms fossilise and which do not would blur our view on the fossil record and of what we think occurred throughout Earth history – these are only one example of ‘unknown unknowns’ that, were they to become known, might affect the validity of the abstractions made in EM.

These and other biases certainly play a role when trying to associate certain observations with constructs from the theory. However, there are two arguments for why the formulation of EM can be regarded as reasonably sound, based on the presented observations.

First, the level of abstraction that Evolution Mechanics uses is primarily on the grounds of logical argumentations. For instance, it follows directly from the definition of the DAT processes that – once all of them occurred – a new organisational level will have emerged: (i) diversification brings novel structures into the system, (ii) aggregation forms new modules out of existing structures, and (iii) transformation makes the new module a self-replicator. Thus, by design, the result of one such iteration is conceptually equivalent to the starting point (with the implicit assumption that underlying levels play no dominant role).

Second, it is oversimplifying to regard the Unfolding of Life on Earth as a *single* unfolding. Yes, it can be seen as one large system; but given the vast time spans and the many conceivable subsystems evolving in parallel, we actually have a huge number of observations for the Evolutionary Unfolding. In addition, the observations presented in Chapter I.2 are reasonably general – eukaryotes, multicellularity, eusocial organisms: these are not isolated observations from Earth history, but wide-spread phenomena. Specifically, these examples are not cherry-picked to fit the narrative of Evolution Mechanics, but the abstractions made by Evolution Mechanics are what allows to associate many of the observations with EM concepts.

However, what Evolution Mechanics does not represent is the multi-level nature of these systems and the coevolution of all of their parts: EM makes the assumption that these can effectively be regarded as separated, while many of the presented observations suggest that this is specifically not the case. While this is a necessary abstraction in order to isolate individual processes, it puts the self-modifying, intertwined nature of evolutionary systems into the background without assessing how relevant this aspect may be compared to the highlighted processes. In the wider context of Evolution Mechanics, it is important to shed light on the multi-level nature of the system, which I want to address in the following section.

I.4.1.3 The multi-level, quasi-continuous, and self-modifying nature of evolving systems

As already stressed during the formulation of Evolution Mechanics, one key assumption is that the processes taking place in an evolutionary system are essentially the same on all organisational levels and that these can hence be generalised. While stressing that lower levels continue to exist and play a role, the ramifications of this coupling are not discussed.

Furthermore, Evolution Mechanics extracts the DAT processes and describes them individually. While noting that the processes occur in parallel and aiming to represent this in the visualisation of the Evolutionary Spiral, EM does not further clarify what higher-order effects result from the concurrency of processes and their effect on other parts of the system.

Ultimately, the Evolutionary Unfolding results in a multi-level structure with entities interacting across levels, distinctions between entities being fuzzy, and all processes having an effect on the nature of the system itself. Hence, Evolution Mechanics can offer only an approximation of the Evolutionary Unfolding.

I would argue, that it is very difficult to adequately represent all these aspects of evolutionary systems in EM. For instance, the existence of the self-modifying nature or the concurrency of processes in the system can be (and is) postulated. However, both these aspects can have very diverse effects on the whole system – simply postulating their existence is not of great value as long as their effects are not abstractable as well. Consider a system with multiple positive and negative feedback loops: which attractor this system would run into depends on the exact configuration of the system, and both the processes and their effects can be highly diverse. Therefore, trying to make abstractions of these processes is not feasible without further knowledge of the systems.

Making assumptions and approximations are necessities for theory formation, where tangible and discernible concepts are needed[33]. Evolution Mechanics tries to make as few assumptions as possible, while trying to extract the fundamental aspects of the Evolutionary Unfolding of systems. Given the scope laid out for the framework, I would argue that it does so in a rather minimalistic fashion.

Overall, despite its shortcomings, Evolution Mechanics can be useful: The framework allows to place the observations of the Unfolding of Life (and of other evolutionary systems) on the trajectory of the Evolutionary Spiral, highlights the fundamental processes that play a role in the formation of modular hierarchies, and perhaps inspires to think about the possible future evolution of a system.

I.4.1.4 Evolution Mechanics in relation to other theoretical frameworks

Having argued in which ways Evolution Mechanics contributes a novel perspective, it is useful to compare EM to existing theories and isolate where they differ. A comprehensive review of theories is not the aim here; instead, I will describe a selection of theories that are most interesting in this context.

Extended Evolutionary Synthesis

The *Extended Evolutionary Synthesis* (EES) [*Laland and Uller and* M. W. *Feldman* et al. 2015; *Pigliucci* 2007] is a recently developed theory of (primarily biological) evolution which aims to describe a wider range of observations of evolutionary systems. To that end, it includes concepts from evolutionary developmental biology (sometimes referred to as *evo-devo*), environment-evolution interactions like niche construction (*evo-eco*), and takes into account that inheritance and selection occurs can occur in multiple ways (*multi-level selection*).

The EES can be seen as an overhaul of the preceding *Modern Synthesis* [*Huxley* 1942] – it is debated whether such an extended theory of biological evolution is necessary at all [*Laland and Uller and* M. *Feldman* et al. 2014], with proponents of the Modern Synthesis claiming that there is lack of evidence for suggestions made as part of EES and that any extensions could better be made as part of the Modern Synthesis. One of the biggest differences between the two is in the role of the genome in evolution where EES allows for evolutionary interactions on multiple layers of selection while the Modern Synthesis takes a gene-centric view [*Laland and Uller and* M. W. *Feldman* et al. 2015, Fig. 1].

The largest difference in relation to Evolution Mechanics is the focus of EES on biological systems, while EM aims for a larger and more abstracted perspective that focusses on evolutionary transitions. Either way, EM should not be seen

---

[33] In a sense, this only reflects the limited capacity of humans to make sense of large and intricately linked networks of processes – and this statement in itself is an example of how we form mental models of systems we cannot fully understand.

as aiming to replace EES, but as integrating and abstracting some the ideas put forward in EES. For instance, the many ways in which species can diversify is well-described by EES and is the foundation for the EM process of diversification; similarly, the developmental plasticity that gives rise to division of labour is a key property of the convergent branch of EM.

Evolutionary Transitions

There are a number of theories that – like EM – aim to extract fundamental concepts of evolutionary systems, specifically in what leads to hierarchy formation. Often, the observation of evolutionary transitions is put into focus for these theories.

Perhaps most influential are the *Major Evolutionary Transitions* (METs) first described in [*Szathmáry and Maynard Smith* 1995] and reworked and extended in [*Szathmáry* 2015]. *Szathmáry and Maynard Smith* make the following observations about commonalities within transitions in biological evolving systems: (i) modularisation: smaller modules joining together to larger ones; (ii) loss of individual replication path: smaller modules no longer being able to replicate outside the larger module; (iii) new ways of information storage and transmission; (iv) functional diversification: smaller modules functionally diversifying within the larger module; and (v) new energy sources becoming available. These observations are integrated into EM as part of the DAT processes, with most of them playing a role in the transformation. As can be seen in this list, the focus for METs is primarily on the transition itself while EM tries to include the preceding diversification in an equally prominent manner.

Nevertheless, a comparison between the ETIs proposed by EM and METs is needed at this point. Overall, the overlap between the two is large, and the differences are mostly in what is being focussed on: In essence, ETIs refer to the change in individuality alone and categorise this more broadly and as successive changes, while METs refer to specific events and include more detailed considerations of the accompanying changes. For instance, METs separate the origin of eukaryotes and of plastids into two METs, while they are associated with a single ETI: simple cells coming together into a more complex self-replicating structure. The ETI thus need not refer to a single event, but primarily to the emergence of a new individuality; it ignores specific changes in the exact manifestation of a transition, as long as it can be classified as a new individuality[34] – something which is explicitly distinguished in METs. A similar classification of ETIs as made in EM is presented in [*West* et al. 2015], focussing on the levels of individuality.

The notion of evolutionary transitions inspired a number of interesting other works which contain conceptualisations of the Unfolding of Life in some form or the other. While I cannot give an in-depth review, I at least want to briefly mention them here.

---

[34] As shown in the the tree of life representation in Figure I.3.3.

Coming from the realm of social evolution, *Bourke* formalises mechanisms that lead to stable biological grouping, the maintenance of these groups and how they may modularise. It inspired the aggregation and transformation phase of EM, where many of the ideas are integrated.

Lenton and A. *Watson* [2011, Fig. 20.1] also go along the story line of the Unfolding of Life. They describe this unfolding in the shape of four 'revolutions', namely: the inception (of life), the oxygen revolution (leading to eukaryotes), the complexity revolution (multicellular life), and the emergence of humans. The shared characteristics they isolate are categorised into changes in information storage and transmission, organisation, energy and matter flow, and prevalent Earth system instabilities; they then go on to characterise the effect of a revolution on the recycling mechanisms in the Earth system and the new stable state the system may attain.

These revolutions are picked up by *Judson* [2017], but with a focus on different energy forms throughout Earth history: geochemical energy, sunlight, oxygen, flesh, and fire. *Judson* suggests viewing the evolutionary history of life on Earth through the lens of the different energy forms, successively being discovered through evolutionary events and transforming the planetary environment as a result.

R. A. *Watson and Mills* et al. [2015] take a completely different, more conceptual perspective on the unfolding of evolutionary systems: they describe algorithmic principles that represent underlying processes of evolutionary systems. Building on this, they highlight similarities to learning processes and the self-modifying nature of evolving systems [R. A. *Watson and Szathmáry* 2016, Fig. 2].

As the last work to mention here, *Solé* [2016] argues that not only the METs observed in nature can be used for gaining an understanding about the Evolutionary Unfolding, but also the 'synthetic' METs which occur, for instance, in engineered cooperative systems, designed multicellular systems, evolved communicating robots, or the pattern recognition and language processing capabilities of artificial neural networks. Comparisons between technological evolution and biological evolution are made in [*Solé* et al. 2013].

Artificial Life

A research field not mentioned so far is that of *Artificial Life* (ALife), where the origin, potential, and limits of life are investigated through biochemical or robotics experiments or with the help of computer models [*Taylor* 2014]. At its heart, Evolution Mechanics aims to answer how complex structures may emerge from simple ones; similar questions are posed in the context of Artificial Life research. Thus, I want to give a slightly more detailed view into that field.

Early work that may be attributed to the field of Artificial Life is that of von *Neumann* [1966] on self-replicating systems. Using cellular automata, he demonstrated how a universal self-replicating machine may be constructed and grow in complexity, akin to biological evolution. While being an interesting proof of concept, these self-replicators proved to be far too slow and fragile for studying evolutionary dynamics [*Taylor* 2015].

As another example, the work by *Braitenberg* [1986] was influential: He explored how relatively simple 'cybernetic vehicles' could behave in complex ways, despite just being made up of a bunch of sensors and motors. The cybernetic vehicles highlighted how even these organisms may be *perceived* as having intent and how few mechanisms would be required to construct them.

One often-posed question in Artificial Life is that of open-endedness: What is needed for a system to exhibit open-ended evolution? This is a question that has large overlap with one that is of interest in EM: What underlies the transition to a higher-level self-replicator? If such a question is sufficiently answered, open-endedness could be a follow-up question.

As stated in the formulation of EM, it is the structure of the entities and their interactions which determines how complex the evolving structures may become and whether they are capable of organising into higher-level structures. How a sufficiently complex structure can be provided to a system in order to evolve is often studied in software-based Artificial Life – and can thus help to substantiate which capabilities are necessary to allow further evolution.

The major difficulty in addressing this question is that the physical world provides a wide range of entities and interactions and their form-function relationship is 'just there', whereas any software-based representation of a system is defined by the deliberately made choices when constructing it. Basically, all form-function relationships need to be provided by an engineer, directly or indirectly. This might not pose a limitation if there would be sufficiently many states and interaction capabilities between entities – but typically, virtual worlds are low-dimensional and are particularly limited in the range of interaction between entities. In short: it appears to be difficult to provide a substrate to *in-silico* evolutionary systems that allows for the system to evolve novelty [*Taylor* 2015] – or something that would be associated with being novel.

Same as in EM, the self-replicating entities play a central role in Artificial Life research. For instance, *Taylor* [2015] claims that the following capabilities of a self-replicator are 'necessary and sufficient for a system to exhibit open-ended evolution': (i) robustly reproductive individuals, (ii) a medium allowing the possible existence of a practically unlimited diversity of individuals and interactions, at various levels of complexity, (iii) individuals capable of producing more complex offspring, (iv) an evolutionary search space which typically offers mutational pathways from one viable individual to other viable (and potentially fitter) individuals, and (v) a drive for continued evolution. This list is very close to statements made in the context of EM, specifically that of robust self-replicators and a sufficiently diverse medium. (Interesting further thoughts along this line can be found in [*Taylor and Dorin* 2020].)

Despite the similarities, one key difference between the two frameworks is their focus: The way EM is structured, the focus is not on open-endedness but on a single transition, spelling the involved processes out in more detail. Furthermore, while the self-replicator is a central component of EM, the environment and the interaction and aggregation with other self-replicators play an equally important role and cannot be reduced to the capabilities of the self-replicator or a general notion of a medium.

Overall, it appears that the more conceptual questions asked in EM – specifically pertaining to the structure of the environment and the self-replicators – are overlapping with questions in the field of Artificial Life. Research in Artificial Life also illustrates the difficulty of studying Evolutionary Unfolding using computer models, something which would also be desirable for further studying Evolution Mechanics. It might thus be interesting to consider how models of artificial life may inform the formulation of Evolution Mechanics and where the conceptual work done in Evolution Mechanics and Artificial Life may inspire each other.

I.4.2 Open questions

What are open questions in the understanding of the Evolutionary Unfolding of systems?

Questions pertaining to this topic go beyond the specific formulation of Evolution Mechanics, but it is useful to formulate them in the language of the framework. In the following, I will pick up some of these questions, spell them out, and try to answer them as far as possible using EM. This discussion will not bring ultimate answers to these questions, but it may help to isolate in which ways future research could further address these.

I.4.2.1 How important are compartmentalisation & dispersion for the Evolutionary Unfolding of system?

In the emergence of prokaryotes, confinement to separate compartments and subsequent dispersion appears to have played a crucial role; in the other transitions, this is less obvious. Hence, it might be worth reflecting: Which general properties of compartmentalisation and dispersion influence the Evolutionary Unfolding?

First off, compartmentalisation should be interpreted in a very general sense: It may refer to the spatial confinement to a certain volume which increases concentrations of the system's particles, but it may also be a process that leads to an interaction becoming more selective by some other means (like when entities specialise for a certain niche or an innovation making a process more efficient). Hence, a compartmentalised system is one that is not well-mixed, but is heterogeneous in its environment or the interaction structure it allows – essentially: a system that has distinguishable niches.

Dispersion is the process that is counteracting compartmentalisation: it causes entities that were previously not interacting to begin to interact in a competitive, cooperative, or neutral fashion. This can be caused by spatial dispersion, but also by reducing the selectivity of an interaction. Furthermore, it may lead to reduced interaction rates, for instance because local particle concentrations were reduced. Generalising this to the picture of niche space, dispersion is the process by which niches become less defined or even merge.

Both these processes have strong effects on the Evolutionary Unfolding of a system. For an illustration, consider again the case of simple reaction networks (Section I.2.2.3 and appendix A.1). As demonstrated in [*Vasas* et al. 2012], such a system in a well-mixed form can only ever select one attractor – it may show competition within that system, but will ultimately fall back onto the single attractor. By allowing compartments and low exchange rates between them, competition between compartments becomes possible and the notion of a viable core *replicating* to another compartment becomes sensible. Through novel reactions, variation can be aggregated, which is not possible in a single compartment. This multi-compartment system thus supplies multiple niches, which are sufficiently separated to allow a parallel exploration of state space.

The parallel diversification in state space, enabled by some form of compartmentalisation, generalises to all evolutionary systems and is perhaps the most important effect of compartmentalisation. Instead of having a single system with a single configuration of selection pressures, several smaller systems will show many different selection pressures[35], which allows for more diversity. Subsequently, when then brought into contact through some form of dispersion, there is a wider diversity of innovations to compete against each other.

For the overall evolution of a system, these compartmentalisation-dispersion-cycles can thus speed up evolution if innovations in one compartment are also beneficial in another. In spatially compartmentalised systems this is often the case for nearby compartments, exceptions being very specific adaptations to the local environment which are not beneficial enough in other environments; something similar can be said for general innovations.

As a remark, compartmentalisation and dispersion are particularly relevant in combination with the effects of radiation and extinction events (Section I.3.3.1) – these are the processes that can be used to describe how a system explores its state space.

The exact degree to which these processes play a role and how important they are in relation to other processes cannot be answered in a general fashion. It can be presumed that there is some optimal balance between them which may lead to an optimal exploration of state space, but how exactly this would play out depends strongly on the assumptions made for a particular state space.

---

[35] Even if they start with identical selection pressures, they will evolve to a more diversified set of pressures.

One option to investigate the importance of these processes would be by imposing compartmentalisation-dispersion cycles on a system and study how they affect the movement in state space. For spatially resolved systems, there are many studies along that line, particularly revolving around agent-based models. An obstacle would be to abstract this to higher-dimensional state spaces and include generalisable heterogeneities which can be associated with the heterogeneities of the state space of actual systems.

In summary, and to directly address the initial question: Compartmentalisation and dispersion are no necessary conditions for a system to evolve, but they drastically increase the space of possibilities, how quickly the system can explore that space, and thus the complexity of the structures that can evolve within the life time of the system. Further studying these processes may help to understand their overall role in the Evolutionary Unfolding of a system and whether an optimal balance between them exists.

I.4.2.2 Are egalitarian transitions 'more difficult'?

From the observations of the Unfolding of Life on Earth, it appears that egalitarian transitions only occurred (basically) once: in the endosymbiosis that led to the origin of eukaryotes (see Section I.2.4). Even if the presumed second endosymbiosis that incorporated proto-chloroplasts and created plant cells would be counted separately, the observed fraternal transitions to multicellularity or in eusocial organisms are far more numerous.

This prompts the question whether there are some inherent difficulties to egalitarian transitions in contrast to fraternal transitions and, if so, what they are caused by. These questions pertain to the processes on the convergent half of the Evolutionary Spiral, aggregation and transformation, and the differences between the two.

As stated in [*Bourke* 2009, Table 1.2] and throughout the formulation of the DAT processes (see Section I.3.3), egalitarian and fraternal transitions differ primarily in their diversification at the point of their initial aggregation and the potentially required specialisation during their transformation: Egalitarian transitions need to evolve mechanisms to reduce the conflicts of interest between the involved parties and then coordinate their replication as a collective. For fraternal transitions, there is an initial difficulty to establish an evolutionary advantage against intra-specific competition; once that has happened, there is the need to coordinate functional diversification within the collective.

The formation of collectives of dissimilar entities is generally easier, because entities of different species typically do not compete for common resources. Thus, if some cooperative interaction between them arises spontaneously, it will directly yield an evolutionary benefit. Even if that benefit is minimal, it will directly have a positive effect on the involved species' fitness. These expectations are matched by the many examples for mutualism in nature.

In contrast, for competing entities of the same species to cooperate, the benefit from cooperation needs to be larger than the cost of their cooperation, otherwise they will not have an evolutionary benefit against their non-cooperating competitors. Of course, these dynamics become intensified if intra-specific competition is large. This makes the formation of fraternal collectives more difficult than in egalitarian collectives.

However, once the fraternal collective is robustly cooperating, it will also have aligned its internal interests, because this is a necessity for robust cooperation. The egalitarian collective, in contrast, may be a group of cooperating entities, but they may still have conflicts of interest which prevent them from becoming more unified.

The example of the origin of eukaryotes can be useful at this point. One particular aspect of that transition is the role of HGT in allowing to outsource functions of the mitochondria to the enclosing cell, thus becoming more dependent on each other (and more unified). What can be seen as facilitating such a process is that both entities of the symbiotic relationship use a compatible language: Both their cell processes consist of essentially the same biomolecular machinery, with a DNA-based inheritance channel; while mitochondria still have their own DNA, they rely on the enclosing cell functions for replication. The similarities of cell processes, owed to the close evolutionary history of bacteria and archae, made it possible to become unified without the need for a novel communication channel between the two to evolve.

Thus, a compatible language between the constituents of an egalitarian collective appears to be an important factor for a subsequent transition, as it alleviates the need for such a language to emerge within the collective. This would also explain why egalitarian transition are not observed for more complex structures than eukaryotic cells: For an egalitarian collective of more complex organisms, the language of the biomolecular machinery is no longer accessible for coordination; hence, another inter-specific language needs to evolve, which can indeed be seen as a difficult step. For fraternal collectives, the coordination mechanism only needs to evolve within the species, which I would speculate to be an easier step.

Have holobionts performed an egalitarian ETI? An adjacent question is whether the emergence of holobionts constitutes an ETI. Shedding light on this question is also illustrative for the differences between egalitarian and fraternal transitions in terms of the evolution of a reproduction path, a crucial part of a transition.

In holobionts, there are often mechanisms in place which facilitate that the microbiome of a parent can propagate to the offspring; for instance, the birth channel in mammals or close contact with members of the same species. Furthermore, host organisms often require the existence of symbionts for their functioning; as such, the symbiosis is not facultative. At the same time, symbionts may require the environment provided by the host for their survival.

These observations make it appealing to consider the possibility that the holobiont is in fact an egalitarian collective, arising from the egalitarian transition of a multicellular eukaryote and microbes of several other species. However, holobionts – or: their hologenomes – are typically not considered to be a unit of selection [*Moran and Sloan* 2015]. While there are examples where the holobiont can be seen as the unit of selection (including eukaryotes, but also some insect species), these cases appear to be rare and do not allow to consider all holobionts as forming an evolutionary unit. The main reasons are that (i) the inheritance of the microbiome generally shows a large variance; (ii) the microbiome can also be picked up from the environment, not requiring any facilitating mechanisms; and (iii) the evolution of the microbiome often does not contribute strongly to the fitness of the host. *Moran and Sloan* stress that even if the fitness benefits of host and symbionts may align, this does not imply that they form an evolutionary unit; it may simply reflect the fact that the host has evolved to depend on a reliably available part of its environment. Combined with the diverse host-symbiont interactions [*Foster* et al. 2017], including the hosts ability to keep the microbiome ecosystem in check, this is in some parts reminiscent of the idea of the extended phenotype [*Dawkins* 1982], albeit without the strict focus on the genes of the organism.

Having discussed the multi-level nature of evolutionary systems above, I would like to offer a perspective on this question where the holobiont could be seen as having performed some kind of 'weak' egalitarian ETI: a transition that does not change the level of selection, but causes a sustained host-symbiont relationship. Such a weak ETI describes the process that causes an organism to become a holobiont. To be clear, I am not proposing that the hologenome becomes the unit of selection, but I am illustrating how a holobiont can be described in the language of multi-level Evolution Mechanics.

What characterises a weak ETI for holobionts is that they (i) have some general dependence on each other; (ii) have evolved an inheritance path (parallel to the existing inheritance path) that facilitates transmission of the microbiome to the offspring; and (iii) show coordination mechanisms that suggest a coevolutionary history. The dependence of the two is characterised by the host requiring that some *general* form of microbiome fulfils certain functions for the host; for the symbionts, the dependence is in form of a habitat they may be more suitable for the typical population of symbionts. Neither host nor symbiont depend on a completely specific interaction, but both have benefits from taking part in a roughly similar interaction.

The parallel inheritance path makes it easier to maintain such a situation, even if it is lossy or introduces high variance. But again, neither host nor symbiont fully depend on the fidelity of the inheritance path.

Lastly, the coordination mechanisms show that the host-symbiont relationship is not a mere spontaneous aggregation but a part of at least one of the two parties life cycles. Given that the host is more complex, control mechanisms like those classified in [*Foster* et al. 2017] will often appear in the hosts; in Evolution Mechanics this would still be associated with a specialisation that occurs in the transformation stage of a transition, but one where only one part needs to adapt. Of course, as *Moran and Sloan* argue, the intimate relationship of host and microbiome does not imply that any adaptations in the two are the result of aligned fitness benefits in their evolutionary history. However, I would argue that the *general* capacity of a host to control its *general* microbiome can be seen an adaptive process; this would be a weaker requirement than the one referred to in [*Moran and Sloan* 2015].

For the picture painted above, the multi-level view and the weak requirements are crucial, as they allow to see the weak ETI as a parallel process that does not substantially change the level of selection. As such, a weak ETI is also different from an 'unfinished' ETI, because the latter would already have moved the level of selection towards the group. This is not to say that the interaction arising from the weak ETI may not evolve to become more important at some later point[36], given that selection pressures are never fully restricted to a single level.

In summary, this is a proposal for a multi-level view of holobionts that is (with some extensions) compatible to the Evolution Mechanics framework. Additionally, it is an example of frequently and robustly occurring egalitarian ETI – with the limitation that they have to be considered 'weak'.

Summary I tried to isolate the different factors playing a role in egalitarian and fraternal ETIs and which hurdles there may be in such a transition. The answer to the initially posed question is that there are two main conceptual reasons that would explain why egalitarian ETIs have been observed far less often than fraternal ETIs: With rising complexity of organisms, their communication layers decouple more strongly from each other such that the constituents of egalitarian groups will likely not have a shared language; this is hindering the emergence of a coordination mechanism that aligns the conflicts of interest within the group. In fraternal groups, this appears to be easier primarily because once the initial advantage of cooperation is established, the entities benefit from being similar and whatever coordination mechanism evolves is robustly inherited. In addition, the larger number of entities in fraternal collectives allows even for simple coordination mechanisms to cause collective behaviour; something that is not the case for egalitarian collectives.

Thus, transitions that shift the level of selection to that of the newly-formed collective appear to be easier in fraternal collectives. To account for the many examples of holobionts, I proposed a perspective by which they can be understood as having undergone a 'weak' ETI, which highlights the multi-level evolutionary nature of these complex life forms.

---

[36] Something that would need to be discussed for the case of eukaryotes or leafcutter ants. For the latter, which depend on a specific fungi and require its inheritance through a parallel inheritance path [*Hölldobler and* E. O. *Wilson* 1990, ch. 17], the symbiosis is obligatory for both parties, which makes this an interesting – yet anecdotal – example of parallel reproduction paths (via the queen genome *and* the fungi it carries along when forming a new colony).

I.4.2.3 What determines evolvability?

At several stages throughout the description of Evolution Mechanics, the capabilities of a system to evolve and form complex structures was of relevance. For instance, a system of auto-catalytic reactions was brought forward as an example of limited evolvability. At the opposite end, the capabilities of the eukaryotic cell were posited as a key requirement for complex multicellular life to evolve.

Evolution Mechanics is a conceptual framework and, as such, it cannot directly address this question. In general, questions of evolvability are much more accessible on a mechanistic level, for instance by studying the proteins that may be synthesised from a gene and trying to isolate the effect of this protein on a larger organism: This is a direct investigation of the relationship between the encoded information, the resulting structure, and its functionality in the environment it is embedded in. By analysis of the genome, the space of possibilities could then be mapped out and used to assess which neighbouring structures could emerge.

Of course, such an analysis is not as easy as it is made to sound here. But regardless of how precise such a method may become, mechanistically studying a system will only assess some form of 'local' evolvability and does hardly allow inference of its capacity to successively evolve into a hierarchically structured system: Looking at mechanistic aspects of a prokaryote alone would not allow to infer that this would evolve to plants, animals, or culture. At this point, Evolution Mechanics has a wider perspective because it isolates which processes need to occur in order for evolutionary transitions in individuality to occur – EM cannot answer if or how exactly this would occur, but it proposes abstraction that allow associating mechanistic processes with each of the individual steps towards such a transition.

With both these approaches only coming marginally closer to an answer to the initially posed question, perhaps it is worth reflecting what kind of answer would make sense for that question. It is clear that highly reductionistic answers ('The Big Bang determined all evolvability') or highly detailed, retrospectively compiled lists of properties of complex life are not particularly insightful in the larger context. Furthermore, it needs to be clarified what kind of potential is measured by evolvability: the number of transitions, the evolved capabilities, some form of dominance over other organisms, or some other measure? Depending on what potential is of interest, widely different answers can be expected.

Evolution Mechanics only allows making some general statements regarding the number of transitions. Once a SR grows to approximately the size of the system, it will become more difficult to form aggregates and transform into a higher-level SR. To evolve to higher levels in such a scenario would either require to get access to further resources or reduce its size while maintaining the hierarchy. For the latter response, lower-levels would need to become 'compacted', a process which is limited by the minimal structures that are needed to maintain the level's function within the hierarchy (essentially: the replication and coordination mechanisms within and between modules, which cannot easily be compacted).

The former is more difficult to spell out, because a system's boundary depends on the available innovations and their emergence cannot be foretold. Consider a world in which oxygenic photosynthesis was not discovered – life would most certainly look drastically different. Similarly, humankind would presumably not have gained a larger influence than other hominids were it not for their capacity for language and the ensuing abstract communication capabilities. It is hardly conceivable how future transitions would look like, especially if they occur in the realm of cultural evolution.

While it is clear that truly open-ended evolution (in the sense of evermore unfolding hierarchical levels) is not possible in a finite Universe, how many organisational layers may appear in between appears to hinge on key innovations that either allow compactification of existing levels or unlock new niches. This connects back to the questions posed above in the context of Artificial Life.

Overall, evolvability remains an elusive concept; the topics discussed above should be seen in the corresponding light. While Evolution Mechanics does not offer answers that go beyond fairly general statements, it can be used to characterise such a scenario and assess where the difficulties are in answering such a question.

I.5

Conclusion & Outlook

The preceding chapters were a journey not only through Earth history and the Unfolding of Life, but also along a number of unexplored paths with the aim of finding a conceptual description for these observations. The framework at which we arrived at and which I formulated in this thesis finds abstractions for these observations and defines a set of processes that, in combination, allow a general representation for the evolutionary unfolding of a system.

Throughout this journey, the view was directed primarily into the past. Naturally, it is appealing to turn the view into the future and attempt to apply the concepts of Evolution Mechanics to humankind's cultural evolution and perhaps its potential future development. Starting from trying to position early and more recent human societies' on the Evolutionary Spiral, it may be tempting to project which capabilities human societies would still need to acquire to go through an Evolutionary Transition in Individuality.

While appealing, thoughts along these lines remain highly speculative: Evolution Mechanics may assist in isolating the abstract necessities for human societies to 'traverse along the spiral trajectory', but – as argued before – it does not allow projections of how, when, or even *if* these can be attained, neither does it draw a feasible path to drive such a development.

In addition, these thoughts touch on a fundamental philosophical question: Even if aware of the abstract necessities of an ETI (which in itself requires a certain degree of consciousness), can a system ever be capable of engineering its own transition in individuality? A question of that sort is certainly out of the scope of this thesis – and may perhaps elude a scientific consideration altogether.

However, there are certain aspects of our world and its near future where the mechanisms comprising Evolution Mechanics play important roles. For such applications, the language provided by this framework and the conceptual links that are made within it may contribute to their investigation.

For instance, as motivated initially, human civilisation became a global force on the planet. The actions of the next years will strongly influence the trajectory of the Earth system (at least on the scales relevant for us as humans) – however, the degree to which humankind can coordinate decisive actions on the scale of the global community is unclear. By studying how the entities of a system may begin to cooperate and form groups, including the effects on other layers of the system, beneficial situations may be isolated. These topics are core aspects of the converging branch of the Evolutionary Spiral and hence allow a description using EM.

In addition, in light of the changes to small- and large-scale ecosystems, understanding their evolution and response to these changes becomes increasingly important. With ecosystems and the interactions within them being the product of an evolutionary process, continuing to diversify and adapt, the concepts of Evolution Mechanics may provide a novel perspective on the evolution and behaviour of ecosystems – the topic of the next part of this thesis.

Part II

Ecosystem Evolution

II.1

Introduction

For the second part of this thesis I focus on the evolution and behaviour of ecosystems. Specifically, I use a numerical model of food web evolution to study the resulting ecosystem structure and its resilience against abrupt changes.

Building on the remarks I made in the prologue of this thesis, I will first motivate this research topic and set it into relation to Evolution Mechanics. I then formulate a set of research questions for this part of the thesis and give an overview of the modelling approaches used to describe ecological interaction networks.

II.1.1 Motivation

In the context of this thesis, there are two main motivations for studying ecological interaction networks:

First, the Evolution Mechanics framework presented in Part I and the study of ecosystems are strongly linked. Ecosystems, in the widest sense, may be seen as comprising all life on the planet; after all, the Earth system can be understood as a planetary-scale ecosystem. Equally so, the interconnected pores of a hydrothermal vent can be understood as an ancient small-scale ecosystem.

In either case, ecosystems are biological evolutionary systems, a subset of the structures and interactions that emerged throughout the Unfolding of Life. The state we observe an ecosystem in is the result of their evolutionary history, thus linking the investigation of ecosystems with the concepts put forward in Evolution Mechanics.

Given this connection between the two, increasing our understanding in one of these fields may feed back to the other. The observations of ecosystems already informed central parts of the conceptualisations made in Evolution Mechanics. In turn, the ideas of EM provide a language to describe and locate evolutionary processes in ecosystems. Specifically, for the study of ecosystems in this thesis, EM emphasises the importance of the evolutionary history of a system if aiming to understand its further development; studying these processes may then inform the formulation of EM.

The second motivation is before the backdrop of the Anthropocene (as briefly characterised in Section I.2.8.2) and the situation humankind finds itself in at this time: in the midst of a ‘climate emergency’ and a period of mass extinctions, both of which are caused by human activity [*Ripple* et al. 2019]. In the light of the effects of large-scale habitat loss, a strong decline in ecological diversity, and accompanying feedback cycles that have wide-ranging repercussions, a better understanding of the dynamics of ecosystems may help to focus efforts of keeping the Earth system within a safe operating space [*Bak-Coleman* et al. 2021; *Steffen and Rockström* et al. 2018]. In particular, better understanding the behaviour of ecosystems when being subjected to anthropological effects can contribute to isolating the important processes for maintaining biosphere integrity [*Steffen and Richardson* et al. 2015].

One of the approaches to study ecosystem dynamics is by representing them as ecological interaction networks, as done in the field of food web ecology. This representation is an approximation of the complex structures comprising ecosystems and the interactions taking place within them, putting a focus on trophic interactions, and it allows to be approached with numerical simulations of these models. As argued above, this approach should take into account the evolutionary processes giving rise to the structure of ecosystems.

With my background in physics – and not ecology –, I will focus on the general processes that play a role in the evolution of food webs, including those mechanisms that are relevant when aiming to study the response to anthropogenic effects. Necessarily, this will be at a level of abstraction that cannot be *directly* applied to the above context; yet, it may contribute to the approaches taken in modelling ecological systems using evolutionary food web models.

Condensing the above, the research questions studied in this part of my thesis are the following:

1. In which ways can Evolution Mechanics and the study of ecosystems inform each other?
2. How do food webs respond to locally or globally introduced perturbations?
3. Which mechanisms make a food web resilient to these perturbations?
4. Are evolutionary food web models a suitable approach for studying these questions?

II.1.2 Background

Here, I provide some background information on studying ecosystems, focussing on the aspects that are relevant in this thesis. I begin with a brief historical overview of how these systems have been investigated, up to recent approaches on the modelling of the Earth system. What follows is a description of the models typically used for representing ecosystems and the assumptions they make for the representation of species, their environment, or their interactions. Finally, I describe a framework for more concisely formulating the often-vague notion of *resilience* in the context of ecosystems.

II.1.2.1 History

The study of ecosystems has a long history – especially if including its roots in the general observations of natural systems, something which has always fascinated humans. I want to highlight some of the important points on the way, with a focus on more recent research in theoretical ecology and the study of food webs in the Anthropocene.

Population dynamics One historically influential aspect was the study of population dynamics. An early observation was that of the tendency of populations to grow exponentially [*Malthus* 1798], leading to the observation that, under limited resources, this would necessarily cause resource shortages. This description of population growth was later expanded to include exactly these resource limitations, leading to the logistic growth model [*Verhulst* 1838] which forms the basis of many descriptions of population dynamics.

Another important mile stone was in the study of predator-prey systems: systems where one species feeds directly on another, thus coupling their dynamics. In the 1920s, Alfred J. *Lotka* and Vito *Volterra* – separately from each other – studied such systems and created the mathematical formalism to describe them, the *Lotka-Volterra equations* [*Lotka* 1920; *Lotka* 1925; *Volterra* 1926].

On a more conceptual level, the study of competitive systems by *Gause and Witt* in the 1930s brought forward the formulation of the *competitive exclusion principle* (CEP) [*Gause* 1932; *Gause and Witt* 1935] which had a big impact. It can be seen as a key principle of competitive systems, even beyond ecology; its generality and validity was also the topic of fierce debate [*Hardin* 1960].

Complexity-stability debate In the early days of food web ecology, it was intuitively argued that larger ecosystems (with more species and more links between them) would be more stable than smaller ones, e.g. because the displacement of one species could be buffered by another or strong population growth would become more unlikely [*Landi* et al. 2018].

In the 1970s, however, work by *May* challenged this argumentation and sparked the so-called *complexity-stability debate*, revolving around the question whether 'a large complex system [will] be stable' [*May* 1972]. Using a random community matrix, representing interaction strengths between species, and linear stability analysis, he demonstrated that small perturbations of these systems' equilibrium state will more likely lead away from the equilibrium if the system contains more species or a larger number of links.

Since then, a large body of work in theoretical ecology focussed on shedding light on this question, not only by making more realistic assumptions about the form of the community matrix, but also by choosing other forms of representing food web dynamics. For informative reviews about this question, see [*Brännström* et al. 2012; *Fritsch* et al. 2021; *Landi* et al. 2018].

Earth system modelling With anthropogenic effects having become more prominent over the course of the last century, there are fewer and fewer scenarios in which an ecosystem may be regarded as sufficiently isolated to disregard them. Subsequently, there are an increasing number of models that attempt to include additional dynamics of the Earth system, including climate and ecosystem models, but ranging into the sphere of human societal and economic dynamics and their effects on ecosystems [*Donges and Lucht* et al. 2017; *Donges and Winkelmann* et al. 2017; *Heitzig* et al. 2018; *Lade* et al. 2017].

Approaches like this are motivated by a more holistic perspective of the Earth system. Naturally, such endeavours need to include a better understanding of ecosystems on a global scale [*Purves* et al. 2013]. One such attempt is made with the so-called Madingley model presented in [*Harfoot* et al. 2014], which represents ecological dynamics on planetary scale and with high detail, aiming to make quantitative predictions.

In addition, another important aspect is to better understand the interfaces between domains that were previously studied in isolation. For instance, in cases where human exploitation of local ecosystems is the most relevant process, the choice between a sustainable or unsustainable resource use is affected by social dynamics. Thus, depending on the social dynamics, shift to sustainable exploitation strategies may be facilitated [*Barfuss* et al. 2017]. This process may be further influenced by governance structures [*Geier* et al. 2019].

Given the interconnectedness of the Earth system and the non-linear dynamics of its subsystems, an important aspect of current Earth system modelling is the study of tipping elements, past and present [*Brovkin* et al. 2021; *Steffen and Rockström* et al. 2018]. These elements may be in the climate system, but can also be of an ecological nature [*Dakos* et al. 2019] or occur in social dynamics [*Müller* et al. 2021]. While abrupt changes in the Earth system are not particular to the present time (recall the GOE, Section I.2.3), their dynamics have changed with the presence of humankind, with time scales being drastically shorter and the system being influenced by the complexities of human socio-economical dynamics. This further emphasises the need for a holistic investigation of the Earth system.

II.1.2.2 Modelling ecological interactions

Which modelling approaches are there to represent and study ecological interactions? In general, ecosystem models need to find suitable representations for (at least) the following aspects of natural ecosystems: (i) a suitable level of representation for the model entities; (ii) a trait representation, i.e. how the entities' functionality is encoded; (iii) their possible interactions; (iv) a representation of an environment; and (v) evolutionary processes. Each of these categories hosts a wide variety of possible choices, ranging over many spatial, temporal, and conceptual scales. These modelling choices need to be aligned with the question that is to be studied using the model.

Noting the importance of evolutionary processes on the structure of ecological interactions, *Brännström* et al. [2012] isolated a group of approaches that aim to study ecological systems with models that combine ecological dynamics with evolutionary processes. These so-called *community evolution models* are characterised by population dynamics, representing the ecological dynamics, and trait-based interactions that model evolutionary processes in the form of mutating traits. The level of individuality in these models is that of the species; this allows to abstract the behaviour of individual organisms to that of a large population of individuals, hence making population dynamics applicable. Given the aim to also represent evolutionary processes, this is a necessary compromise due to the large time scale on which evolution acts.

Apart from the similarities in assumptions for the representation on the species-level and the use of population dynamics, a wide variety of modelling approaches exists within this group of community evolution models. These differences pertain to the topology of the trait space (discrete or continuous; low- or high-dimensional), the kinds of interactions (trophic, competitive, mutualistic), the foraging behaviour (dynamic or static), the mutation mechanisms (gradual or large-step), and more – for a more detailed comparison of modelling choices, see [*Brännström* et al. 2012, Table 1].

Notably, these models do not aim to represent the Evolutionary Unfolding by successive modularisation, but try to find a heuristic process by which the expected ecosystem structures can arise. Using the picture of the Evolutionary Spiral, the focus of these models lies on the diversification processes; any form of aggregation, transformation, or multi-level representation is abstracted away.

Niche-based food web models

Among the community evolution models reviewed in [*Brännström* et al. 2012] I want to highlight one approach which not only led to a large number of adaptations and extensions but which formed the basis for the model that I will use throughout this part of the thesis. The model in question is that by *Loeuille and Loreau* [2005] and is often referred to as an *evolutionary niche model*.

The central idea of these models is to model trophic interactions by locating species' and their feeding preferences in a one-dimensional niche space. Each species is then characterised by two properties: its own feeding preferences (a volume of niche space) and the position in niche space at which other species perceive it (and subsequently feed on it). In the model by *Loeuille and Loreau* [2005], the one-dimensional niche axis represents the mean size of individuals of a species; this is an appealing choice, as it can be easily related to actual species properties and their feeding preferences. A typical assumption is that species feed primarily on smaller species, thus introducing a size-ordering between predators and prey.

The system becomes evolutionary by randomly or periodically adding new species to it which have slightly mutated traits, e.g. a larger size or a shifted niche. Additionally, more traits can be introduced that influence the width of the niche, thus distinguishing generalist species, which may feed on prey of a wide variety of sizes, and specialist species, which require very specific prey sizes.

These dynamics typically lead to multi-layered food webs emerging in these models: species evolve to larger sizes, eventually being able to feed on a cluster of species at a smaller size. Each feeding relation is subject to biomass conversion processes, which have a non-optimal efficiency. Subsequently, the depth of the food webs is then limited by the amount of energy that is available to the species at the end of the food chain.

The idea of representing trophic interactions this way was not new at the time – in fact, it goes back to the 1970s and the consumer-resource model by *MacArthur* [1972] (but see [*Ackermann and Doebeli* 2004] for a more recent description). However, the novel idea of *Loeuille and Loreau* [2005] was to escalate this approach such that consumers can become the resource for other consumers. In addition, they included further ecologically motivated mechanisms like optimal foraging theory and allometric scaling[37] to simplify the dynamics and make them more comparable to real food webs.

The model by *Allhoff and Ritterskamp* et al. [*Allhoff and Ritterskamp* et al. 2015] that I will use in the rest of this thesis builds on these ideas – it is described in more detail in Chapter II.2. For a list and comparison of further models see [*Fritsch* et al. 2021]; the authors also propose an encompassing framework of niche models. Also note [*Girardot* et al. 2020], which addresses some of the questions that are also of interest in this thesis.

---

[37] The observation that physiological properties of organisms appear to be related to each other by power laws. For instance, the metabolic rate of an organism of mass m scales with roughly $m^{3/4}$ for most animals [*Brown* et al. 2004]. Note that this is a heuristic description: the exact value of the exponent lacks a theoretical explanation.

II.1.2.3 Ecosystem resilience

The term *resilience* was already mentioned a few times in the previous sections, yet lacked a proper definition. While commonly used, it is often interpreted in widely different ways, thereby making it more difficult to compare research investigating this topic. A precise formulation of *resilience* is needed when studying questions revolving around the resilience of ecosystems (and other systems), not only to sharpen one's own view of the term but also to communicate more clearly which aspects of the system are studied.

To that end, I will follow a framework proposed by *Tamberg* et al. [2020] that aims to sharpen the use of the term and the surrounding research questions. The authors present a conceptual framework of how to communicate research questions and models surrounding the study of resilience of socio-technical-environmental systems. They present a checklist of questions, which may serve as a guideline for more precise language about models investigating resilience:

1. Resilience of what? – The *system*
2. Resilience regarding what? – The *sustainant*
3. Resilience against what? – The *adverse influence*
4. Resilience how? – The *response options*

Here, the term *sustainant* is a neologism that denotes 'the feature of a system that should be maintained in order to call the system resilient', which may for instance refer to the system's state, its structure, or some function it fulfils. *Tamberg* et al. reflect on the precision of this approach and weigh it against the conceptual flexibility needed when discussing and comparing different studies. They suggest to use the checklist for refining both model formulations and research questions, which in turn determine the components of the checklist; through an iterative approach, all these involved aspects become further refined.

In the following, I will go through the guideline proposed in [*Tamberg* et al. 2020] and apply it to ecosystems *in general.* At this point, these can only be general considerations, because many answers will depend on the actual modelling approaches and the particular research question and can thus only be addressed once the model was formulated.

What is the system? The ecosystem itself can be regarded as the system of interest in this context: the ecosystem's resilience is the subject of this investigation. What consistutes the system of interest can and needs to be specified further depending on the choice of representation of the ecosystem in the model, mainly: its substructure (entities like species and resources), the interactions between these, and the environment. In models where human societal aspects play a role, processes like agency would also need proper definition.

The authors note that many choices regarding the system are already decided when the model representation is designed. These aspects are typically not informed by the questions around resilience. Nevertheless, it can be seen as useful to precisely formulate system constituents.

What is the sustainant? *Tamberg* et al. describe the sustainant as follows:

> Which feature or property of the system is supposed to be sustained or maintained in order to call the system resilient? Its state or structure, its pathway? Some long-term equilibrium? Its function, purpose, or utility for some stakeholder? Some quantitative or qualitative aspect of the system? What this 'sustainant' is is no objective feature of a system but is normatively chosen by the observer from their perspective, which should be clearly communicated. Especially what the 'function' or 'purpose' of a system is can be seen differently from different perspectives. *Tamberg* et al. 2020, Section 2.2

Furthermore, they stress:

> Considering the normativity of the sustainant, it becomes clear that its choice can be subject to power relations, inequality, and competing interests. Many authors therefore demand to consider the question 'Resilience for whom?' to account for these aspects. This question is located on a meta level above that of the checklist. It can help to both choose the sustainant and to criticise this choice, for example from an inequality perspective. *Tamberg* et al. 2020, Section 2.2

Questions to ask and investigate thus relate to threshold values, acceptable ranges of key observables, or their acceptable recovery times, but have to be seen as inherently normative. The study of ecosystem resilience is not exempt from such normative choices, which makes the question of the sustainant a more difficult one and may in some scenarios call for a wider perspective – especially in the context of socio-ecological systems, these questions become pressing.

In more abstract models, pragmatic choices for sustainants may suffice. Examples could be the biodiversity in a model or some other measure of ecosystem functioning that allows quantifying whether it can be held in an acceptable range or in which time it may recover back into that range.

What is the adverse influence? The adverse influence describes any effect that acts on the system adverse to the sustainant – the term alone is not meant to convey whether the influence is something undesirable or not, that aspect is evaluated depending on the choice of sustainant. The adverse influence may be in the form of a perturbation or a temporary shock, but could also be a permanent change to some system property (on the studied time scales). There may also be several adverse influences affecting the system and their influence on the sustainant may be direct or indirect.

As an example from the context of ecosystems, fast changes could be in the form of volcanic eruptions or fluctuating solar power. Depending on the time scales of interest, the anthropogenic rise in greenhouse gases would be classified as fast or slow.

What are the system's response options? The response options refer to the mechanisms that determine the system's response to the adverse influence. These determine in how far the sustainant will be affected and, if so, how it *Tamberg* et al. note the following questions for better specifying the response options: 'At which levels can or does a system react to adverse influences? [...] What is the range of possible reactions? Which reactions are endogenous as a consequence of the systems structure and rules?'

I would make the additional remark that a system's response options are a consequence of its capabilities. While this is not a limiting consideration in highly diverse natural systems, it is of substantial importance when investigating ecosystems (or other systems) using models. Models necessarily represent only a subset of the involved entities and processes and have to use drastically simplified representations of both. To remain understandable, reductionist approaches are often successful, which include only the processes that are deemed most relevant, inadvertently excluding relevant response options. In effect, response options in a model are severely limited, are influenced by the focus set in the formulation of the model, and may show vastly different characteristics than those observed in natural systems.

It is not a problem per se that models have limited response options: If the primary response of the to-be-represented system is well-known and a good model representation of the mechanism can be found, it can be considered a sufficiently good approximation of the system. However, if a system's response is an interplay of many processes where the recovery of a sustainant is an emergent property, reductionist models may not capture this property of the system. This needs to be taken into account when evaluating a system's response options.

II.2

Model Description

In this chapter, I motivate my choice of a model to study the presented research questions and then describe the model mechanisms as well as its implementation.

II.2.1 Choice of model

My studies of ecosystem evolution in the context of this thesis are based on an evolutionary food web model proposed by *Allhoff and Ritterskamp* et al. [2015]. The model can be categorised as a niche model (see Section II.1.2.2), with species being defined primarily by their feeding interactions. Unlike many other food web models, the evolution mechanism in this model leads to a continuous species turnover, such that the food web structure does not freeze out but new species cause a constant restructuring of the food web.

Based on this model, Korinna T. *Allhoff* and co-authors studied questions of biodiversity and ecosystem functioning. The authors found that the resulting interaction networks more closely resemble those of natural food webs than previous niche models [*Allhoff and Ritterskamp* et al. 2015]. Furthermore, they isolate a set of measures for quantifying biodiversity and ecosystem functioning in the presented model [*Allhoff and Drossel* 2016], which are insightful for characterising the state of the food web and its reaction to changing environmental conditions.

As reviewed in Section II.1.2.2, there are many approaches for modelling ecosystems. Why use specifically this model to study the previously posed research questions?

Given that the research questions revolve around food web resilience, which includes their recovery after perturbations, the model of choice would need to operate on evolutionary time scales and with a suitable evolution mechanism; this was the first criterium for the model. Furthermore, to remain computationally feasible in such a situation, a sensible level of abstraction needs to be chosen. In this case, representing the ecosystem on the level of individuals would not allow studying food web resilience on evolutionary time scales; the level of representation would need to be that of the species instead. In addition, the model should ideally be easy to handle and be sufficiently studied to know that it behaves robustly.

From the range of niche models, many of which are derivations or adaptations of the one presented in [*Loeuille and Loreau* 2005], the model by *Allhoff and Ritter-*

skamp et al. [2015] has a number of favourable properties which match these criteria: (i) it is already well studied with respect to ecosystem functioning – no need to re-invent the wheel; (ii) it features a continuous species turnover, which makes it more akin to the continuously changing evolutionary systems observed in nature; (iii) its formulation is comparably simple, elegant, and can be adapted to other scenarios; and (iv) it uses a logarithmic representation of the feeding dimension, alleviating many of the implementational difficulties of niche models that are restricted to a finite resource domain. These were the main considerations for selecting this particular model for studying the presented research questions.

And why not design a food web model from scratch?

Designing new models appears to be a frequently-taken approach in the study of food webs, leading to a wide variety of published models, few of which are re-used in other studies. While my initial approach was to do the same, the resulting model was ill-suited to address the topic of food web resilience. In contrast, the criteria listed above suggested that the model by *Allhoff and Ritterskamp* et al. would be a suitable basis for studying these questions.

II.2.2 The ECST model

The *evolutionary food web model with continuous species turnover* introduced by *Allhoff and Ritterskamp* et al. [2015] – from now on referred to as ECST model – is a niche model with species traits and allometrically motivated species properties. Its dynamics can be subdivided into two parts, one representing the ecosystem's ecological dynamics and the other introducing new species into the food web, thus bringing variation into the model and making it an evolutionary model.

The following sections are a description of the ECST model. In the notation, I mostly follow [*Allhoff and Ritterskamp* et al. 2015] and add additional symbols where necessary.

II.2.2.1 Ecological dynamics

The ECST model distinguishes two kinds of entities: resources and species. Conceptually, these differ only in resource entities receiving external energy input while the species can extract energy only from the resources or other species.

The abundance of resources and species is measured in their *biomass density* B_i over the model domain. Additionally, the entities are characterised by their specific *body mass* m_i, which allows to compute a *number density* $n_i = B_i/m_i$ which is a measure for the number of individuals in the model domain. In the context of this model, biomass is the universal currency and can be abstracted to a form of energy that is usable by all represented entities of the model.

As in other niche models, the ECST model represents species and resources as interacting exclusively via directed trophic interactions which transfer biomass from one entity to another; throughout this process, some biomass is lost in the biomass conversion process. Subsequently, each entity has additional traits describing these feeding relations: the *feeding centre* f_i and the *niche width* s_i, together referred to as their *feeding kernel*. The feeding kernel of a species determines which species it may feed on; its body mass determines which other species may prey on it.

Population dynamics

The population dynamics are described by a system of ordinary differential equations, consisting of one set of equations for resource entities (Equation (II.2.1), with resources denoted by index r) and one for species (Equation (II.2.2)).

Resources The dynamics of resource entities in this model are described by an influx and an outflux term:

$$\dot{B}_r = G_r B_r - \sum_j g_{jr} B_j \quad . \tag{II.2.1}$$

Here, the additive term describes replenishment of the resource entities, where $G_r = R_r(1 - B_r/K_r)$ is the *replenishment rate* of the resource, parametrised by the *carrying capacity* K_r and intrinsic growth rate R_r. The replenishment rate takes the form of a logistic growth equation. The resource entities are referred to as *external resources* because they are external to the set of species.

The subtractive term represents consumption by the species of the food web and includes the *functional response* matrix g_{jr}, the biomass-density-specific rate with which species j consumes resource r. Which processes are modelled by the functional response will become clearer below.

Species For species entities, the consumption term includes a mirrored counterpart that represents the influx from consuming other entities. Additionally, species have to exert energy to stay alive, represented as an additional outflux term:

$$\dot{B}_i = \sum_j \lambda_j g_{ij} B_i - \sum_j g_{ji} B_j - x_i B_i \tag{II.2.2}$$

Here, λ_j denotes the *biomass conversion efficiency* with which biomass foraged from species j can be converted into biomass of species i. The last term denotes intrinsic losses, or: a death rate. It includes the mass-specific *metabolic rate* x_i, also called *respiration rate*, which follows from allometric considerations as

$$x_i = 0.314 \cdot m_i^{-1/4} \quad . \tag{II.2.3}$$

According to [*Allhoff* 2015, Table 4.1], the constant is from [*Yodzis and Innes* 1992], where metabolic rates are derived from the study of certain organisms' energetics in combination with allometric scaling dependencies. As the exact value of the constant is of no particular importance in the context of this thesis, I decided to use the same value as in the original formulation of the ECST model for better comparability.

Functional response

The crucial part of Equations (II.2.1) and (II.2.2) is the functional response matrix g, which contains information about all species interactions. It is computed as follows:

$$g_{ij} = \frac{a_{ij} B_j}{m_i} \left[1 + \sum_k h_i a_{ik} B_k + \sum_l c_{il} B_l \right]^{-1} \tag{II.2.4}$$

The summation terms in the denominator refer to the attack rate on prey species and the competition with other predators for those prey, respectively. This type of functional response is characterised as BeddingtonDeAngelis functional response, which saturates for large prey populations or high competition. It is similar to a Holling type II response, which is also saturating but does not include a competition term in the denominator.

Handling time The *handling time* h_i is representing the time needed to locate a prey and consume it, motivated by the presumed behaviour of predators according to optimal foraging theory:

$$h_i = 0.398 \cdot m_i^{-3/4} \tag{II.2.5}$$

Again, the scaling is allometrically motivated and the constant is taken from [*Allhoff* 2015] for easier comparability.

Attack rate The so-called *attack rate* a_{ij} describes the rate with which predator species i consumes a prey species or resource j. It is computed by evaluating the gaussian feeding kernel at the position of the prey entity's specific body mass m_j:

$$a_{ij} = a_i \mathcal{N}_{ij} = m_i^{3/4} \frac{1}{\sqrt{2\pi s_i^2}} \exp\left[\frac{(\log_{10} f_i - \log_{10} m_j)^2}{s_i^2} \right] \tag{II.2.6}$$

The species-specific *attack rate factor* $a_i = m_i^{3/4}$ again follows from allometric considerations, denoting an attack rate that grows increasingly slowly with larger predator body mass.

Competition Competition is modelled as *interference competition*, where similar predator species that are competing for a common prey interfere with each other in a way that *both* their consumption rates are reduced [*Zhang* et al. 2015]. In the ECST model, similarity of two species can be conveniently defined because the feeding relations are fully specified via the feeding kernels of the competing prey species. By that measure, fully overlapping feeding kernels should result in maximum competition, whereas far-apart kernels should lead to negligible interference and thus a negligible effect of the competition mechanism.

The *feeding kernel overlap* $I_{ij,k}$ of two predator species i and j competing for prey species k is defined as the integral of the product of the involved feeding kernels. As both feeding kernels are gaussian, their product is again a gaussian (**??**) with mean and variance being a function of the feeding kernels of the predators; the prey properties no longer play a role:

$$
\begin{aligned}
I_{ij,k} &= \int \mathcal{N}_{ik}\,\mathcal{N}_{jk}\,\mathrm{d}(\log_{10} m) && \text{(II.2.7)}\\
&= \ldots = \frac{1}{\sqrt{2\pi(s_i^2 + s_j^2)}}\ \exp\left[-\frac{(\log_{10} f_i - \log_{10} f_j)^2}{2(s_i^2 + s_j^2)}\right] && \text{(II.2.8)}\\
&= I_{ij} = I_{ji} && \text{(II.2.9)}
\end{aligned}
$$

Using this overlap as a similarity measure, the competition $c_{ij,k}$ between two predator species i and j for a prey species k is also independent of the prey's properties, and is thus denoted solely as c_{ij}. With a normalisation by the self-overlap and a special case for self-competition, it is defined as follows:

$$
c_{ij,k} = c_{ij} = \begin{cases} c_{\text{food}} \frac{I_{ij}}{I_{ii}} & \text{for } i \neq j \\ c_{\text{food}} + c_{\text{intra}} & \text{else} \end{cases} \qquad \text{(II.2.10)}
$$

Here, $I_i i \propto s_i^{-1}$ is the self-overlap, which is higher for species with a lower niche width, such that those predators that have a higher specificity in their feeding range are affected more strongly by competing species. Subsequently, I_{ij} is symmetric, but c_{ij} is not.

Vector notation

Equation (II.2.1) can also be written as a single equation using the biomass density vector $\mathbf{B}$ and the functional response matrix g:

$$
\dot{\mathbf{B}} = \mathbf{G} \odot \mathbf{B} + (\mathrm{g}\lambda - \mathbf{x}) \odot \mathbf{B} - (\mathbf{B}^\intercal \mathrm{g})^\intercal \qquad \text{(II.2.11)}
$$

The $\cdot \odot \cdot$ operation denotes the Hadamard product, i.e. element-wise multiplication. Note that only those elements of $\mathbf{G}$ that refer to external resources are non-zero, such that the first term only applies to resources' biomass.

This vector representation has the benefit of being closer to the data structures used in the implementation of this model, where vector-based data structures allow for higher computational efficiency (see Section II.2.3.2).

II.2.2.2 Evolution mechanism

The population dynamics mechanisms described above represents the interaction of species in a *specific* configuration of the interaction network as given by the species traits. By adding a mutation mechanism to the model, the system becomes adaptive and the system becomes an evolutionary system.

In the ECST model, there are two mechanisms by which the species composition in the system can change:

- **An extinction event**: A species' biomass density drops below a species-specific threshold, such that the species is considered extinct and subsequently removed from the system.
- **A speciation event:** A mutation occurs, introducing a new species into the system as offspring of an existing species.

As such, the extinction event models the disappearance of a single species (not like the large-scale extinction events discussed throughout EM). The speciation event, in turn, is a heuristic representation of the process by which a population of individuals within a species bring forward sufficient variation to be considered a new species; these internal dynamics are not represented in the ECST model.

Extinction events

The *extinction threshold* $n^{\min}$ is defined as a number density, i.e. a representation of the number of individuals of a species in the model domain. It translates into a species-specific minimal biomass density, depending on the body mass:

$$B_i^{\min} = n^{\min}\, m_i \tag{II.2.12}$$

Once a species' biomass is below that threshold, it will be removed, thereby modelling an extinction event. Because the threshold is scaling with m_i, species with larger body mass need to acquire more biomass in order to not become extinct.

In [*Allhoff and Ritterskamp* et al. 2015], the extinction threshold is chosen as $n^{\min} = 2 \cdot 10^{-4}$. If not noted otherwise, the same value is used throughout this thesis.

Speciation events

Through speciation, a new species is introduced into the system. In the ECST model, such events occur periodically every $T^{\text{mutate}} = 100$ iteration steps.

Traits of new species The new species' traits depend on that of a randomly chosen parent species p, with trait values of the offspring species o mutated as follows:

$$\begin{aligned}\log_{10}(m_o) &\sim \mathcal{U}\left[\log_{10}(\frac{1}{2}m_p),\ \log_{10}(2m_p)\right]\\ \log_{10}(f_o) &\sim \mathcal{U}\left[\log_{10}(m_o)-3,\ \log_{10}(m_o)-\frac{1}{2}\right]\\ s_o &\sim \mathcal{U}\left[0.5,\ 1.5\right]\end{aligned} \tag{II.2.13}$$

where $\cdot \sim \mathcal{U}$ denotes drawing a random value from a uniform distribution with the given bounds. Notably, only the offspring's body mass m_o depends directly on traits of the parent; the feeding centre f_o is then 3 to 1000 times smaller than m_o and the niche width s_o is selected fully independently from the parent's traits.

The offspring species with the above trait values is then added to the ODE system with an initial biomass density of $B_o = B_o^{\min}$: directly at the extinction threshold. To maintain biomass conservation throughout the system, the biomass density of the parent is reduced accordingly. If this would require more biomass than available in the parent species ($B_p < B_o^{\min}$), the offspring species will *not* be added to the system and the parent's biomass will remain as it was – effectively, no mutation takes place in such a case.

Size of mutations In general, the ECST model assumes comparably large mutations: As can be seen in Equation (II.2.13), only the body mass trait m is directly inherited from the parent and allows a body mass difference of factor 2 between parent and offspring. In particular, the f and s traits are being drawn from the *whole* range of possible values, decoupling strongly (in the case of f) or completely (for s) from the parent species' properties.

Owing to the large mutations and f and s mutating mostly independently from the parent's traits, the shared properties between between parent and offspring are very low. Subsequently, these 'speciation events' in the ECST model are more similar to what is typically referred to as an *invasion*, i.e. a largely unrelated species becoming part of the system. For the sake of consistency with the original literature, these events will still be referred to as 'speciations' or 'mutations' throughout the thesis; this also frees up the term 'invasion' for a dedicated study of species invasion.

An alternative mode of drawing offspring values (based on normal distributions centred around parent traits) is explained and studied in Section II.3.4.1.

Timing of speciation events As mentioned, the ECST model introduces species periodically, which may appear as a departure from the inherently stochastic way in which species appear in real ecosystems. The motivation behind periodic speciation events is that the ecological dynamics have time to equilibrate, representing the separation of time scales between population dynamics and evolutionary processes.

While I did look into an adapted speciation mechanism with randomly occurring events, the qualitative behaviour was overall equivalent. As the analysis of the ecosystem effects caused by individual speciation events became more difficult, I chose to focus on the version with periodically appearing new species.

II.2.2.3 Model initialisation and iteration

Typically, the model is initialised with a single resource ($m = 1$, $B = 100$, $R = 1$) and a single species ($m = 100$, $f = 1$, $s = 1$, $B = B^{\mathrm{min}}$). Notably, the initial species' feeding kernel is perfectly aligned with the external resource and can be considered a generalist consumer due to the comparably large niche width. This simple primordial scenario is sufficient to generate a structured food web via mutations.

In some cases, the initial state will not be this simple situation, but a manually constructed one. Alternatively, it may be loaded from a previous simulation (see Section II.2.3 for more on this), in which case all species traits and biomass densities are used for the setup of the system.

The initial state, after setup finished, is denoted as $t = 0$. After that, the model's iteration step is performed a specified number of times, incrementing the time after each step; this is repeated until $t = T^{\mathrm{max}}$.

The iteration step itself consists of the following procedure:

1. Check for extinctions and remove extinct species
2. (Potentially) add mutant species
3. Integrate population dynamics

In the implementation by *Allhoff and Ritterskamp* et al. [2015], each integration covers the time frame from one mutation to the next one, basically equating one model iteration step with the population dynamics between mutation events. In their formulation, the time between two mutation events is denoted as '10 000 steps'. Here, this time period is split into 100 iteration steps in which the above-mentioned procedure is carried out. Subsequently, to arrive at the same time scale as they use, each iteration step performs an integration of the population dynamics over 100 time units. While this approach introduces an additional 'internal' time scale, it allows for a larger flexibility in terms of data output or other model mechanisms acting at varying times between two mutation events. Throughout this thesis, t refers to the number of *iteration* steps, such that the time between mutations is $T^{\mathrm{mutate}} = 100$ steps.

The time values used in [*Allhoff and Ritterskamp* et al. 2015] are all on a time scale equivalent to the internal time scale used here, making all times shown there larger by a factor of 100 when compared to the values here; this is merely a difference in representation.

II.2.2.4 Summary and overview of parameters

In summary, the ECST model can be characterised as follows:

- Species are defined by three traits: body mass m, feeding centre f and niche width s.
- The food web is provided with energy via an external resource that is replenished to a certain carrying capacity K.
- Trophic interactions are located on a resource axis; the m and f traits refer to positions on that axis. The species traits' describe possible feeding interactions.
- Allometric scaling is used for species-specific parameters; accordingly, interactions are defined in logarithmic space along the resource axis.
- The functional response accounts for interference competition.
- Population dynamics describe the time development of species' biomass density and include a species-specific respiration rate x_i. If a species drops below a certain number density, it is regarded as extinct and removed from the system.
- The evolution mechanism randomly selects a parent species and generates a small offspring species with mutated traits. While the offspring body mass m is related to that of the parent, the remaining traits are largely independent from the parent's traits.

Table II.2.1 gives an overview of symbols and default parameter values. If not specified otherwise, the values given there are used throughout the rest of this thesis.

Table II.2.1 Parameters and notation used in the ECST model

Symbol	Default	Name / Description	Comment
G_r		resource replenishment	logistic growth, Equation (II.2.1)
R_r	1	intrinsic replenishment rate	
K_r	100	carrying capacity of resource r	
λ_r	0.45	biomass conversion efficiency	for predation on resource r
m_i		body mass trait	determines species properties and feeding relations
f_i		feeding centre trait	position of feeding kernel
s_i		niche width trait	standard deviation of feeding kernel (on log scale)
λ_i	0.85	biomass conversion efficiency	for predation on prey species i
x_i		respiration rate	$x_i \propto m_i^{-1/4}$, Equation (II.2.3)
h_i		handling time of predator i	$h_i \propto m_i^{-3/4}$, Equation (II.2.5)
a_i		attack rate factor of predator i	$a_i \propto m_i^{+3/4}$
B_i		biomass density	
n_i		number density	$n_i = B_i / m_i$
g_{ij}		functional response	Equation (II.2.4)
a_{ij}		attack rate of predator i on prey j	Equation (II.2.6)
c_{ij}		competition b/w species i and j	Equation (II.2.10)
c_{food}	0.8	food competition	used in c_{ij}
c_{intra}	1.2	intra-specific competition	used in $c_{ii} = c_{\text{food}} + c_{\text{intra}}$
n^{min}	$2 \cdot 10^{-4}$	extinction threshold	equal to the initial number density after speciation events
T^{mutate}	100 steps	time between speciation events	

II.2.3 Model implementation

The ECST model was implemented and evaluated using Utopia, a modelling framework for complex and evolving systems that was jointly developed in the group of Kurt Roth starting in 2018. As I have co-developed this framework during my time in the group, I will briefly present key ideas and features of Utopia itself before turning to the model implementation itself.

II.2.3.1 The *Utopia* modelling framework

The Utopia framework aims to be a comprehensive tool that can be used throughout the whole workflow of modelling-based research: (i) providing a conceptual language to assist in designing and communicating about a model, (ii) efficiently implementing the model, (iii) performing model simulations, and (iv) evaluating the simulation data in a flexible and reproducible way. To that end, Utopia consists of a backend based on modern C++ and focussing on efficient model implementations, and a Python-based frontend from which simulations can be configured, run, and evaluated. It relies heavily on YAML as a versatile and adaptable configuration language [*Ben-Kiki* et al. 2009].

The backend makes use of modern C++ to provide not only a capable and extendable simulation infrastructure, but also a library to efficiently store data in the HDF5 format, and a wide range of optimised and tested modelling tools for implementing cellular automata, agent-based models, or graph-based models. Via a `Model` base class, simulation infrastructure like configuration files and output groups are conveniently accessible, simplifying model implementation to specialising the methods of a derived class.

Furthermore, Utopia was designed with hierarchical structures in mind: All data structures in the backend are built to support nesting of models, such that one model can include other models and couple its dynamics to it. While this pertains mostly to the infrastructure and not to the (open and difficult) problems of representing a wide range of scales and processes, simplifying the technical aspects of model coupling proved to be a valuable feature for a modelling framework.

The frontend aims for convenient configuration of simulation runs and evaluations. For instance, this is achieved by simplifying the definition of parameter sweeps using the `paramspace` package [*Sevinchan* 2020], which allows to turn an arbitrary parameter in the model configuration into a sweep dimension. In addition, parameter sweeps can be trivially parallelised, further simplifying sensitivity analysis of models.

For data analysis, Utopia integrates with the dantro package, which provides a configuration-based data processing pipeline for the simulation data; this is described in a bit more detail below.

For more information, refer to [*Riedel* et al. 2020] or https://utopia-project.org/. Regarding the experiences we made in using this framework on a larger scale within the working group, including more than 20 individual projects and teaching multiple M.Sc.-level seminars, see [*Sevinchan and Herdeanu and Mack* et al. 2020].

II.2.3.2 Implementation details

The model was implemented using the definitions given above, which in turn are based on the formulation by *Allhoff and Ritterskamp* et al. [*Allhoff and Ritterskamp* et al. 2015] and personal correspondence with her. Technical details differ between her implementation and the one presented here. In the following I will provide the most relevant details of my implementation[38].

Representation of ODE system At its core, the ECST model requires the iteration of the population dynamics, which are also the costliest parts of the simulation. The implementation makes use of the possibility of representing the ODE system in vectorial form, see Equation (II.2.11), which not only simplifies the data structures but also allows efficient computation. The well-established linear algebra library `Armadillo` [*Sanderson and Curtin* 2016], which is already provided as part of UTOPIA, is used for these data structures and corresponding computations.

Subsequently, quantities like the biomass densities $\mathbf{B}$ or the functional response matrix g can conveniently be represented as vectors or matrices. These data structures are also used for all other species-specific properties, like trait vectors $(\mathbf{m}, \mathbf{f}, \mathbf{s})$.

The implementation also distinguishes between species by an ID. For technical reasons, ID values start at 1, which is typically assigned to the external resource.

Integration of ODE system The integration of the ODE system employs the `Boost.Numeric.Odeint` library [*Boost Developers* 2021] and its implementation of the explicit Runge-Kutta Cash-Karp method. Integration uses an adaptive time step with $\varepsilon^{\mathrm{abs}} = 10^{-7}$ and $\varepsilon^{\mathrm{rel}} = 10^{-6}$.

Unlike the implementation by *Allhoff*, the population dynamics from the addition of one species to the next are not integrated contiguously, but split into 100 UTOPIA iteration steps. This proved to be a good compromise between computational efficiency (which is highest when letting the ODE integrator do all the work) and flexibility in data writing or attaching other model mechanisms between the population dynamics (which is highest when doing many UTOPIA iterations).

In effect, the quantitative values for iteration time differ by a factor of 100 between this thesis and figures shown in [*Allhoff and Ritterskamp* et al. 2015]: Where the implementation by *Allhoff and Ritterskamp* et al. has 10 000 steps between two mutation events, the UTOPIA-based implementation has only 100 steps. Other than that, this 'internal' time scale used in the implementation does not play a role anywhere in this thesis.

---

[38] The implementation's source code and evaluation routines are available upon request.

Random number generation For modelling random processes in computer models, so-called *pseudo-random number generators* (PRNG) are used. These can vary greatly in their statistical properties and speed, so care has to be taken to choose a generator with suitable properties.

My model implementation makes use of the `Xoroshiro128**` generator, which not only passes all established statistical tests for PRNGs but is also very fast and has a small memory footprint[39]. The C++ implementation of the PRNG itself was carried out and tested by Harald Mack.

The PRNG is set up and managed by UTOPIA and shared throughout all data structures of an individual model run. This allows for repetitions of simulations, as there is no other source for pseudo-randomness. The initial PRNG state is controlled via a `seed` parameter; if desiring to perform multiple simulations with different random number sequences, a parameter sweep over the `seed` is performed.

Configuration The system is initialised with information from the configuration that UTOPIA passes on to the model. All parameters can be specified via that configuration, including the number of species and resources[40] and their properties.

Furthermore, system parameters can be dynamically changed during the simulation. To that end, the model implementation defines a set of ‘triggers’ and ‘actions’, which can be controlled via the configuration and conveniently change a model value at some time in to the simulation or following some other condition.

Re-loading system states Due to long spin-up times of the model and overall long simulation times, being able to restart a simulation from a previous state is a valuable feature. For the ECST model, the system state is fully defined by the biomass densities $\mathbf{B}$, the traits, and some accompanying information (resource parameters, species IDs). By storing these values at the end of a simulation, the food web can be re-loaded and the simulation continued from that point.

There is one caveat: the PRNG’s state is *not* stored. While a loaded simulation will be deterministic, depending on the `seed`, it will not be identical to the direct continuation of the simulation.

II.2.3.3 Data output

The UTOPIA framework provides an interface for conveniently writing data in the HDF5 format. To let data be self-descriptive, HDF5 metadata attributes are used, the content of which are parsed during data evaluation and turned into labelled multi-dimensional arrays. These labelled arrays contain coordinate information (time, species ID, etc.) and allow for more reliable evaluation routines.

---

[39] See https://prng.di.unimi.it.

[40] In principle, the implementation allows for multiple resources with different properties. However, this scenario was not studied in more detail and is not part of this thesis.

I further expanded Utopia's data writing capabilities with three custom data structures: First, an alternative `DataManager` which is optimised for re-using shared coordinate data and reducing the number of datasets needed to represent variable-length data.

Second, a `PropertyTracker` data structure which collects species-specific properties. These trackers first aggregate this information and then write it to the respective HDF5 dataset upon the next write operation. They ensure that species properties can be retained and stored regardless of the timing of the `DataManager`.

Finally, as another option for storing simulation data, I implemented a so-called `SnapshotManager` which keeps track of system states at different times of the simulation. Snapshots can be compared according to arbitrary conditions, which allows to dynamically collect them during model iteration but only keep those that are of interest (according to the condition). These snapshots can then be used for re-loading a system state and continuing a simulation from that point forward.

Timing Data is stored after the iteration step (see Section II.2.2.3) was finished and the time incremented. That way, species that newly dropped under the extinction threshold may still be part of the stored data, as they would only be removed at the beginning of the next step.

Species-specific data is not stored at a specific time, but first aggregated whenever it appears (e.g. upon creation of a new species or its extinction). Upon the next regular write operation, it is written into a dataset. If there are no write operations left, it will be stored at the very end of the simulation.

Stored data A wide range of data can be stored throughout a simulation run; which data is stored and how frequently is controlled via the custom `DataManager`. To give a brief overview of stored data:

- Time-resolved species-specific data (biomass densities, resulting interaction strengths, . . .)
- Time-resolved macroscopic observables (overall biomass and biomass flow, interaction structure, . . .)
- Time-resolved species- or biomass-based statistics (biomass distribution, trophic level distribution, . . .)
- Extracted network structure in a reconstructible format and with associated node and edge properties
- Species-specific properties (traits, addition and removal times, life time, observables quantifying their effect on the system, . . .)
- Snapshots of system states at different times or for certain conditions
- Time-resolved benchmarking data (for optimisation purposes)

II.2.3.4 Simulation and data evaluation

Simulations were typically carried out on a local machine. Larger simulations were performed on the high-performance cluster 'BwForCluster MLS & WISO' using Utopia's cluster mode feature.

For data evaluation, Utopia couples to the dantro package, which provides a configuration-based data processing and visualisation pipeline. It was developed by Benjamin Herdeanu, Jeremias Traub, and me for use in the Utopia framework, but is an independent Python package [*Sevinchan and Herdeanu and Traub* 2020]. With the pipelining approach of dantro, the process of model evaluation not only becomes faster, but also more flexible and reproducible[41].

Specifically, dantro is optimised for working with hierarchically structured high-dimensional data, as is often the result of complex model simulations with parameter sweeps. To that end, it integrates the `xarray` [*Hoyer and Hamman* 2017] and `h5py` [*Collette* 2013] packages to turn the stored HDF5 data into labelled multi-dimensional `numpy` arrays [*Harris* et al. 2020; *van der Walt* et al. 2011].

In addition, the data evaluation routine needs to be able to handle large amounts of data. As this frequently complicates data evaluation, especially if long computations are required, dantro includes a data transformation framework which represents data transformations as nodes of a directed acyclic graph (DAG) and allows to optimise computations. One such approach is to intelligently cache computation results, thus avoiding unnecessary re-computation.

Transformations can be defined directly via configuration files and alongside configuration-based plot definitions. By internally separating the processes of data transformation and visualisation, evaluating simulations with dantro is very flexible, allowing for declarative definitions of visualisations. Plots are then generated by wrappers around the widely-used `matplotlib` library [*Hunter* 2007].

The evaluation of the ECST model relied heavily on these features provided by dantro, especially when needing to analyse the simulation data from large parameter sweeps. By providing these tools and allowing to automate many parts of the analysis, dantro simplified working with and studying the ECST model and its extensions.

---

[41] For more information, see https://pypi.org/project/dantro/.

II.3

Food Web Evolution

In this section, I describe the basic behaviour of the ECST model with the aim of illustrating how food web evolution occurs in this model. These are meant to introduce the model and its behaviour, widening the perspective given in [*Allhoff and Drossel* 2016; *Allhoff and Ritterskamp* et al. 2015].

I will start with a closer look at the model dynamics on short time scales (1 k steps), starting from a simple situation and introducing relevant observables and visual representations along the way. I will then illustrate how the competition mechanism in the ECST model allows coexistence, and will use artificial food webs to exemplify the population dynamics in a controlled setting. Finally, I turn to the long-term evolution (1 M steps) of food webs and introduce a range of food web measures and evaluation approaches that allow to characterise food web states.

II.3.1 Primordial situation

Food web evolution in the ECST model starts with a single species of mass $m = 100$ feeding on the external resource ($m = 1$) with a niche width of $s = 1$. The ODE system is then integrated according to Equation (II.2.11), leading to changes in the species' biomass density.

In regular intervals, a new species is added via the mechanism described in Section II.2.2.2, briefly: a species is selected at random and a minute amount of biomass is transferred from that species to a newly introduced offspring species. The mass m of the offspring is mutated from that of the parent species; the feeding kernel (traits f and s) are drawn from uniform distributions, dependent only on the offspring's body mass and specifically not dependent on the feeding kernel of the parent species.

Figure II.3.1 shows the biomass density development of this initial setup, starting with a single species and subsequently adding mutant species every 100 steps. Upon each addition, the offspring species' biomass density B rapidly increases from the initial $n^{\min}$ to a value several magnitudes larger. Notable exceptions are the lack of new species at times 200 and 500: these offspring species' traits rendered them uncompetitive in this scenario, leading to their immediate extinction – hence, these *non-viable species* have a life time of 1 step and do not appear in the figure.

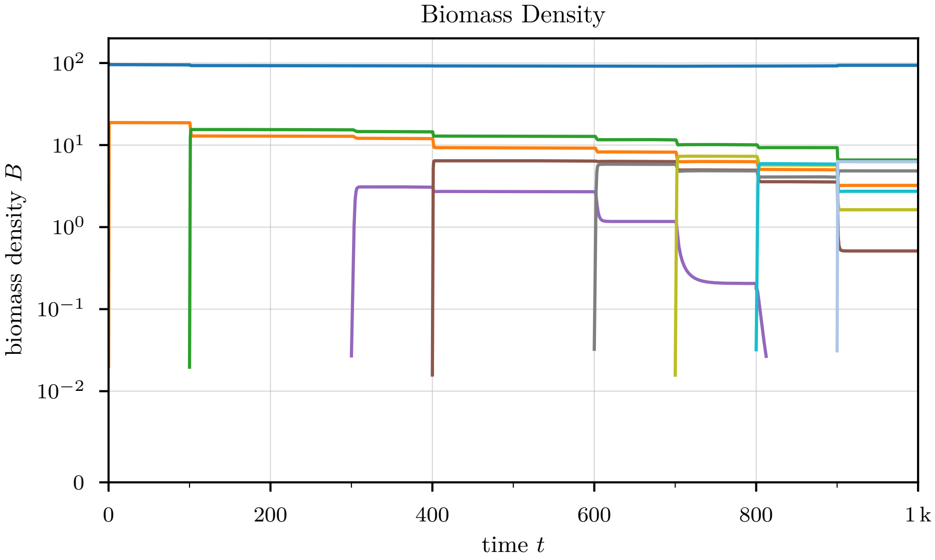

Figure II.3.1 Biomass density of the external resource (blue line) and all species (other lines). Note the symlog y-scale, being linear in $[0, 10^{-2}]$ and logarithmic above.

A characteristic of the ECST model is that new species are typically added into a system that is in dynamic equilibrium. Upon the addition of a new species the system reorganises, with trophic biomass flow rates changing for a set of affected species. The transients during this reorganisation are usually very fast (few time steps), but depend on the traits of the added species. For instance, the addition of the species at $t = 700$ seems to elicit only a small increase in out-flow from or small decrease in in-flow to species 5 (purple line), such that the transient is somewhat slower. Note that while the transients shown here are monotonic, they may also exhibit overshoots and oscillations (not shown).

Figure II.3.1 also shows that – in this early state of the food web – offspring species are often able to coexist with the existing species. However, an example of displacement is observable in Figure II.3.1 as well: through the addition of the species at $t = 800$ (cyan line), the species added at $t = 300$ (purple line) becomes extinct shortly afterwards.

The system dynamics of this event are shown in Figure II.3.2, which also uses a different mode of representation: the species-specific *number density* $n_i = B_i/m_i$. The number density is typically more insightful than the biomass density, as it allows to judge a species' size with respect to the extinction threshold n^{min} (shown as horizontal line at $n = 2 \cdot 10^{-4}$). For that reason, most of the following plots will use the number density instead of the biomass density.

Figure II.3.2 also offers an explanation for the extinction of species 5 (purple line): The discontinuity in its number density is due to biomass being transferred to its

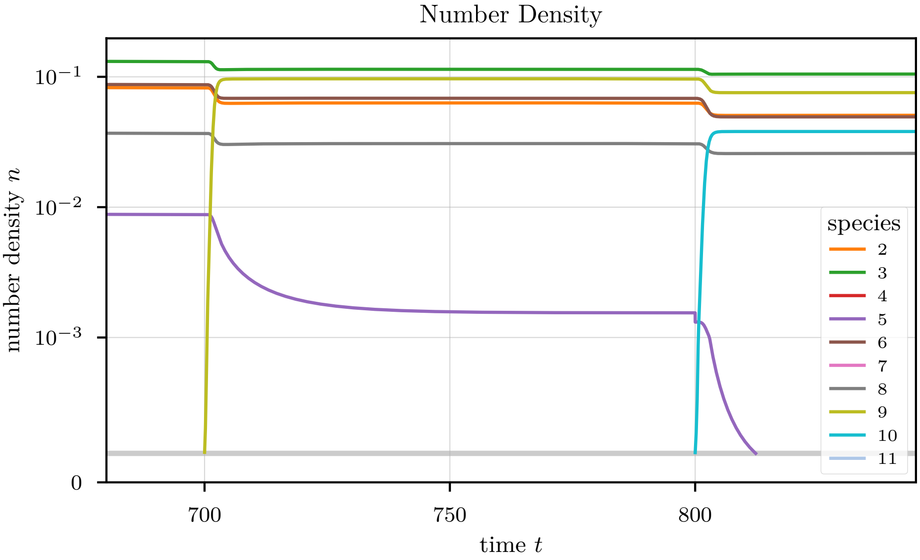

Figure II.3.2 A section of the same simulation shown in Figure II.3.1, but plotting the number density $n_i = B_i/m_i$. The external resource (species 1) is omitted here. The horizontal line denotes the extinction threshold n^{min}, which is also the number density at which offspring species are added. Note the symlog y-scale, being linear in $[0,\, 10^{-3}]$ and logarithmic above.
The discontinuity in the number density of species 5 is due to it being the parent of species 10, thus requiring a biomass transfer to the offspring species (see main text).
Note that species IDs start at 2, while the external resource has the ID 1.

offspring, the newly-added species 10 (cyan line). This event coincides with species 5 dropping under the extinction threshold a short time later. While this suggests that the parent and offspring species were in competition, causing extinction of the parent, these kinds of statements are not as easily made as they appear: Due to the mutation mechanism in the ECST model (see Section II.2.2.2), offspring and parent might have vastly different feeding kernels; only their body masses can be expected to be similar. Subsequently, them being parent and offspring – implying similarity in traits – can *not* be seen as the underlying cause for an extinction.

To better understand the interactions in a food web, it is insightful to look at the actual trait values of the species, as these determine the feeding interactions. Figure II.3.3 shows the traits of all species as well as their life times. In addition to the traits, the *relative feeding distance* m_i/f_i can put a species' feeding preferences into perspective with other species.

The trait overview shows one of the design choices of *Allhoff and Ritterskamp* et al. [2015] for the ECST model most dominantly: Mutations in m depend on the trait value of the parent and are close-by, while mutations of the feeding kernel are largely independent from the parent value and thus typically large. Therefore, the

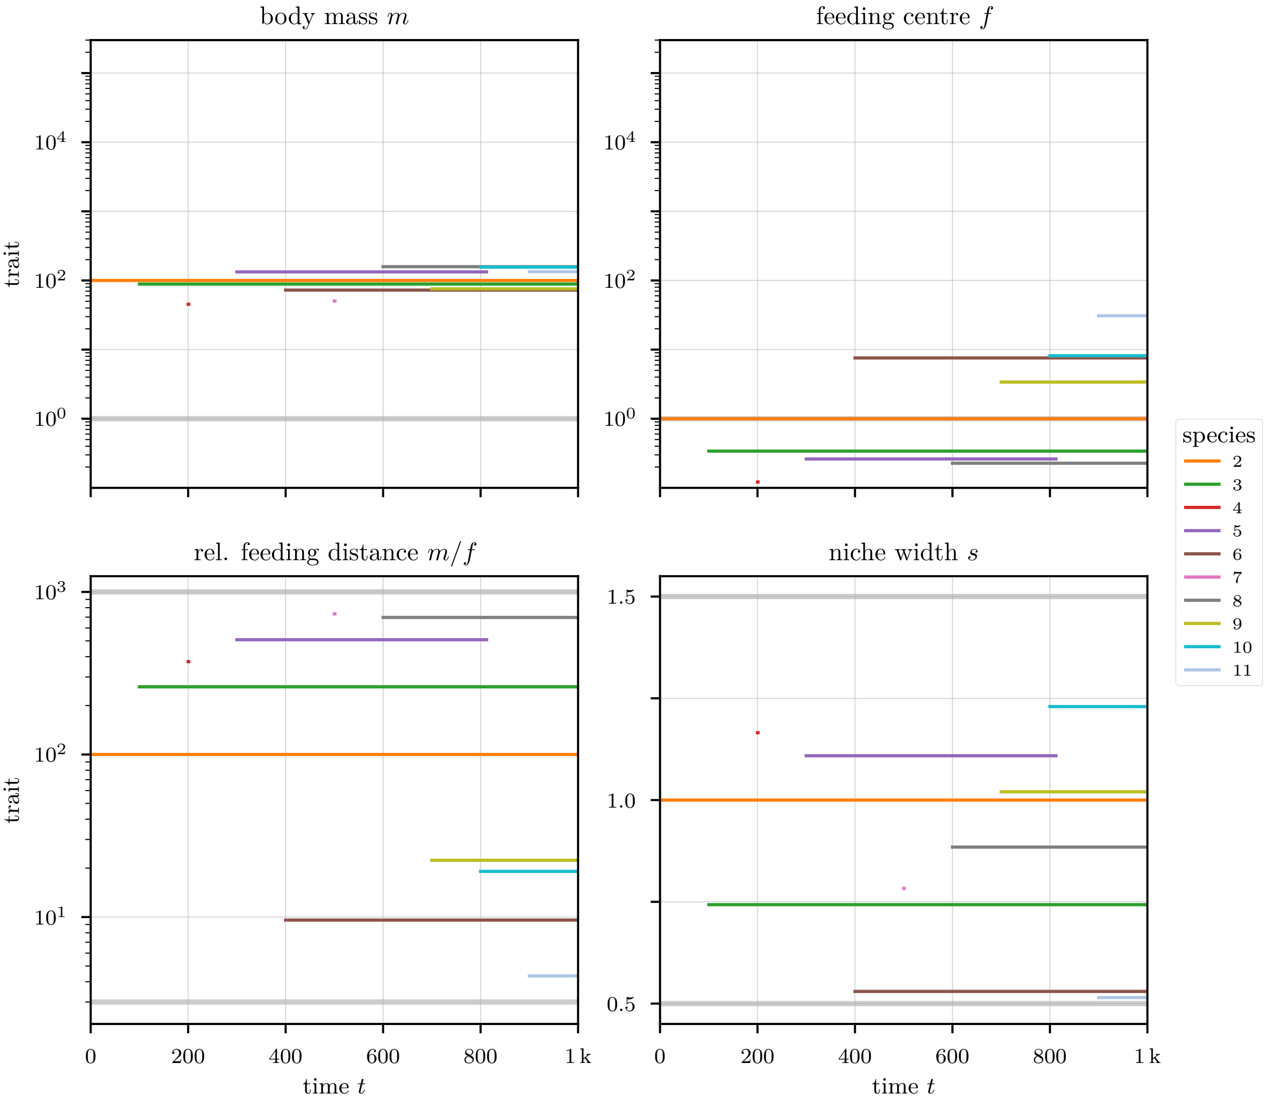

Figure II.3.3 Trait values and life times of all species, corresponding to the simulation shown in Figure II.3.1. The grey horizontal line at $m = 1$ denotes the position of the external resource; in the other subplots, grey lines mark the range within which trait values may mutate. The species with a very short life time were non-viable; they appear slightly larger here to make them visible and allow putting their traits into relation with those of the other species.

spread in m is still rather narrow after $t = 1000$, while the other traits cover much wider ranges.

We may also learn more about the cause of extinction of species 5 (purple) from the trait overview. This requires a holistic view of the system, putting the traits of species 5 in relation to all other species' traits. With the exception of the initial species, all species' feeding centre is roughly half a decade away from the external resource at $m = 1$. Subsequently, a wider niche width s is of advantage for them to consume the external resource; however, if s becomes too large, they become less competitive (see Equations (II.2.7) and (II.2.10)). The addition of species 10 is particularly eventful because it has a feeding centre of $f \approx 10$: right in between (in log space) the external resource at $m = 1$ and the cluster of species at $m \approx 100$. Furthermore – and unlike the other species – its niche width is comparably large, such that it can feed on both the external resource and all other species. The increased competition for the external resource as well as the added consumption by species 10 can be seen as the reason for why species 5 was displaced.

An alternative representation of feeding interactions – and a visually simpler one – is in the form of a directed network like shown in Figure II.3.4. Here, nodes correspond to species and edges represent interactions. The network representation shows that species 11 is not strongly dependent on its trophic connection to the external resource (species 1) but feeds primarily on species 3, making it a species on a higher trophic level than all others in this network. The self-edge denotes a cannibalistic connection.

While the above network is a result from the functional response matrix g, it does not include all interactions but only those edges that provide for a large proportion of a predator's diet. I will go into more detail on the thresholding method and other network-related observables a bit later in Section II.3.5.3.

In summary, the following characteristics of the ECST model can be observed from this short simulation: (i) mutant species may be non-viable and become extinct directly after their addition to the system, (ii) addition of viable mutant species leads to a new equilibrium state of the system, and (iii) parent and offspring species only have a strong similarity in their body mass trait m.

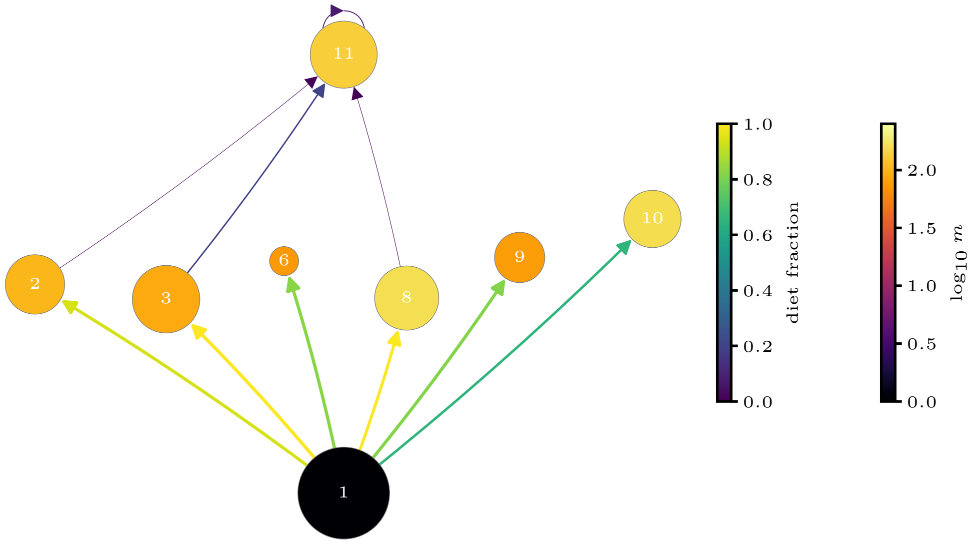

Figure II.3.4 Network representation of the above system (Figure II.3.1) at $t = 1000$. It is generated from the functional response matrix using edge thresholding (explained in more detail in Section II.3.5.3).
Many species and interaction characteristics are encoded in this representation: Node colour (outer colourbar) represents the species' body mass trait m; node sizes scale with the biomass density B_i of the species. The y-position of a node depends on the flow-based trophic level of the species (no axis shown); the x-position is arbitrary. Edge widths hint at the absolute interaction strength between two species, while edge colour (inner colourbar) denotes the fraction of a predator's diet that is satisfied via that edge. Note that the diet fraction may not sum up to 1, which is a consequence of edge thresholding.

II.3.2 Competition & coexistence

As visible in the previous section, species in the ECST model may coexist but also displace each other. This section looks at some of the factors that influence whether species may coexist or become displaced. The constructed scenarios used here all look at the competition of a number of species for a single common resource, while varying the trait values of the species. Furthermore, the effect of the food competition parameter c_{food} (see Equation (II.2.10)) is studied.

Figure II.3.5 shows the time development of three separate food webs, each with five species feeding on the external resource. The species are identical with the exception for one of their traits m, f, or s.

In Figure II.3.5a, species only differ in their body mass m, thus influencing their handling time parameter $h_i \propto m_i^{-3/4}$ (see Equation (II.2.5)) and their respiration rate $x_i \propto m_i^{-1/4}$ (Equation (II.2.3)). With larger m, the handling time decreases and the functional response g_{ij} in turn increases. Furthermore, a larger body mass m reduces the species' respiration rate. This puts the larger species at an advantage in terms of respiration rate and functional response; however, this is counteracted by it having a smaller number density, thus putting it closer to the extinction threshold. However, with the species in this scenario having identical feeding preferences and only slightly different dynamic behaviour, they can coexist. In this scenario, there is also a limited amount of cross-feeding: due to the relative proximity of the species with the smallest m to the resource at $m = 1$, all species feed on that smallest species to some small part.

The picture is a different one in Figure II.3.5b, where the feeding centre f was varied between species. The species that accrues most biomass in this scenario is the one with its feeding centre at $m = 1$, right at the mass of the external resource and thus with optimal foraging capabilities; the species becoming extinct is feeding at $m = 0.1$, thus not being able to forage enough resources to survive. Why does the species with $f = 10$ not become extinct like the one with $f = 0.1$? Such a behaviour would be expected given the fact that the distance of their feeding centres to the external resource at $m = 1$ is equal (in log space, see Equation (II.2.6)). However, the species with $f = 10$ forages biomass not only from the external resource, but also from the other species (and itself) with $m = 100$, subsequently providing an additional energy input that the species with $f = 0.1$ lacks. This additional input is sufficient for it to not become extinct.

When varying the niche width s as shown in Figure II.3.5c, the species that is most specialised on the external resource (has the smallest s) can accumulate most biomass; species that have a wide feeding kernel are at a disadvantage and may become extinct as a result. The benefit from a narrow niche width s follows directly from the normalisation of the feeding kernel, see Equation (II.2.6). As all species' feeding centre trait f matches the position of the external resource at

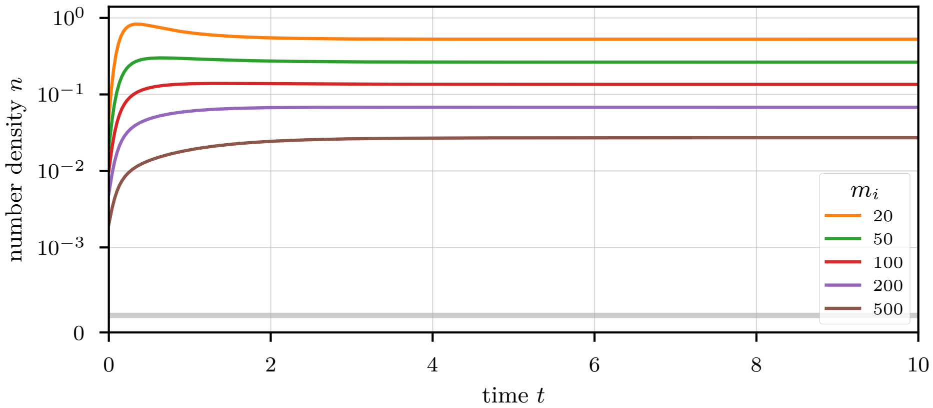

(a) Identical species with varying body mass m.

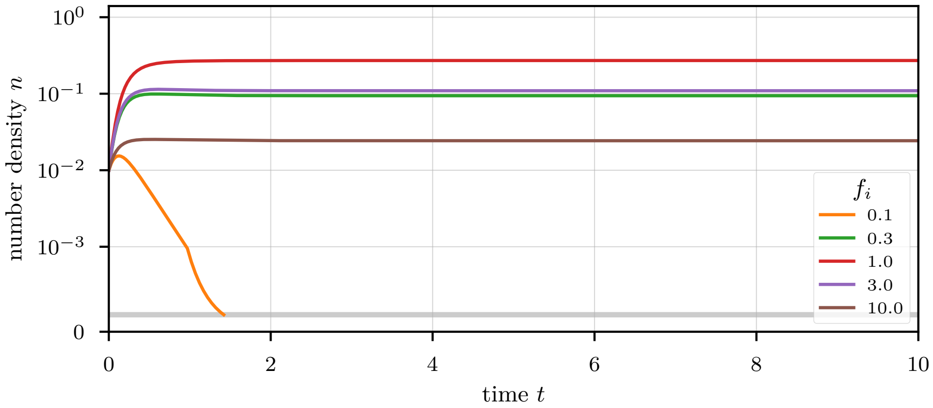

(b) Identical species with varying feeding centre f.

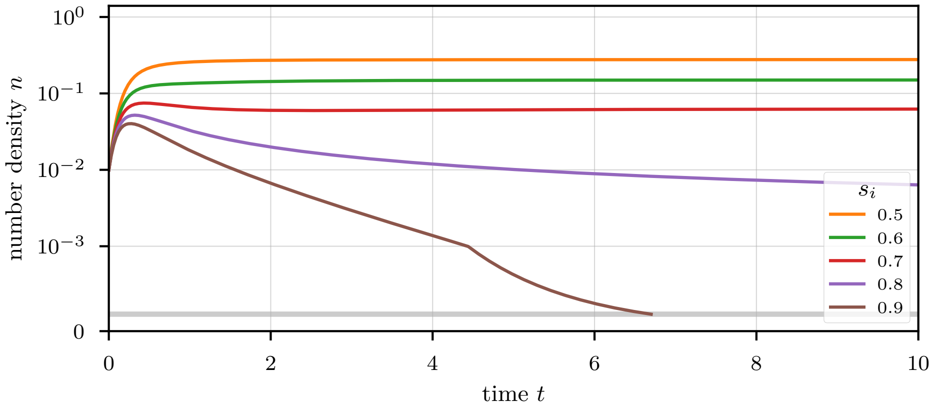

(c) Identical species with varying niche width s.

Figure II.3.5 Three simulations of species that vary only in one of their traits. All other species properties are identical: $m = 100$, $f = 1$, $s = 0.5$, $B(t = 0) = 1$, $c_{\text{food}} = 0.8$. The grey horizontal line at $n = 2 \cdot 10^{-4}$ denotes the extinction threshold n^{min}. Note the symlog y-scale, being linear in $[0,\ 10^{-3}]$ and logarithmic above.

$m = 1$, Equation (II.2.6) simplifies to $a_i \propto s_i^{-1}$; in effect, a narrower niche leads to a higher individual attack rate a_i.

An important parameter in these scenarios is the competition strength, specifically the food competition parameter c_{food}. Together with the feeding kernel overlap I_{ij} of two competing species, it determines all off-diagonal elements of the competition matrix c (see Equation (II.2.10)), and subsequently the equilibrium biomass densities.

Figure II.3.6 shows scenarios like in Figure II.3.5 but with much stronger competition for the external resource ($c_{\text{food}} = 5.0$ instead of the default value of 0.8). Unlike with weak competition, the differences in the feeding kernels lead to the extinction of all but the best-adapted species (small s, f closest to the external resource). Despite the high food competition parameter, coexistence is still possible in these scenarios, but species need to have more similar traits; for instance, a species with $f = 1.1$ would be able to coexist with the species with $f = 1.0$.

Furthermore, Figure II.3.6a shows how species that do not differ in their feeding kernel but only in their body mass m can coexist, albeit with different equilibrium number densities. The overall number densities are smaller compared to Figure II.3.5a due to the higher competition value reducing the functional response g_{ij} of all interactions; because this term is identical for all species, it does not lead to one species being in a more favourable position than another.

The scenarios shown in this section exemplify how species with similar-enough properties may coexist in the ECST model. Their coexistence shows that the *competitive exclusion principle* (CEP) applies in a weaker form in this model: the definition of a *niche* and a *species* is rather soft. Only if there are stark differences between species' feeding kernels do they actually lead to the displacement of a species. How large the trait differences need to be depends on the competition parameter c_{food}.

Aside the competition parameter, the extinction threshold n^{min} determines if a species is actually removed from the system. In fact, Equation (II.2.11) is defined such that a system with positive biomasses will always reach an equilibrium state with positive biomass densities for all species. Specifically, biomass densities will reach zero only asymptotically. The extinction threshold fulfils the role of breaking this asymptotic nature of the ODE system and allows to regard a species population as having become small enough to actually be regarded as extinct. Subsequently, the choice of this value has a strong influence on the level of coexistence possible in the ECST model. Other choices for the extinction threshold are conceivable, like making them depend on body size as well; in such a case the scaling would affect the demographics of the food web. These are not discussed in more detail in this thesis.

The small, artificial food webs compared in Figures II.3.5 and II.3.6 are an illustration of competitive behaviour in the ECST model. In contrast to these small food webs, evolved food webs exhibit many more interactions and species traits of

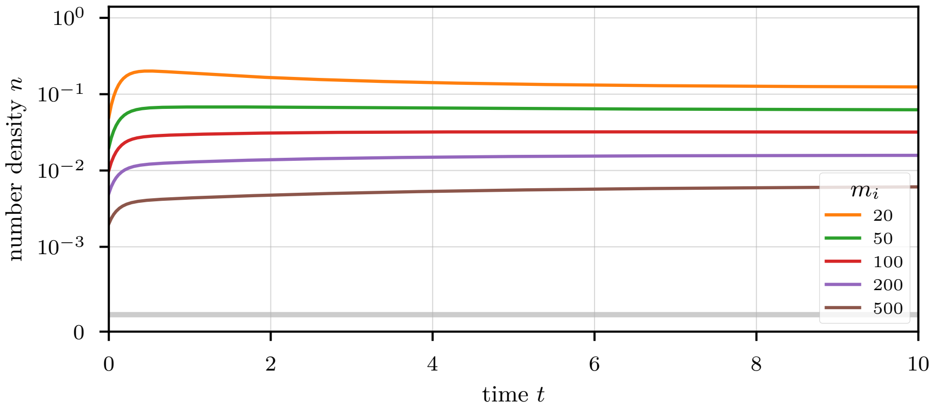

(a) Identical species with varying body mass m.

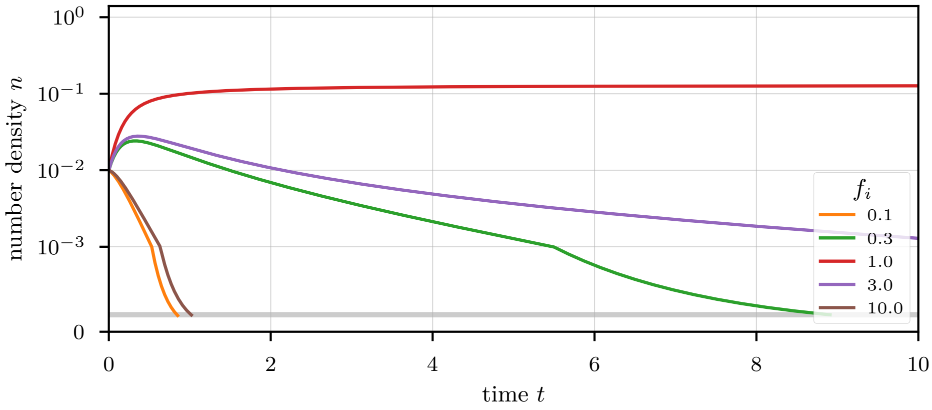

(b) Identical species with varying feeding centre f. The species with $f = 3.0$ becomes extinct at $t \approx 24$ (not shown).

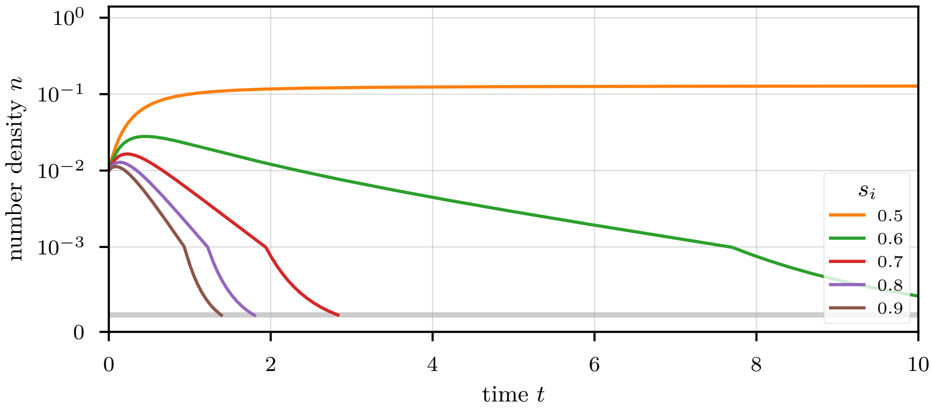

(c) Identical species with varying niche width s.

Figure II.3.6 Competition scenarios as in Figure II.3.5, but with a strongly increased competition of $c_{\text{food}} = 5.0$ (instead of 0.8).
Again, note the symlog y-scale, being linear in $[0,\ 10^{-3}]$ and logarithmic above, leading to the 'kinks' in the time series.

a far higher diversity, making the cause for extinction or coexistence of species less transparent. Yet, the observations made here can be applied to the larger scale, namely that species competing for a common resource *may* coexist as long as their feeding kernels are similar enough; otherwise, they get displaced and are removed from the system.

II.3.3 Artificial food web

The previous section showed species competing for access to the external resource. In general, species may not feed solely on the external resource, but also on each other, thus forming food chains or food webs with multiple layers. While the system's biomass is still supplied by the external resource, the dynamics change when species are dependent on each other through trophic connections.

Below, the effect of adding species to or removing them from a small artificial food web is studied. The food web is referred to as *artificial* because it is constructed manually to serve as an example; it has not evolved via the evolution mechanism described in Section II.2.2.2 and its initial state is not one that would appear in a simulation based on that evolution mechanism.

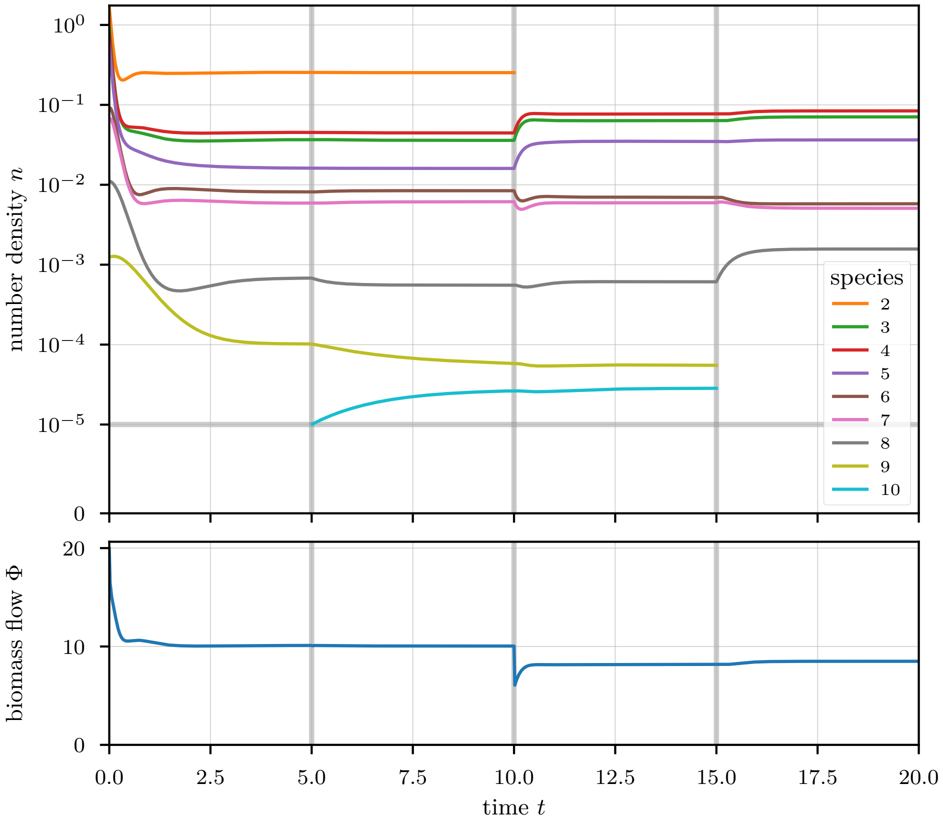

Figure II.3.7 The time development of key observables of an artificially constructed food web. The vertical lines denote active manipulations of the food web. The horizontal line in the upper plot denotes the extinction threshold of $n^{\min} = 10^{-5}$ (which is smaller in this simulation than the default value).

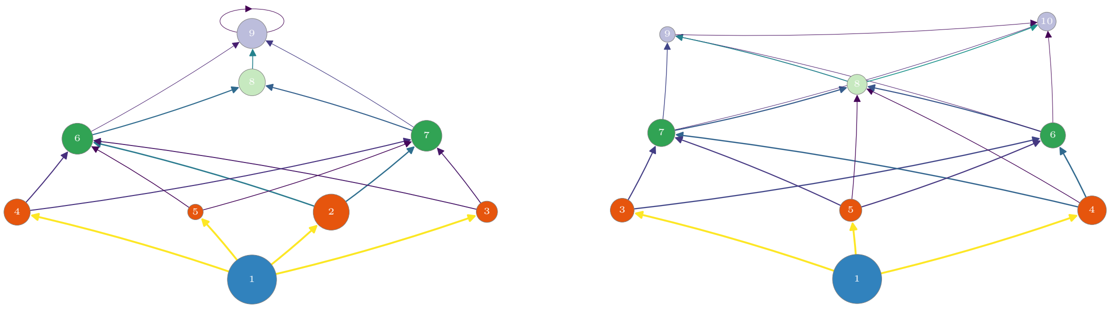

(a) $t = 5$

(b) $t = 15$

Figure II.3.8 Artificial food webs at two different times of the simulation shown in Figure II.3.7. The colour allows a rough association with a trophic level, starting from blue (TL = 1) for the external resource to purple (TL = 4) for the top-predators.

Figures II.3.7 and II.3.8 show key characteristics and network representations of the artificial food web. It is set up with eight species distributed over three trophic levels as shown in Figure II.3.8a. Each species is initialised with a large biomass density of $B = 100$ and the extinction threshold is reduced to $n^{\text{min}} = 10^{-5}$. These changes are important in this artificial scenario to allow species on a higher trophic level to survive until those on lower trophic levels have equilibrated.

The simulation also includes three artificial changes to the food web structure:

- the addition of a new species with a high trophic level (species 10 at $t = 5$);
- the removal of a primary producer (species 2 at $t = 10$);
- and the removal of all top-level predators (species 9 and 10 at $t = 15$).

Species densities change in reaction to these events which illustrates further aspects of the ECST model dynamics and food web dynamics in general. Note that species are added or removed every 5 steps in order to make the dynamics more easily visible; in regular operation of the model, species would only be removed when dropping below the extinction threshold and new species would appear only every $T^{\text{mutate}} = 100$ steps as a result of a speciation event.

The addition of species 10 – a top-level predator – has little effect on species on lower trophic levels. This is mainly due to it being in competition with species 9 for the resources provided by species 6, 7, and 8 (see Figure II.3.8b). Effectively, these lower-level species behave like the external resource in previous simulations for this particular predator.

The removal of species 2 leads to a reduction in competition for other primary producers (orange nodes in Figure II.3.8) and cause an increase in their density. However, the remaining species 3, 4 and 5 cannot sustain the same biomass flow from the external resource, as visible in Figure II.3.7, thus reducing the energy available to all downstream species. In effect, species 6 and 7 equilibrate at a lower density.

Finally, with the removal of the top-level predators 9 and 10 at $t = 15$, species 8 becomes the top-level predator. Not being foraged on any longer, the species density increases, in turn leading to a higher foraging rate on lower-level species and thus reducing their density.

These artificial food webs demonstrate the characteristic dynamics of food web models like competition effects or the cascading of changes throughout the food web, showing how the species are interconnected.

II.3.4 Long-term evolution

While looking at short simulations may elucidate the model's behaviour in itself, models like the ECST model offer little insights into the behaviour of food webs on the short time scale. This is because they do not include mechanisms relevant on short time scales, like a spatial representation of habitats or agent-based representations – they directly abstract to the population dynamics resulting in the equilibrium biomass densities. For that reason, the focus will be on long time scales (100 k to 10 M steps) henceforth; the time scale that food web evolution happens on.

Figure II.3.9 shows how the addition of new species leads to a constant reorganisation of the network, with new mutant species being integrated into the network and others becoming extinct – the constant species turnover that is a key char-

Figure II.3.9 Number density of individual species and the total number of species in the food web for a longer simulation with default parameters.
The data output resolution is $t = 10$, such that only species with a life time greater than 20 are visible in the number density plot; non-viable species do not appear there. Species colour holds no information and is only meant to make lines more easily discernible.

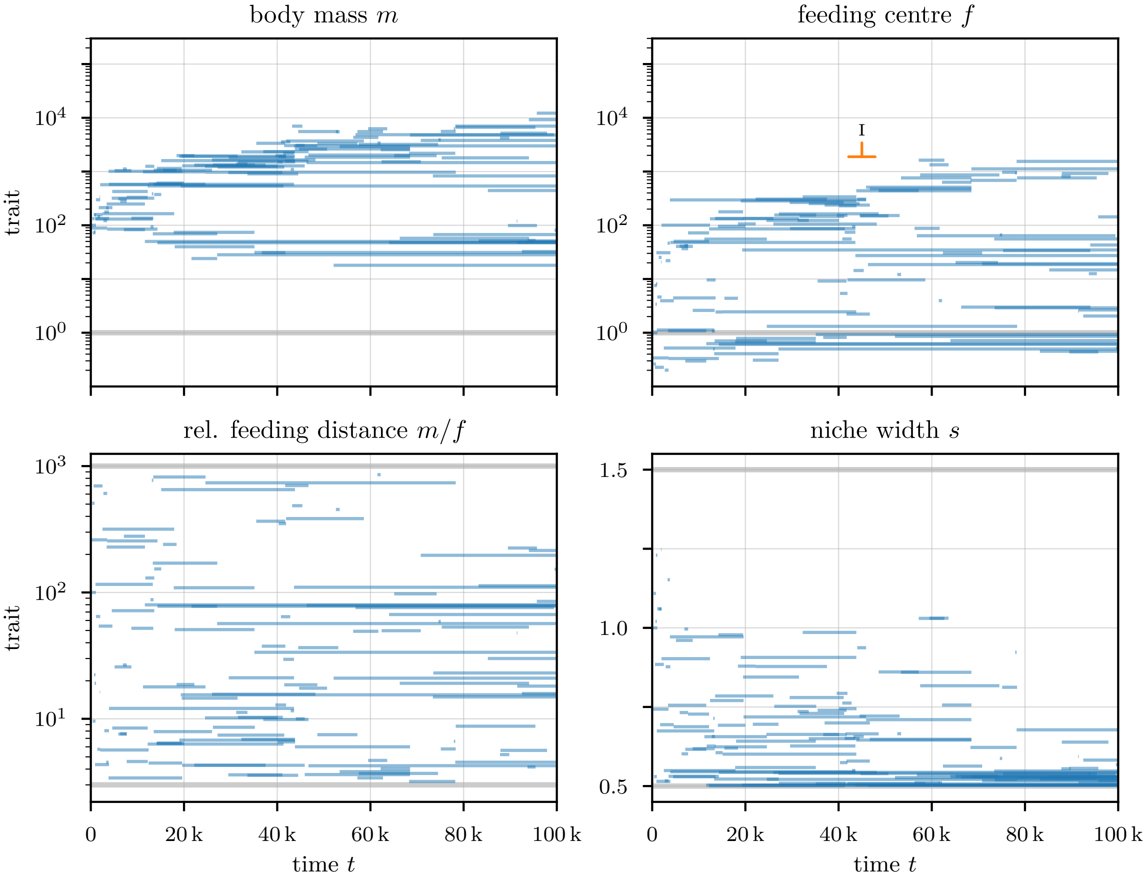

Figure II.3.10 Trait evolution of the simulation shown in Figure II.3.9. The time range marked ‘I’ shows a number of close-by extinction events, followed by the emergence of a species cluster with high-valued feeding centre f. This cluster evolves to feed on species with $m \approx 10^3$ and is thus, following Equation (II.2.13), comprised of species with $m \in [3 \cdot 10^3, 10^6]$.

acteristic of the ECST model. This includes very drastic changes to the system’s composition, like the extinction of a cluster of species at about 45 k steps, with a cluster of new species emerging shortly after. In between these reorganisation periods, the addition of a species typically has little impact and merely leads to the system equilibrating at slightly different values for the species’ number density.

While plots of this kind offer insights into model behaviour on this time scale – e.g., the existence of strong and fast changes in the number density distribution – representations on the species level become increasingly hard to interpret for longer simulations.

In turn, trait evolution plots like Figure II.3.10 become more informative and allow to observe patterns in the collective trait space of all species. For instance, cluster of species with similar traits become apparent and their time development shows branching points. While each species has fixed traits, the continuous species turnover in this model allows for new species with different traits.

In particular, Figure II.3.10 shows how the interaction network evolves from the simple primordial situation to a structured food web with multiple trophic levels. This structure is observable particularly in the m and f traits: there is a cluster of species feeding on the external resource at $m = 1$, basically representing the group of primary producers at $\mathrm{TL} = 2$. Another cluster feeds on species with $m \approx 30$ and yet another at $m \approx 1000$. Furthermore, species with a small niche width s seem to be more successful, going towards the lower boundary of the range of values allowed by the mutation mechanism (Equation (II.2.13)).

The emergence of the species clusters and the trend towards small niche widths is a consequence of evolutionary pressures and the foraging mechanics of the ECST model: On the one hand, species can forage more biomass if their feeding kernel aligns well with many large species. In such a situation, small niche widths are of benefit in a competitive scenario (as seen earlier in Section II.3.2), explaining the evolutionary pressure towards small s. On the other hand, species with a wider niche width may tap into multiple species clusters, if these clusters are not too far away from each other. The s trait basically denotes the standard deviation of the gaussian feeding kernel $\mathcal{N}_{ij}$ of a predator species i on a prey species j, in units of $\log_{10} m$. Thus, a species with $f = 10$ can feed equally well on species with $m = 1$ and $m = 100$; but the attack rate will only be large enough if s is not too small.

Figure II.3.11 shows the continuation of the previous simulation and allows to observe trends on a longer time scale: The feeding centre trait continues to evolve into four clusters, with a large number of species feeding on the external resource and a long-lived species with high f forming the top predator. Furthermore, the trend towards small niche widths is continued, letting species with small s dominate the food web.

To briefly (and qualitatively) compare how the structure of the extracted networks changes over time, Figure II.3.12 shows the networks at $t = 100\,\mathrm{k}$ and $t = 500\,\mathrm{k}$. During that time, the network not only grew considerably in size, but also in depth, now allowing for roughly four clusters of species at similar trophic positions (along the y-axis). Also note how the lowest level becomes much more populated – this wider base layer is what allows the emergence of the higher-level species; to that end, the low-level species need to extract sufficient biomass flow from the external resource to accommodate the top-predators.

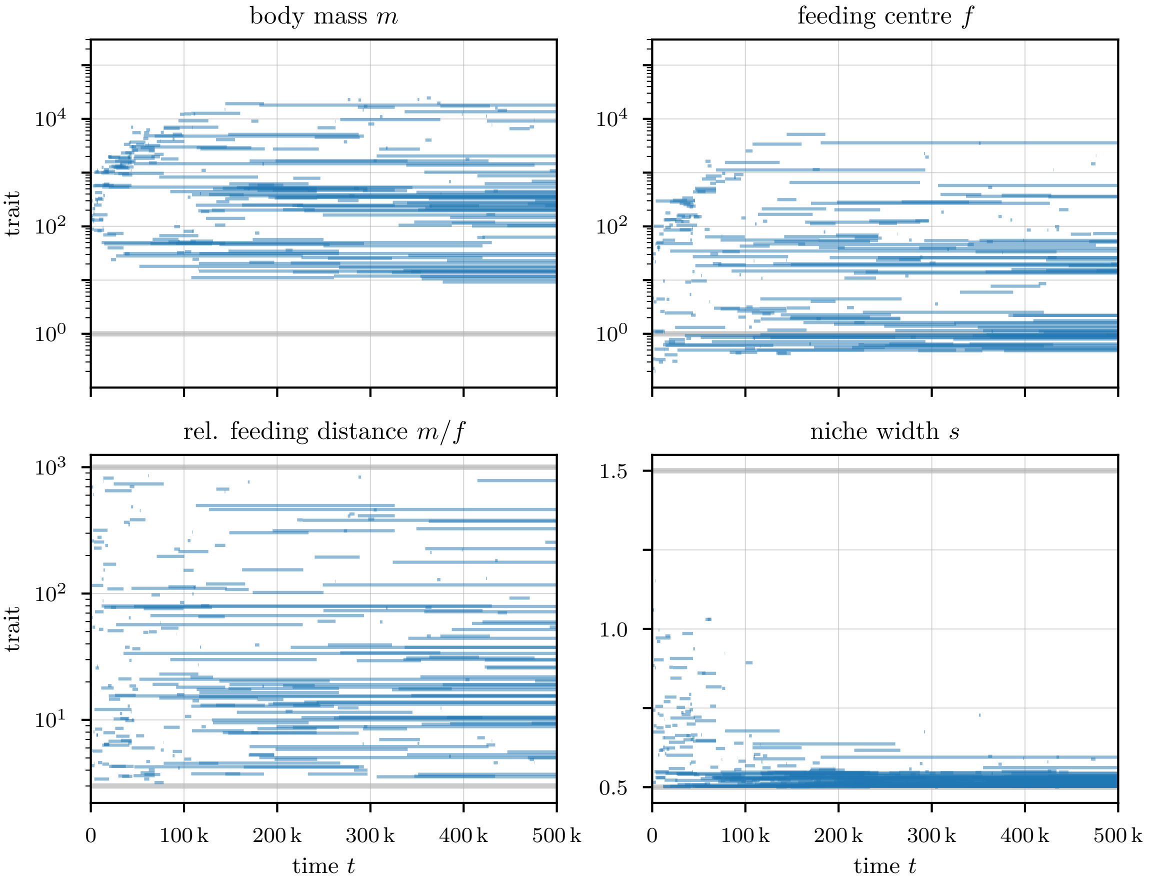

Figure II.3.11 A continuation of the previous simulation over 500 k steps.

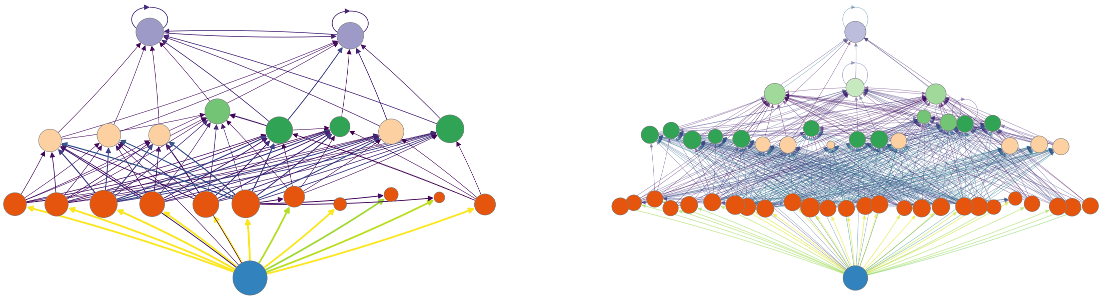

(a) $t = 100\,\mathrm{k}$ **(b)** $t = 500\,\mathrm{k}$

Figure II.3.12 Network representations of the simulation shown in Figure II.3.11. Node colour qualitatively encodes the species' trophic level.

II.3.4.1 Alternative mutation mechanism

The mutation mechanism implemented in the ECST model introduces a large amount of stochasticity. This is primarily because only the offspring's body mass trait depends on the parent's body mass; both the feeding centre and the niche width are drawn from uniform distribution. Even though the feeding centre depends indirectly on the parent's feeding centre, the strong mutation in that trait practically remove any dependency between the two (see Section II.2.2.2). As proposed in [*Allhoff* 2015, ch. 4.5], other mutation modes are conceivable where offspring inherits more properties from the parent species, for instance by drawing the values from gaussian distributions centred around the parent species' value.

I further explored the idea of stronger inheritance links and want to briefly illustrate how this affects the evolution in trait space. While interesting, there are some reasons why a stronger inheritance mechanism is problematic in the ECST model, which is why this will only be an illustration here – in the rest of the thesis, the regular mutation mechanism is used.

In the gaussian mutation mode, the inheritance strength between parent and offspring is larger, meaning that their traits are more likely to be similar. However, unlike the proposition made in [*Allhoff* 2015, ch. 4.5], I implemented a version where not the traits (m, f, s) are inherited but $(m, {}^{m}\!/_{f}, s)$, which turned out to be more robust. With the respective standard deviations σ, the alternative mutation mechanism is defined as follows:

$$\begin{aligned} \log_{10}(m_o) &\sim \mathcal{N}\left(\log_{10}(m_p),\ \sigma_m\right) \\ \log_{10}\left({}^{m_o}\!/_{f_o}\right) &\sim \mathcal{N}\left(\log_{10}\left({}^{m_p}\!/_{f_p}\right),\ \sigma_{m/f}\right) \\ s_o &\sim \mathcal{N}\left(s_p,\ \sigma_s\right) \end{aligned} \tag{II.3.1}$$

However, it is important that the inherited values stay within certain bounds, otherwise the food webs frequently collapse or the species turnover comes to a halt. To that end, the values for ${}^{m_o}\!/_{f_o}$ and s_o are re-drawn if they were not within the bounds; while this may skew the effective distributions, it maintains the correlation between parent and offspring traits[42]

Figure II.3.13 shows the trait evolution of two simulations with different mutation modes. It illustrates why the regular mutation mode can be characterised as 'low-fidelity' inheritance: There is a very high variance in trait space and only the body mass trait shows weak branching patterns. For the alternative mutation mode, the branching pattern is not only more pronounced in the body mass trait, but also observable in all other traits – even for the niche width. This is caused by the

---

[42] In [*Allhoff* 2015, ch. 4.5], a different approach is used to achieve this: The respective gaussians are multiplied with another gaussian that is centred around the middle of the valid bounds such that values are typically within the bounds. I also explored this approach but found it difficult to balance, as it introduces additional parameters in form of the second gaussian.

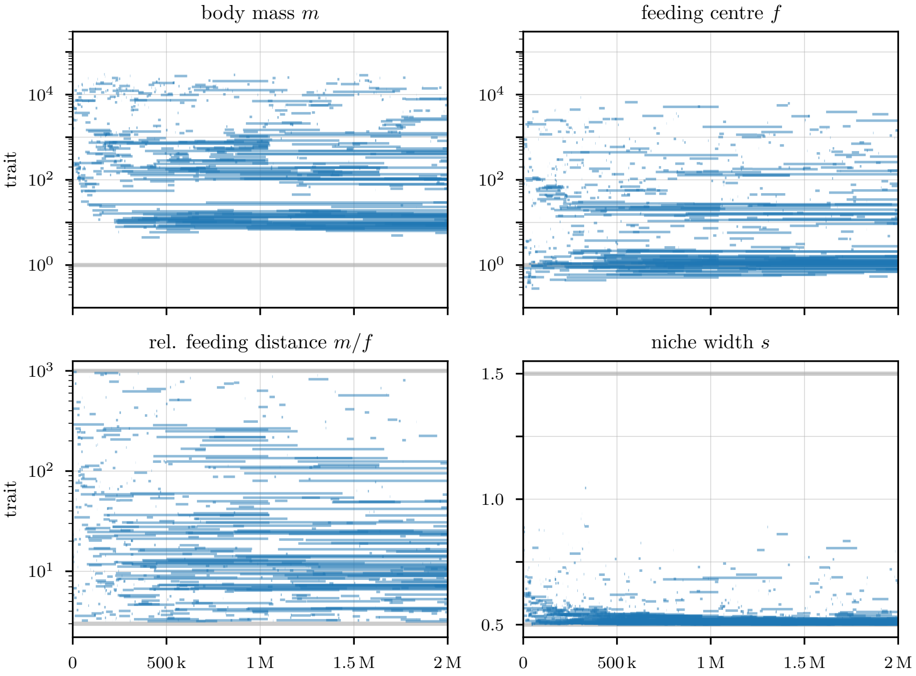

(a) Regular mutation mode

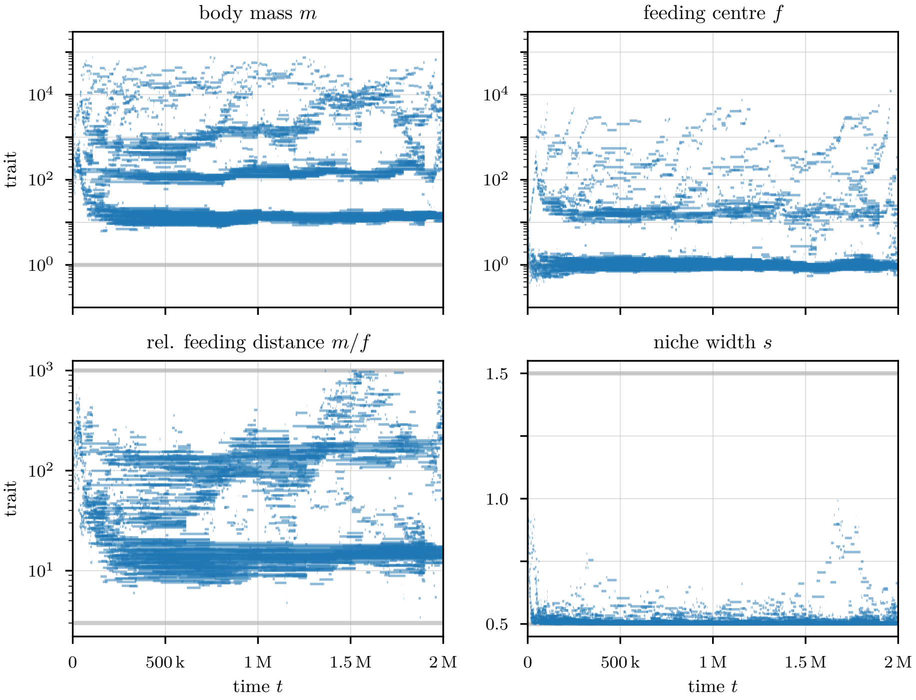

(b) Alternative mutation mode

Figure II.3.13 Trait evolution of two simulations with different inheritance modes. For the regular, low-fidelity inheritance mode **(a)**, offspring species appear in wide ranges of trait space. With the alternative mutation mode **(b)**, offspring and parent traits are more similar, causing contiguous structures in state space and a characteristic branching pattern. (Parameters: $\sigma_m = 0.2$, $\sigma_{m/f} = 0.1$, $\sigma_s = 0.05$)

higher similarity between parent and offspring and visualises how certain changes in the traits are more evolutionary successful.

In the feeding centre trait, the branching forms a particularly intricate pattern, with many branches growing towards larger values and then typically ending. This is a result of the mutation occurring on $^m/_f$ and not the feeding centre f directly. Thus, due to the absence of intricate branching in $^m/_f$, they can be seen as being caused by offspring with larger m being more successful; however, this development is limited because high m species will be closer to the extinction threshold.

Same as in the regular mutation mode, species with low values of the niche width s are generally more successful; the fact that the alternative mutation mode still enforces the bounds of [0.5, 1.5] is clearly observable here. As can be seen in the time after $t = 1.5\,\mathrm{M}$, branching towards higher values is still possible, including the development back towards smaller values.

Other effects of the alternative mutation mechanisms are a much larger number of species, frequently surpassing 150 after roughly 3 M steps. In contrast to the regular mutation mode, where the species number equilibrates somewhere between 70 and 100 species, it continues growing with the alternative mutation mode. This is not surprising: Due to the offspring traits being drawn using gaussians, it is not unlikely that an offspring will be almost identical to a parent. As shown in Section II.3.2, species with very similar properties will likely coexist; species will only be displaced if their differences are large enough. In effect, successful species accumulate, but are almost identical in their traits; instead of few high-B species in the regular mutation mode, the alternative mutation mode results in many low-B species.

The high number of species is problematic, because computational costs for the simulations increase strongly[43]. Unfortunately, species number cannot be easily reduced by adjusting the competition parameters, as that mechanism acts more strongly on dissimilar species with the same feeding centre; while causing the displacement of some of those species, it primarily causes a narrower trait distribution within which a large number of species coexist.

For longer simulation times, this aspect of the dynamics lead to the food webs ‘freezing out’, causing an irrecoverable reduction in diversity. Triggered by a larger extinction event, the number of branches is reduced; with parent species being selected randomly, clusters with a high number of species are more likely to reproduce – but as the offspring will be similar to the selected parent, it is the same cluster that is growing in size, not contributing to a larger diversity of traits in the food web. Thus, this alternative mutation mode may run into scenarios where too little variation is introduced through mutations, effectively bringing the evolutionary development to a halt.

For these reasons, the alternative mutation mode as proposed above is not feasible for reliable investigation of food web properties. While it provides an inheritance

---

[43] After all, the ODE solver operates on the functional response matrix which has N_s^2 entries.

path with higher fidelity, it introduces operationally difficult behaviour into the model, which would require further changes or extension to the mechanism in order to become a robust model of a mutation mechanism. I will further discuss this in **??**.

Evolutionary trend towards small niche widths I want to briefly remark on the niche widths trending towards low values for both mutation modes. As noted alongside the discussion of Figure II.3.10 above, there are evolutionary advantages of specialisation, because more specialised species have a higher attack rate than less specialised ones with the same feeding centre. This trend is limited by the enforced boundary of both mutation mechanisms.

Ideally, an evolutionary model would aim to alleviate parameters that enforce certain trait ranges and rely fully on the model mechanisms to keep traits within reasonable (i.e., non-pathological) ranges. In practice, this turned out to be very difficult for the ECST model.

One remedy I explored was to introduce another mechanism which makes it disadvantageous to have small niche widths (not shown). The motivation behind this would be that high specialisation on prey becomes costly; it was implemented accordingly: For smaller niche widths, the species' respiration rate was multiplied with an exponentially growing factor, thus increasing their living costs for very small niche widths. An alternative approach I tested was increasing the handling time h, leading to a lower functional response.

However, both cases are conceptually equivalent to that of the strictly enforced bounds in that they require parameters that determine which range of values is sensible for s. While they make the boundary softer and wash out the trait distribution in s, the bound is still prescribed rather directly by the newly introduced mechanism and does not emerge indirectly. Thus, imposing costs for low niche width values is qualitatively equivalent to directly enforcing boundary values, even if the former leads to seemingly more ‘natural’-looking distributions and differs in that the range of reasonable values is mediated via the respiration rate or the handling time.

To keep the model's complexity low, I did not enable this mechanism throughout the other simulations made for this thesis.

II.3.5 Measures of ecosystem functioning

Observing the behaviour of individual species in the ECST model illustrates the model mechanisms on a microscopic level. However, with the aim being to better understand ecosystem evolution as a whole, macroscopic measures of ecosystem functioning are required. In the following, I introduce the relevant measures and illustrate their behaviour for the data from the simulation shown in Figure II.3.11. These measures were partly studied in [*Allhoff and Drossel* 2016], but are extended here with network-based measures.

II.3.5.1 Biomass-based measures

The dynamics of the interaction network evolving in the ECST model essentially represent a flow reactor: The basal resource regrows using some external energy input which is the singular energy input into the system. Outflow occurs at each node (species' metabolic losses and death rates) and edge (conversion losses) of the network. In between, species accumulate biomass and a food web structure emerges where trophic interactions transport biomass between species. Subsequently, observables based on the aggregated biomass in the system and the biomass flow rates at different interfaces are a useful characterisations for the functioning of the ECST model.

Measures based on the biomass density are the *resource biomass* B_0, the *species biomass* B_S, and the *total biomass* B^tot of the system. The latter two are defined as follows:

$$B_\mathrm{S} = \sum_{i=1}^{N_\mathrm{S}} B_i \tag{II.3.2}$$

$$B^\mathrm{tot} = B_0 + B_\mathrm{S} \quad . \tag{II.3.3}$$

In addition, three kinds of biomass flow rates are useful: The *resource output* Φ_0^out, the *intra-guild biomass flow* Φ^ig, and the *metabolic losses* Φ^x:

$$\Phi_0^\mathrm{out} = \sum_{i=1}^{N_\mathrm{S}} g_{i0} B_i \tag{II.3.4}$$

$$\Phi^\mathrm{ig} = \sum_{i=1}^{N_\mathrm{S}} \sum_{j=1}^{N_\mathrm{S}} g_{ij} B_i \tag{II.3.5}$$

$$\Phi^x = \sum_{i=1}^{N_\mathrm{S}} x_i B_i \quad . \tag{II.3.6}$$

These flow rates are all formulated as *out*-flows Φ^out from the respective prey species (or the resource) and thus do not include conversion losses (the λ factor in Equation (II.2.11)). This is because any kind of conversion losses are attributed to the predator species, where the foraged biomass needs to be converted into a usable form.

Using the above rates, a relation for the equilibrium state of the ODE system – where outflows balance inflows – can be formulated:

$$0 = \sum_{i=1}^{N_\mathrm{S}} \dot{B}_i = \sum_{i=1}^{N_\mathrm{S}} \left(\Phi_i^\mathrm{in} - \Phi_i^\mathrm{out} \right) = \lambda_0 \Phi_0^\mathrm{out} - \Phi^x - (1 - \lambda_\mathrm{S}) \, \Phi^\mathrm{ig} \quad . \tag{II.3.7}$$

Note that the specific equilibrium values for $(\Phi_0^\mathrm{out}, \Phi^x, \Phi^\mathrm{ig})$ will change each time a mutant species is added to the system.

Figure II.3.14 depicts these biomass-based measures for the simulation shown in Figure II.3.11. It shows that the system's biomass is roughly equally distributed between the external resource and all species. The resource has a much smoother time series than the species collective due to it being consumed by many species and the species being subject to extinctions; the fast changes in B_S can be attributed to a species with high biomass becoming extinct without another species replacing it. Among the measures of biomass flows, it is notable that the species lose more biomass via respiration and intrinsic death rates Φ^x than via intra-guild predation Φ^ig.

Two regions in the shown plot may be highlighted as they exemplify part of the model dynamics. In the region annotated with 'I', a sudden increase in biomass extracted from the resource can be observed, with no such change in Φ^ig. Subsequently, the increase in Φ_0^out must have been caused by a new primary producer that was better adapted than existing species, hence extracting more than the others. This is in contrast to the change observed at 'II', where species biomass B_S drops and Φ^ig grows, hinting at a case where a new species was added to the system that feeds primarily on the other species, not the external resource.

II.3.5.2 Functional diversity

While measures based on the biomass or biomass flow are able to characterise the system's ability to retain and turn over biomass, they do not represent the species' properties or the structure of their interactions.

A key metric in this context is called *functional diversity* and aims to quantify the breadth of functions that the species in the system fulfil. For instance, assuming a species' function to be to consume resources, then two species that feed on very similar resources would only contribute a small degree of functional diversity: In case one of these species became extinct, the other could easily take its place, given how similar they behave in the context of the whole system. In turn, species that have truly unique feeding preferences would contribute strongly to the system's functional diversity: if one of those species was displaced, no other species could compensate its function (or it would first have to evolve again).

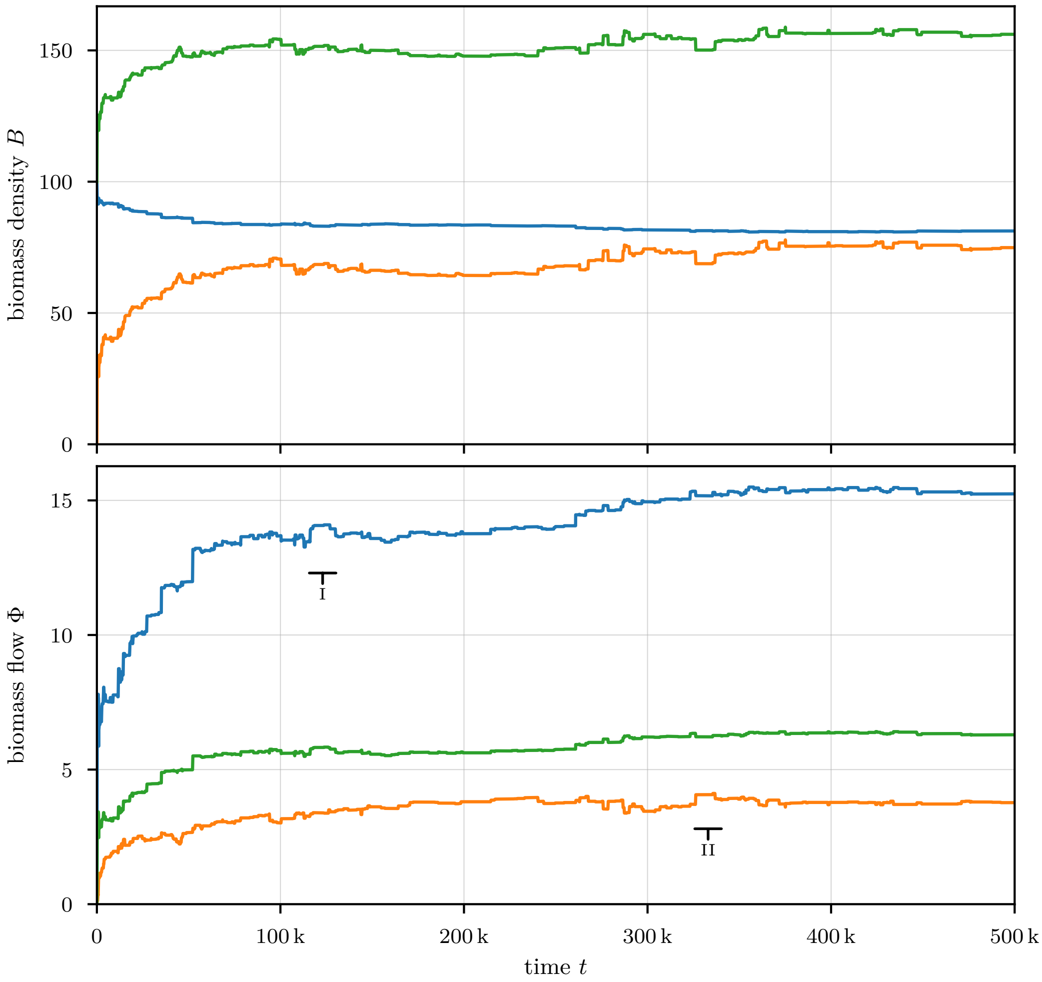

Figure II.3.14 Biomass-based measures of ecosystem functioning; see main text for definitions. The regions annotated with I and 'II' denote interesting developments in these observables. Note that while the biomasses add up to $B^{\mathrm{tot}} = B_0 + B_{\mathrm{s}}$, the biomass flows shown here are outflows and will only add up if weighted as in Equation (II.3.7).

When modelling ecosystems, the functionality of a species in the ecosystem is determined by the model. Depending on the model, a wide range of metrics are used to quantify functional diversity; for a review see [*Petchey* et al. 2009].

In this thesis, I will use the same functional diversity measure as *Allhoff and Drossel* [2016] do for the ECST model, simply denoted as FD and defined as the integral over the envelope of all feeding kernels $\mathcal{N}_i$:

$$\mathrm{FD} = \int\limits_{-\infty}^{+\infty} \max\left(\mathcal{N}_{1j}, \mathcal{N}_{2j}, \ldots, \mathcal{N}_{N_\mathrm{S}j}\right) \,\mathrm{d}\log_{10} m_j \quad . \tag{II.3.8}$$

As the gaussian feeding kernels are normalised to 1 (see Equation (II.2.6)), two species with mostly overlapping feeding kernels will have FD ≈ 1 while two species with feeding kernels far apart would lead to FD ≈ 2. Given the structure of the ECST model, this is a pragmatic way to quantify functional diversity; it is a measure that describes how much of the model's niche space (basically, the $\log_{10} m$ dimension) is covered by the species in system.

II.3.5.3 Network measures

In a food web, naturally, network-based measures are an important set of observables to quantify the interaction structure between species.

In the ECST model, the functional response matrix g, defined by Equation (II.2.4), represents all trophic interactions in the food web and constitutes the most important part of the population dynamics. As a square matrix, it can be interpreted as the adjacency matrix of all species and hence: a network. Due to the nature of g, it is a directed network with weighted positive edges, as previously shown in Figure II.3.12.

Instead of g, the *interaction strength matrix* w is used as the network representation. It is defined from the perspective of the predator species i, hence requiring to include the conversion efficiency λ_j when feeding on prey species j:

$$w_{ij} = \lambda_j g_{ij} \quad \propto \frac{a_{ij} B_j}{m_i} \tag{II.3.9}$$

As such, it represents the biomass in-flow to the predators, which is important due to varying conversion efficiencies between different species. Note that the interaction strength is therefore not directly comparable to the biomass flows introduced earlier, which are formulated as biomass *out*-flows from the prey and thus do not include the conversion efficiency factor.

Based on this interaction strength matrix, further observables can be defined in the following.

Trophic level

One important observable characterising a species' interactions within a food web is its *trophic level*, which relates to its distance from the basal resource of the web. In simple food chains, a basal resources is assigned the trophic level of 1, species feeding on the basal resources would then be on level 2 and so on. However, species often feed on more than one species and those species may be on different trophic levels, such that integer-valued trophic levels are typically not sensible – at the very least, they discard structural information about a species' feeding relations. To that end, a *fractional trophic level* can be defined, which takes into account the composition of a predator's diet [*Pimm and Lawton* 1978; R. *Williams and Martinez* 2004] and the trophic levels of its prey.

A specific interaction's contribution to a predator's diet is referred to as the *diet fraction*

$$v_{ij} = \frac{w_{ij}}{\sum_j w_{ij}} \tag{II.3.10}$$

with predator i and prey j. As $w_{ij} B_i$ constitutes the biomass inflow to the predator, the *flow-based trophic level* is defined as

$$\mathrm{TL}_i = 1 + \sum_j v_{ij} \mathrm{TL}_j \tag{II.3.11}$$

which specifies a system of linear equations with the species' trophic levels TL_i as its solution. In essence, a species' trophic level is the weighted mean trophic level of its prey plus one.

Throughout the literature, other forms of trophic level definitions can be used; see [R. *Williams and Martinez* 2004] for an overview. However, for the ECST model the flow-based TL is most meaningful because the biomass flow is the central part of the population dynamics.

Thresholding

As mentioned earlier, the network visualisations shown so far (e.g. Figure II.3.12) resulted from applying a thresholding procedure to the interaction matrix. For a sensible visualisation, such thresholding is necessary to reduce the number of edges to the most relevant ones: Due to the unbounded gaussian feeding kernels of the ECST model, all species interact with all other species at least to *some* non-zero degree, such that $g_{ij} > 0$ for practically all species, leading to a fully connected network. However, many of these connections are *effectively* zero, motivating exclusion of those edges from the network representation.

The edge thresholding algorithm used by [*Allhoff and Ritterskamp* et al. 2015] is as follows: For a predator i, an edge w_{ij} only becomes part of the thresholded adjacency matrix $\bar{\mathrm{w}}$ if the interaction strength is greater than a factor q of the average link strength $\langle w_i \rangle$:

$$\bar{w}_{ij} = \begin{cases} w_{ij} & \text{if} \quad w_{ij} \geq q \langle w_i \rangle = q \frac{1}{1+N_\mathrm{S}} \sum_{j=0}^{N_\mathrm{S}} w_{ij} \\ 0 & \text{otherwise} \end{cases} . \tag{II.3.12}$$

In [*Allhoff and Ritterskamp* et al. 2015], the thresholding factor is set to $q = 75\,\%$. As briefly discussed there, the choice of thresholding procedure has strong effects on the value and distribution of observables; however, the authors argue that this approach leads to network properties that better match those of empirical food webs than a threshold value proportional to the maximum link strength.

Connectance and trophic incoherence

With a thresholded adjacency matrix, further network measures can be defined. From the wide range of available measures, I will focus on two that have previously been found to relate to the robustness of modelled food webs: the *connectance* as a measure of link density [*Romanuk and Zhou* et al. 2017], and the *trophic incoherence* [*Johnson* et al. 2014] to quantify homogeneity of the trophic level distribution.

The *connectance* is the ratio between realised (i.e., non-zero) links L and theoretically possible[44] links:

$$C = \frac{L}{(1 + N_\mathrm{s})^2} \quad . \tag{II.3.13}$$

For the *trophic incoherence*, the distribution of trophic distances $\mathrm{TL}_i - \mathrm{TL}_j$ is studied. As proposed by *Johnson* et al. [2014], the variance of this distribution acts as a measure of the inhomogeneity of trophic distances. Due to the links in the ECST model having vastly different link strengths, it is important to weigh the distribution accordingly, leading to the following definition for the *trophic incoherence*:

$$\mathrm{TI} = \frac{\sum w_{ij} (\mathrm{TL}_i - \mathrm{TL}_j - 1)^2}{\sum w_{ij}} \tag{II.3.14}$$

In other words, TI is the edge-weighted variance of trophic level differences deviating from the expected mean of 1. It is 0 iff all trophic interactions are between species that are exactly one trophic level apart, leading to a maximally coherent situation[45]. Note that TI does not require the aforementioned thresholding process, because it includes the varying link strengths directly as weights, while the connectance requires the discretisation procedure.

II.3.5.4 Illustration of diversity and network measures

Figure II.3.15 shows the newly introduced functional diversity measure and the network measures over time. Unlike the number of species, which will continue to grow, all measures begin to reach a *relatively* stable level. As can be seen with the annotated regions, they react differently to certain events. Also, the network measures are fluctuating much more strongly; for the connectance that is due to the computation via the thresholded adjacency matrix and for TI this is caused by the fluctuating entries of the adjacency matrix which are used as weights in the computation of trophic incoherence.

---

[44] Due to the resource not being able to be on the receiving end of a trophic interaction, the maximum number of links is $N_\mathrm{s}(1 + N_\mathrm{s})$ in practice.

[45] See [*MacKay* et al. 2020] for an interesting overview of properties of the trophic incoherence and comparison to similar formulations.

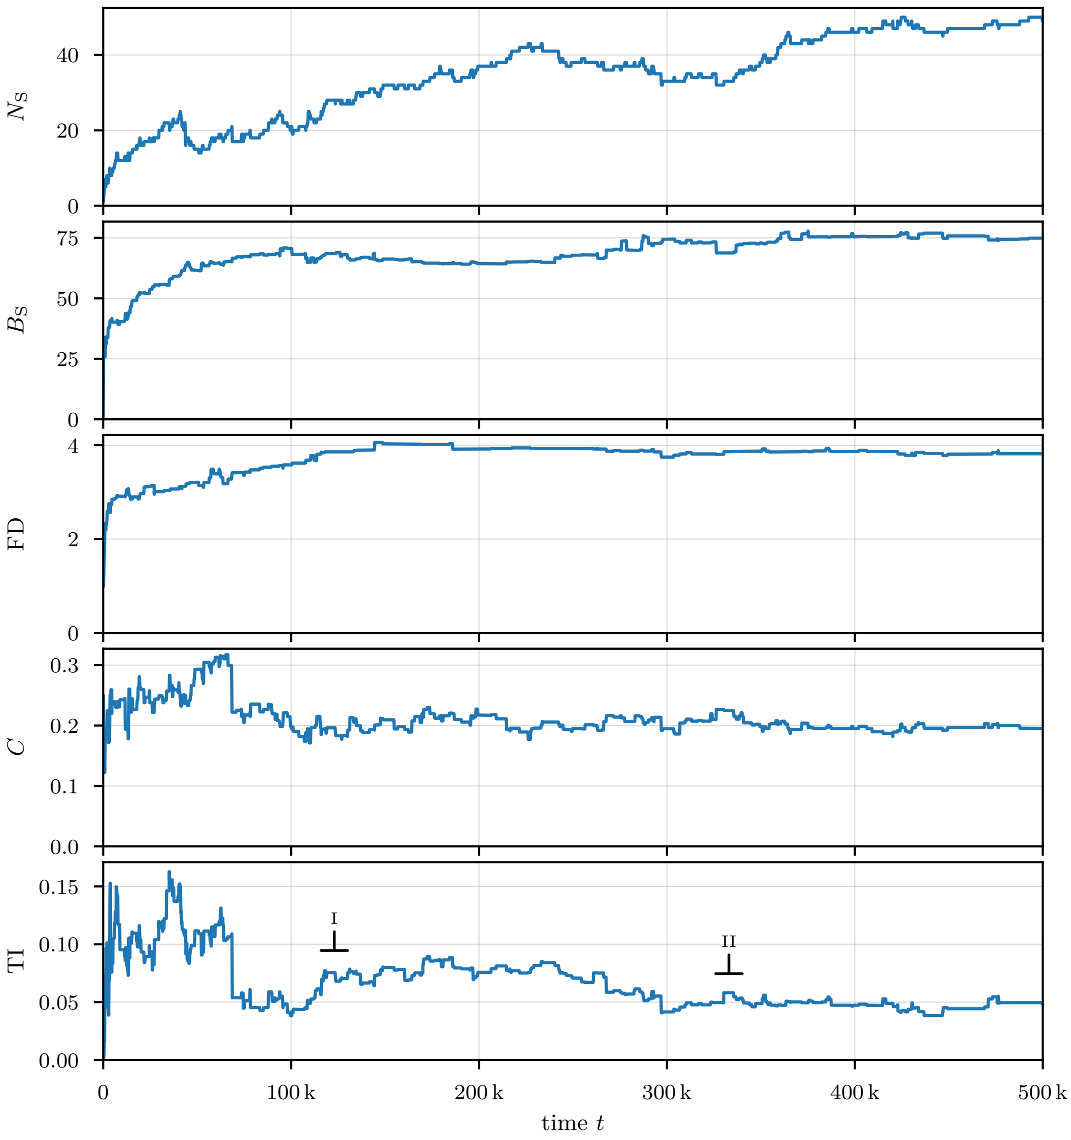

Figure II.3.15 Diversity and network observables over time alongside previously shown observables. The annotated regions in the TI time series correspond to those discussed for Figure II.3.14.

In general, the trophic incoherence value is very low, even for longer simulations, which is due to the edge-weighting in Equation (II.3.14): The strongest links in the network are typically those between the resource and primary producers. At the same time, trophic level differences are only slightly above 1 for this lower level, thus contributing many strongly-weighted entries of small variances to TI. Despite the TI values being small, Figure II.3.15 shows sufficient variance to make this observable informative about the higher trophic levels.

Overall, these measures of ecosystem functioning can give insights into the state of the food web and its time development. They are best seen as complementary pieces of information about the system.

II.3.6 Food web ensembles

So far, figures and analyses considered only the results of individual simulations. While this is important to elucidate general characteristics of the ECST model, it does not allow to reach more general conclusions about the behaviour of this model, because the model dynamics depend strongly on the randomly drawn species trait values. To isolate effects that go beyond the variability caused by a specific random number sequence, simulations and analyses need to include many instantiations of the model with different random numbers. This is achieved by seeding each simulation's pseudo-random number generator (PRNG) with a different value, i.e. varying its `seed` parameter[46].

In Figure II.3.16, key observables for an ensemble of 24 simulations are shown. As can be seen there, N_S may fluctuate strongly and may include fast and large extinction cascades.

Unlike N_S, the other observables shown in Figure II.3.16 reach relatively stable values within a few 100 k steps and the ensembles generally have a low variance. Specifically the biomass-related observables (total biomass B^tot and biomass flow from the external resource Φ_0^out) exhibit less variable time series, while the functional diversity FD and the connectance C may abruptly change.

Despite the biomass-based measures and other observables reaching a quasi-stable state after a few 100 k steps, the number of species typically stops growing only after some 2 . . . 5 M steps; see Figure B.2.

II.3.6.1 Choice of competition parameters

The system size, biomass turnover, and diversity depends strongly on the choice of the competition parameters c_food and c_intra. As compared in Figure B.3, high intra-specific competition c_intra leads to a higher number of species coexisting in the system. This is caused by all species' functional response being reduced by the larger c_intra (see Equation (II.2.4)), hence reducing their consumption rates; as a consequence, more species are able to coexist, but all species have smaller number densities. Apart from that, biomass flow rates and other observables are largely unaffected by changing c_intra.

The food competition parameter c_food (or: *inter*-specific competition) is more influential: Lower food competition leads to larger N_S, biomass flow, connectance, and functional diversity. Ultimately, this is caused by reduced adverse competition effects of species on each other, leading to overall larger consumption rates. Furthermore, through reduced adverse effects, species become viable that would be displaced in a scenario with stronger competition.

---

[46] These `seed` values are typically chosen explicitly in order to make simulations reproducible. In effect, simulations can be regarded as fully deterministic for a specific choice of a PRNG seed. See Section II.2.3 for more information.

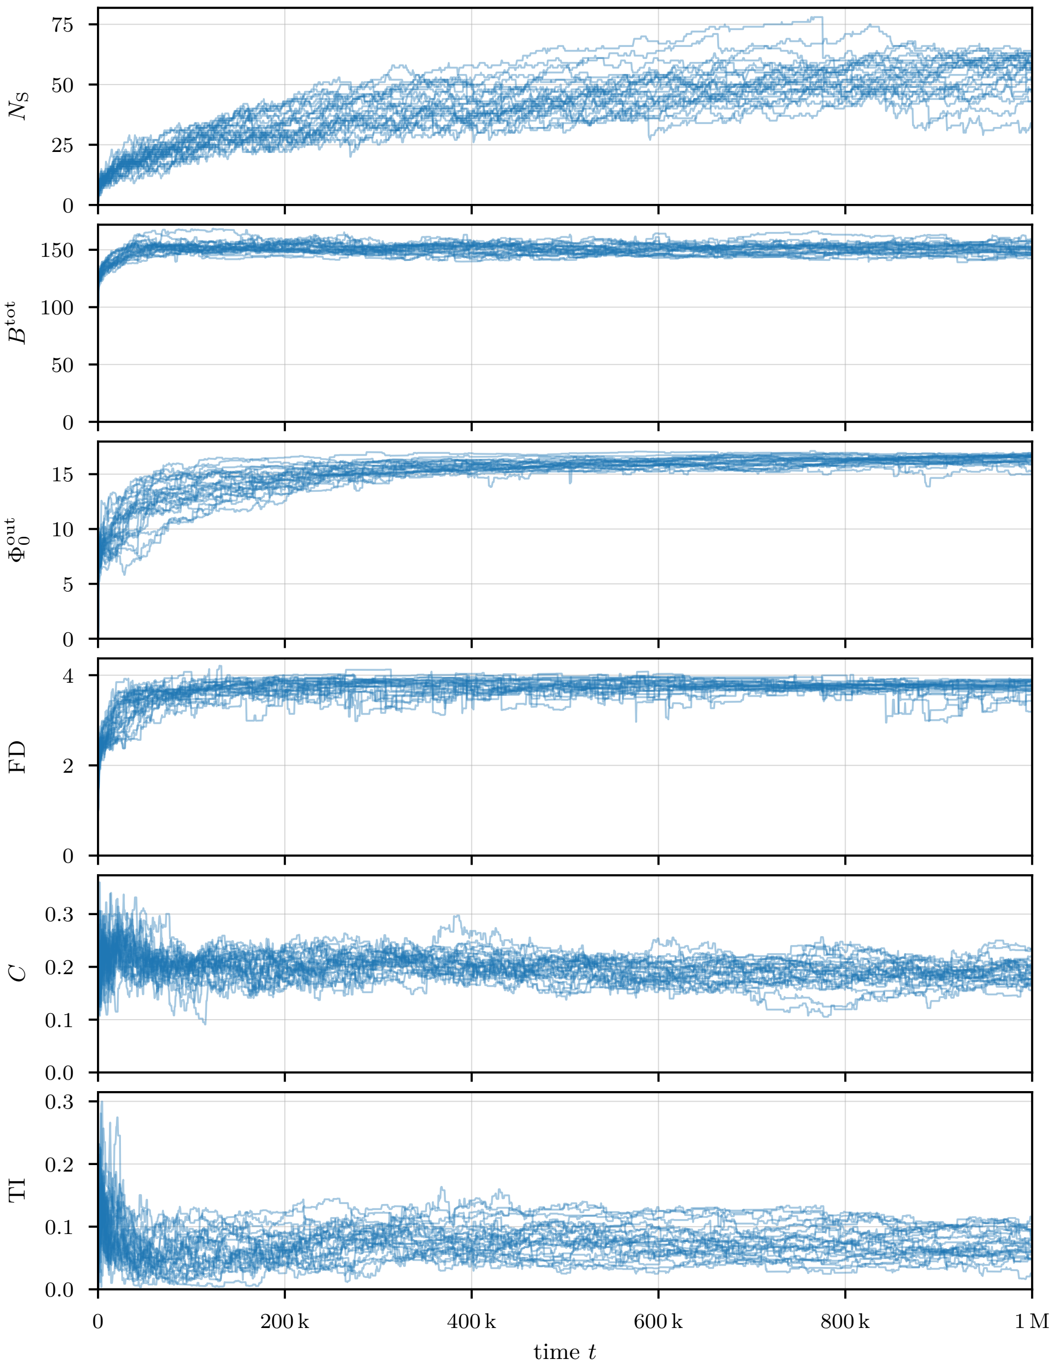

Figure II.3.16 Global observables of 24 instantiations of the ECST model with default parameters (differing only in their `seed` parameter). The output resolution is 1 k steps, equivalent to 10 mutation events.

While interesting, the study of the effect of the competition parameters on the system's behaviour is not an objective of this thesis and is already addressed in [*Allhoff and Ritterskamp* et al. 2015]. Instead, I chose sensible default parameters such that the number of species and the evolving diversity is sufficiently large, yet small enough to make long simulations and large parameter sweeps computationally feasible. Another factor in the choice of parameters was the speed with which the system observables reached a quasi-steady state – which was found to take much longer for large systems, if such a state was reached at all.

II.3.6.2 Species-level statistics

Apart from the time series or trait overview plots shown above, aggregated species statistics can be insightful for the understanding of the model dynamics. Similar to the trait overview, which visualises species-specific properties, other observables relating to single species are collected throughout model iteration. These include scalar quantities like the species' life time (time between its emergence and extinction), whether it was viable (life time greater than 100 steps), the number of extinctions caused by the addition of the species to the food web (the so-called *cascade size*), or its mean trophic level[47].

Figure II.3.17 shows the life time distribution of a large number of long-running simulations with default parameters. The logarithmic representation clearly shows that most offspring species added to the network have a very short life time of $T^{\text{life}} < 100$ and are considered non-viable, with most of these becoming extinct very shortly after they were added to the food web. With the next mutant being added at 100 steps, newly added mutants that are close to extinction are likely to become extinct as a result of this next mutant, hence leading to the abrupt onset in that part of the histogram. From then on, the life time distribution follows a smooth form with a peak at roughly $T^{\text{life}} = 10^4$ steps and many species surviving longer than that. The linear representation provides more detail on the tail of the distribution and shows that there are a substantial number of species with life times on the order of the simulation time, with the distribution dropping off more quickly when nearing the maximum life time.

What are the properties of these non-viable species compared to viable ones?

For one, non-viable species are predominantly on low trophic levels, as can be seen in Figure II.3.18. There are two reasons for this: Due to the higher occupancy of low trophic levels and the uniform probability of all species to be selected as parent species in the evolution mechanism, it is more likely for low-level species to create offspring. In addition, low-level species benefit with a small niche width s are more competitive than those with a larger niche width; this is especially important for the low levels, where there is a lot of competition and only little advantage of being a

---

[47] With the trophic level changing throughout the species' life time depending on food web structure, the mean value is only an approximation. Typically, the time-dependent values of TL are within ±0.2 of the mean value, making it a sufficiently good approximation.

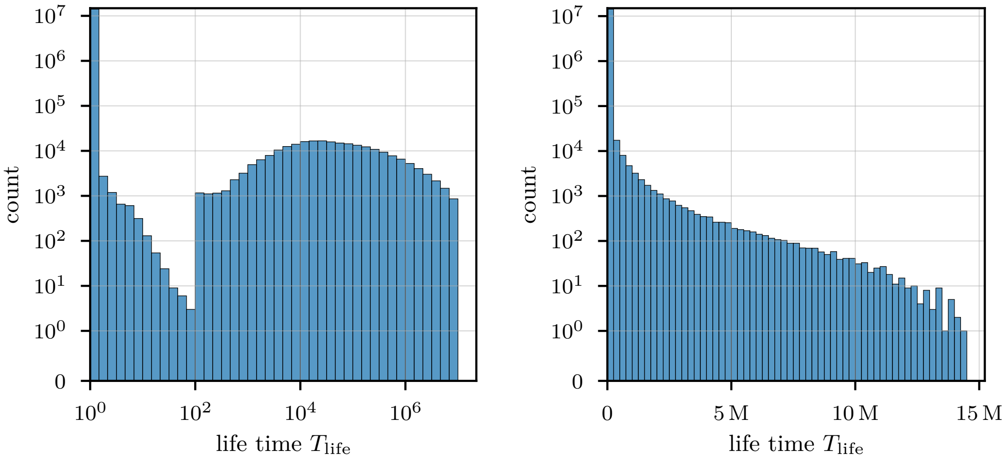

Figure II.3.17 Species life time distribution in logarithmic **(left)** and linear **(right)** representations from 96 instantiations of the model, each simulated for 15 M steps. From $14.4 \cdot 10^6$ species that were added throughout all these simulations, those that were still alive at the end of each simulation were discarded such that only 'full' life times were included; this left approximately $14.3 \cdot 10^6$ counts in total.
The logarithmic representation shows the distribution of short-lived species, most of which become extinct before the next species is added after 100 steps; these species are referred to as *non-viable* and comprise about 95 % of all species added to the system.

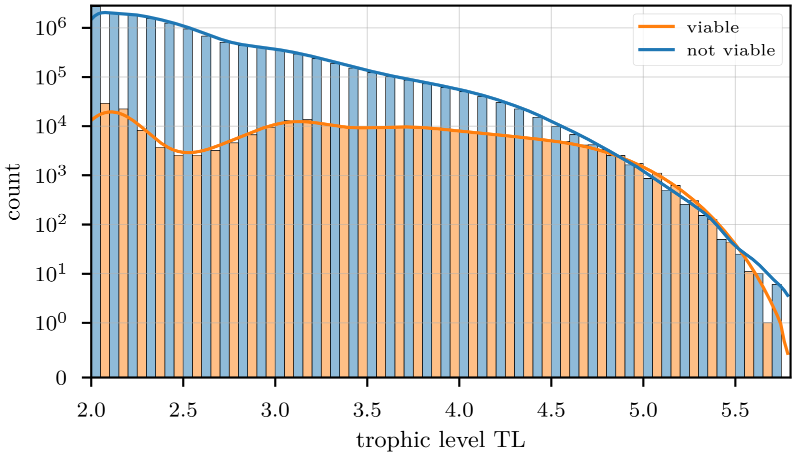

Figure II.3.18 The distribution of species' mean trophic level from the simulation shown in Figure II.3.17, distinguished by species' viability. Viability is assessed by life time: species with a life time smaller than 100 steps are categorised as non-viable. Lines show corresponding kernel density estimations and are only meant as a visual aid.

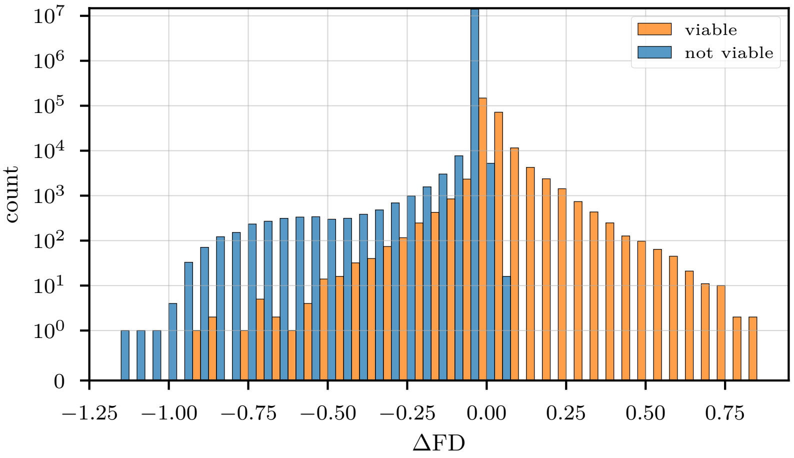

Figure II.3.19 The change in functional diversity, $\Delta\text{FD} = \text{FD}(t_1) - \text{FD}(t_0)$ induced by each mutant species, with t_0 being the time the mutant was added and t_1 being the time before the *next* mutant was added. The histogram is grouped by whether the mutant was viable or non-viable.
(Technical remarks: while binning intervals are right-open, counts in the $\Delta\text{FD} < 0$ bin may also result from inaccuracies in floating point arithmetics. The larger values are due to a miscategorisation of species that became extinct *exactly* after 100 steps; given the low counts and the known cause of this error, the simulations were not repeated.)

more generalist consumer due to the strong clustering of species. Subsequently, the viable volume of niche space in s is strongly reduced. As the mutation mechanism (see Equation (II.2.13)) determines an offspring's s trait from a uniform distribution, many mutants be non-viable alone due to a large niche width.

Species on higher trophic levels are more tolerant towards large s, which also shows in the number of viable and non-viable species being roughly equal. Taken together, this can be seen as an indirect measure of the viable volume in trait space.

Viable and non-viable species also differ in their effect on the food web, for instance in its functional diversity. Figure II.3.19 shows the change in functional diversity directly induced by a mutant species; see the figure caption for information on how these quantities were acquired. As visible there, non-viable species generally lead to a reduction in FD, while viable species can cause both a reduction or an increase. Increases in functional diversity can only result from the mutant species being viable and (thus contributing their feeding kernel to Equation (II.3.8)); this is also the reason for the higher number of counts for $\Delta\text{FD} \geq 0$. In turn, decreases can only occur if the addition of the mutant caused the extinction of other species. With the bins around zero containing most counts (by a margin of more than one order of magnitude), it appears that many of the viable species feed on roughly the same parts of the food web that other species already feed on.

As shown here, species specific measures can be insightful for understanding the model dynamics, especially because they allow to evaluate the direct response of a mutant being added to the food web. They allow attributing a change in food web properties to a specific species – which is the direct cause of these changes in such a setting. Due to these statistics being collected during the regular iteration of the model and the high number of mutant species being added, sample sizes are large and thus simplify the detection of systematic effects in these observables.

Further analyses of species-specific measures will be conducted at a later point in this thesis (see Chapter II.5).

II.4

Food Web Perturbation & Resilience

The previous chapter looked at the general behaviour of the ECST model and how it leads to continuously evolving food web structures. With these observations as a foundation, the focus in this chapter will be on the broader goal of better understanding the reaction of this model to different perturbations and its ability to recover from them.

I will distinguish two kinds of perturbations: Through *local* perturbations, a microscopic part of the system is changed; these perturbations act on the species-level and can take the form of a species being added; this is studied in Section II.4.1. In contrast, *global* perturbations act on macroscopic components of the system, like the rate of resources made available through the external resource; this is the topic of Section II.4.2.

These studies of the perturbation response focus on the population dynamics on short time scales. In effect, the evolution mechanism will be disabled entirely throughout these simulations and species are only added to the system as part of the perturbation itself.

For the investigation of ecosystem resilience, the mutation mechanism is of central importance, as it represents the system's ability to reorganise after a perturbation; this is the subject of Section II.4.3.

II.4.1 Species invasion

Typically, new species are periodically added to the ECST model via the evolution mechanism described in Section II.2.2.2. As seen before, the addition of such mutant species may lead to a reorganisation of the food web. Such a reorganisation may be minute (with only the equilibrium biomass densities changing) or more drastic (with species becoming extinct and even triggering larger extinction cascades). In other cases, the newly added species is not viable and becomes extinct without leading to any notable change.

Scenarios like this can be associated with the process of a non-native species being introduced into an ecosystem. While this may have positive effects, species introductions have been found to pose a considerable threat to biodiversity and ecosystem functioning [*Ehrenfeld* 2010] – hence the typically negative connotation of species *invasion*. Species invasion is further found to be exacerbated by anthropogenic effects, either by plant or animal species that are directly introduced by humans, sometimes inadvertently, or by the reduction in habitat that drives species [*Romanuk and Zhou* et al. 2017].

In this section, I will analyse the effect of adding a single new species to an evolved food web and how this action – which can be understood as a small, local perturbation – leads to changes in the structure and properties of the system. Motivated by the above context, I will refer to the newly added species as the 'invading' species (or short: 'invader') and to those that exist before the invasion as 'resident' species; these terms are not used in a normative fashion here.

To study species invasion in an isolated setting, the regular evolution mechanism (randomly selecting a parent species, strongly mutating traits of the offspring) is replaced by deterministically choosing trait values of the invading species. In effect, all simulations in this section are governed by the population dynamics (Equation (II.2.11)) and the species extinction mechanism. By subjecting food webs with different properties to this invasion scenario, dependencies on the food web structure can be studied, including their robustness against species invasions: their capacity to integrate a new species without many other species becoming displaced.

In the bigger picture, invasions represent not only a perturbation, but are more or less equivalent to the regular speciation mechanism of the ECST model: While they are lacking a parent-offspring relationship, the study of the evolution mechanism in the previous chapter (see Section II.3.4) showed that mutations in the regular speciation mechanism are relatively large, thus allowing only for a weak similarity between parent and offspring traits anyway. As such, the investigation of species invasion also informs about the mechanism of structure formation in the ECST model and comparable models.

In the following, I will first describe the methodology in more detail. I will then qualitatively illustrate this method by looking at the effect of different values for the niche width s of the invader species and varying functional diversity of the initial states. Based on these qualitative observations, I will then quantitatively study larger sets of evolved food webs with varying degrees of functional diversity FD, connectance C, and trophic incoherence TI.

II.4.1.1 Methods

A species' effect on a food web depends on two factors: the species' trait values (m, f, s) and the state of the food web at the time the species is added to the system: all other species' trait values and the corresponding equilibrated biomass densities $\mathbf{B}$. This is true for the addition of mutant species during the regular operation of the ECST model as well as for this invasion scenario. However, the mutant species' traits depend on the randomly selected parent species (see Section II.2.2.2) and mutations from the parents' traits, while the study of invasions requires a more systematic choice of trait values: sampling from the three-dimensional trait space.

Extensively studying such a scenario quickly becomes computationally costly: Given the sensitivity on a species' m and f traits, the sampling resolution would need to be sufficiently high in those dimensions, say 10^2 samples each, already bringing the minimum number of simulations to 10^4. This needs to be repeated for a large enough set of s trait values and initial conditions. Hence, to make this approach feasible, individual instantiations need to be sufficiently fast and the number of instantiations needs to be reduced as far as possible.

To address the former point, the spin-up phase of the system is not repeated for every instantiation but a sufficiently evolved food web – iterated for at least 5 M steps – is loaded; the species is then added to this evolved food web. The population dynamics and extinction mechanism is then run until 4 k steps; ample time for biomass densities to equilibrate, which typically is a very fast process (< 100 steps).

To reduce the number of instantiations, the number of samples along the s dimension can be reduced quite a bit: As seen before, most species with large niche widths are not viable in evolved food webs; many species have niche widths $0.5 < s < 0.6$, so focussing on small niche widths is typically sufficient unless a direct comparison between generalist and specialist species is desired. If not stated otherwise, the niche width will be set to the minimum value of $s = 0.5$, giving the invader a benefit over the existing species in terms of specialisation.

For sampling along the body mass dimension, values $m \in [10^1, 10^5]$ are chosen. For an easier visual representation, the relative feeding distance $^m/_f$ is used instead of the absolute values of the feeding centre f, and samples are selected as $^m/_f \in [1, 1000]$. Note that this minimum relative feeding distance of $^m/_f = 1$ is smaller than the one used in the regular evolution mechanism with $^m/_f \in [3, 1000]$. These low-valued feeding distances lead to predominantly cannibalistic interactions, which are problematic in the regular operation of the model but are interesting to include in these invasion experiments.

The set of initial conditions consists of sufficiently evolved food webs. These can be chosen either randomly or by some process that takes into account their properties, like their size, connectance, functional diversity.

Another aspect needs to be taken into account at this point: The addition of the new species affects the system not only by the presence of the new species and its trophic connections, but also by the change in the biomass composition of the system.

For a new species with large m, its biomass density $B = n^{\mathrm{min}}m$ may become large, hence acting as a biomass-induced perturbation of the whole system; this is not desirable if the aim is to study the effect of the species traits alone. To avoid this perturbation, these simulations use a reduced initial number density of $n^{\min} = 2 \cdot 10^{-6}$ instead of the default $2 \cdot 10^{-4}$, such that the invader species' biomass is sufficiently small[48].

The initial number density $n^{\min}$ is equal to the extinction threshold of the system. Thus, this departure from the default value has the effect that the newly added species starts with a number density at least two magnitudes smaller than all existing species (which evolved in a food web with the default extinction threshold), and all extinctions are measured against the reduced extinction threshold. While this represents a relaxed scenario compared to the default, this approach of introducing a low-density species is preferable when compared to the alternative of inducing a biomass-related perturbation, which would intermix two kinds of perturbations: adding a species *and* significantly changing the biomass composition of the system.

Summarising the above, (i) the species will be added to a sufficiently evolved food web; (ii) the species will have explicitly chosen trait values rather than random ones; (iii) it will have a smaller initial number density $n^{\min}$ than all other species; and (iv) the system will have a smaller extinction threshold.

II.4.1.2 Response to species invasion depending on niche width

As mentioned before, the choice of invader niche width s is important as it strongly influences the viability of the species. In Figure II.4.1, results from an invasion experiment with different niche widths illustrate this. Each plot shows the $(m, {}^{m}\!/\!{}_{f})$ trait space of the invader species while the heatmap represents different observables, here: whether the invader was able to survive and the change in the number of species in the food web.

Figure II.4.1a shows that a larger niche width leads to fewer trait combinations in which an invader can survive in the invaded food web. With increasing s, survival is possible only for invaders with an increasingly high body mass m. Furthermore, the onset of the viable regime moves to higher values of m for increasing relative feeding distances ${}^{m}\!/\!{}_{f}$; this happens with a slope of about 1, which can be explained by the structure of the invaded food web: Invader species along the onset line will feed on the cluster of species at $m \approx 10^2$, which appears to be large enough to support an additional species beside the already existing ones. In effect, the new species is in competition with the other species in that regime.

Structures along the $m \propto {}^{m}\!/\!{}_{f}$ lines are very common in these invasion experiments and denote that the invading species feeds on the same cluster of species,

---

[48] As a technical remark, the invader species does *not* introduce new biomass into the system. The model is implemented in a way that enforces biomass conservation when new species are added by transferring that biomass from the parent species. For the invasion scenario, choosing a parent species does not make sense; instead, the biomass is transferred from the external resource, which is a very large biomass reservoir. In effect, this transfer does not act as a significant biomass perturbation, even if the invader's body mass m is large.
While one may argue that an invasion scenario should in fact introduce new biomass, the focus here is on analysing trait-induced effects. Subsequently, the aim here is to reduce all biomass-related perturbations, which is achieved by the above measures.

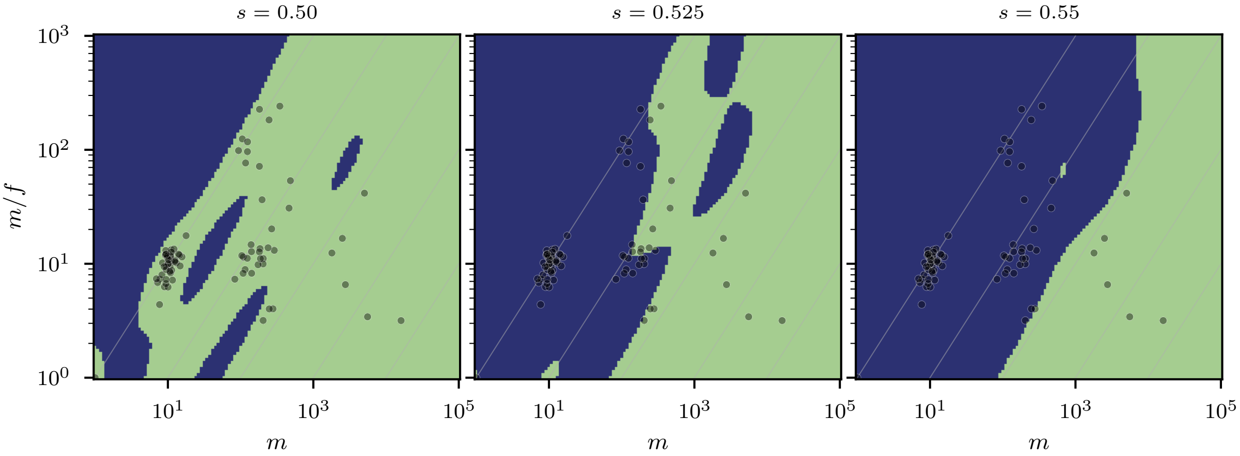

(a) Whether the new species was able to survive (green) or not (blue).

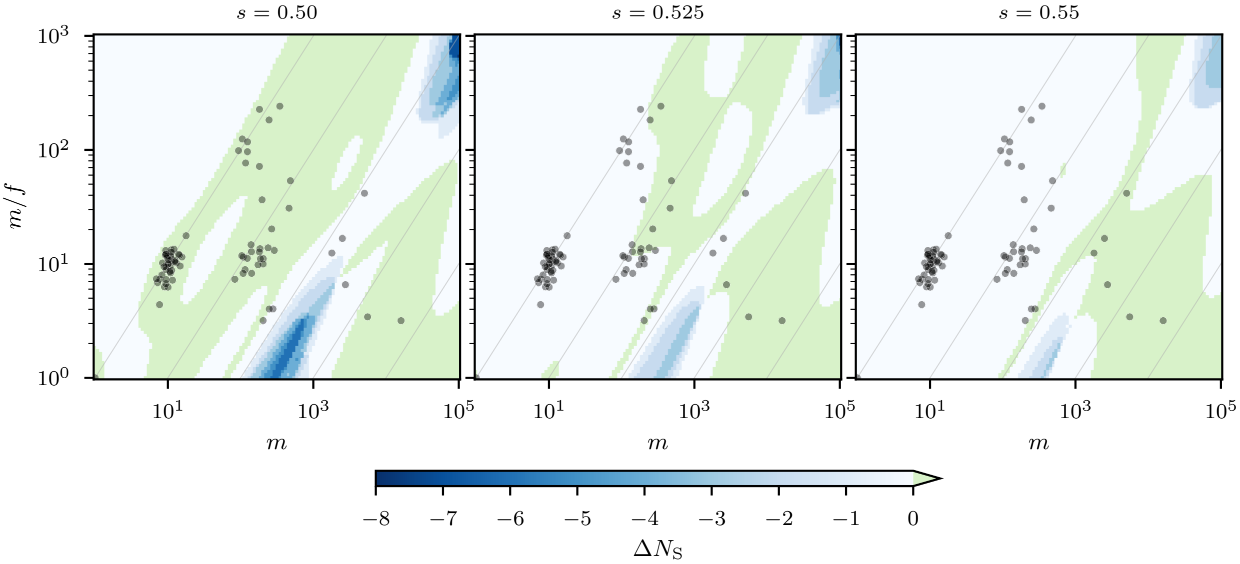

(b) Change in N_S induced by the addition of the new species. In the green regions, the new species with those trait values was able to coexist with all existing species; no extinctions occurred. In areas coloured in the lightest shade of blue, the total species count remained the same, meaning that *either* the invader or one of the previously existing species became extinct.

Figure II.4.1 Results of an invasion experiment for different values of the niche width s. The x- and y-coordinates denote species traits and the separate heatmaps show the effect on the respective observable when adding a species with such a trait combination. The black dots represent the species (and the external resource) in the evolved food web, allowing to estimate the interactions in the food web and with the invader species. To make locating prey species easier, grey lines indicate positions of constant feeding centre f. A predator's feeding centre f can be deduced by following such a line to $m/f = 1$, i.e. the x-axis; its prey species are then those with a body mass m of an approximately equal value.

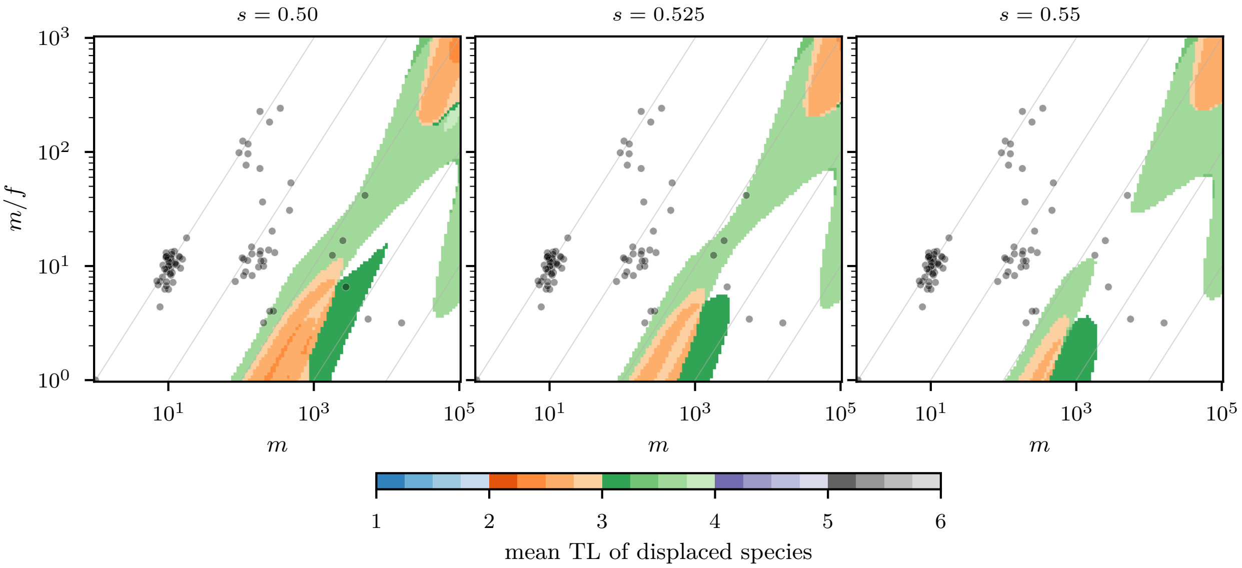

Figure II.4.2 The mean trophic level TL of displaced resident species. In white areas, no resident species became extinct. Note that the flow-based TL may produce values closely scattered around integer values; hence, areas with light shades are better associated with the next-higher integer value.

only with a different $(m, {}^{m}/_{f})$ combination. In other words, the diagonal lines are those of a constant feeding centre trait f and species along those lines are in direct competition for prey; see the caption of Figure II.4.1 for an explanation of how to locate the corresponding prey clusters for a certain predator.

Feeding on the same cluster can also be observed when comparing the cases for $s = 0.50$ and $s = 0.55$: In the former, the viable regime is larger because the invader is competitive enough to survive feeding on the cluster at $m \approx 10$ and on the external resource at $m = 1$.

Going beyond the mere survival of the invader, Figure II.4.1b allows to see the effect on the whole food web in terms of species becoming extinct. There are again stark differences when varying s, but this plot shows additional structure in the regions where the invader survived: While the invading species is able to coexist with resident species when feeding on the $m \approx 10$ cluster or on the external resource, competition becomes stronger for higher-valued feeding centres and, in some cases, leads to a larger number of extinctions.

To better understand which species become displaced by the invading species, Figure II.4.2 shows the mean trophic level of those resident species that became extinct as an effect of the invasion[49]. It shows that apart from the trait combinations

---

[49] To be exact: if the invader species became extinct as well, its TL is not taken into account for the calculation of the mean.

where large extinction cascades occur, it is the species with TL ≈ 4 (light green) that become extinct, i.e.: the direct competition of the invading species which can also be placed at TL ≈ 4. For higher s, these regions are much smaller, meaning that the specialisation of resident species gave them an advantage over the invader. In regions with higher number of extinctions, the mean TL is lower; this is caused by the invader feeding on the $m \approx 10^2$ cluster and affecting more and more species with decreasing TL (species in that cluster that have large $^m/_f$ values and hence a lower TL), in turn reducing the mean value.

In summary, if the invader is able to displace resident species it is by feeding on the $m \approx 10^2$ cluster in a more competitive way than resident species. The large cascades result from its top-down control on that cluster, where a successful predator species reduces the equilibrium biomass densities of the layer below and thus pushes species below the extinction threshold.

As can be seen from these qualitative observations, there are many mechanisms at play when an invader species is added to an evolved food web and the impact the new species has strongly depends on its trait values. In general, invaders seem to have less impact when increasing their niche width s as this leads to a reduction of its viability and a subsequently reduced effect through species invasion. This reduction in impact is already strong at $s = 0.55$, and quickly increases with higher values (see Figure B.4), thus further justifying the choice of a single s value for the following experiments, where the focus is not on comparing specialist and generalist species with each other.

II.4.1.3 Response to species invasion depending on initial state

After having varied the niche width s of the invader species, this section will qualitatively point out the differences between individual invasion experiments and how these depend on the choice of initial conditions. Furthermore, specific scenarios are studied in more detail to understand the underlying cause of the invasion response.

To that end, three evolved food webs with spread-out values for the functional diversity ('low': FD = 2.98, 'medium': 3.68, and 'high': 4.04) were chosen for the starting point of the simulations. While studying only these three food webs will not allow to make claims about the general dependencies on the functional diversity, these choices of FD lead to structurally different networks which highlight potential differences in the effect of an invader species.

Functional diversity

To start, how is the functional diversity itself affected by the invading species? The change in FD follows directly from its definition via the envelope of all feeding kernels, Equation (II.3.8): if the invader is able to survive and has a feeding centre f in a previously unpopulated area along that dimension, FD will increase. These unpopulated areas in f space are by definition larger if the invader is introduced into

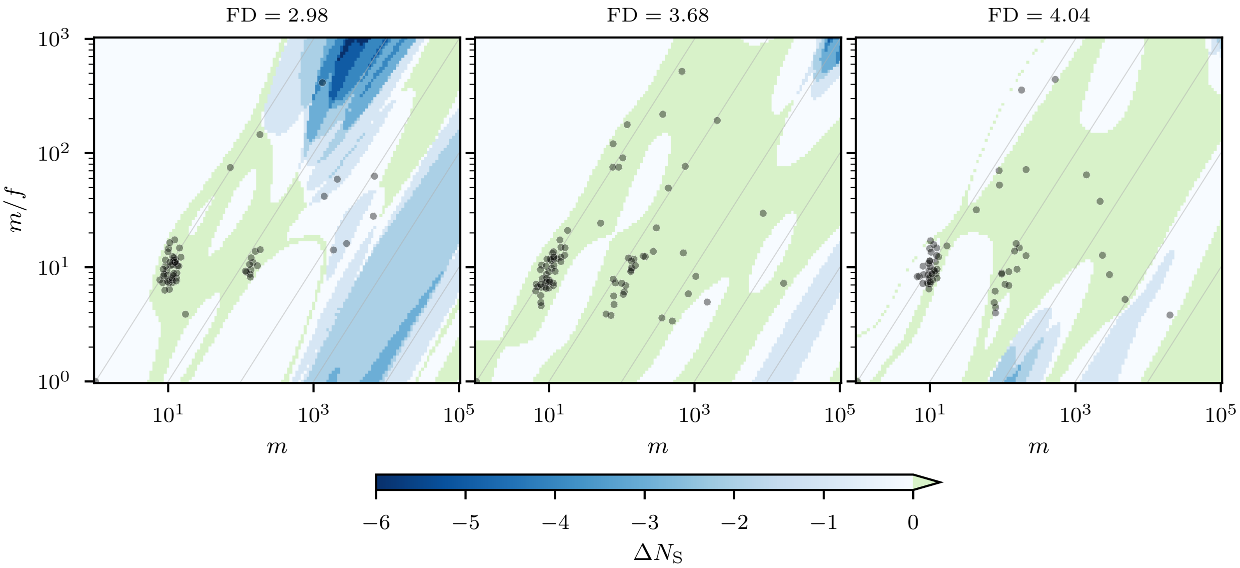

Figure II.4.3 Change in N_S induced by the addition of the new species.

a low FD food web, where an high f predator will always lead to an increase in FD. If the invader causes an extinction cascade, FD may decrease; the magnitude of functional diversity loss depends on which species are affected and whether they are part of a cluster where other species may compensate for the loss. The simulation results match these expectations, see Figure B.7.

Species extinctions

Figure II.4.3 shows the change in N_S. As can be seen there, the initial state has a strong influence on the patterns resulting from the invasion: The low FD food web exhibits a comparably small region of coexistence and at the same time includes large extinction cascades for high m or high m/f values. Additionally, a large area of the low m/f regime does not seem to support coexistence, because the invader does not survive there (cf. Figure B.6). In this particular case, said region is uninhabitable for the invader because it competes with higher m predators, which are at an advantage in that situation.

In the medium FD scenario, the invader was able to coexist in a much larger region. Incidentally, the coexistence regime seems to be spanned by the location of resident species, particularly in the low f part of trait space. This means that an invader with similar traits to a resident species was not only likely to survive itself, but that it did not exert sufficient additional competition on resident species to displace them. The coexistence region being comparably large suggests that resident species were well adapted to each other and could be regarded to be in a more robust configuration than in the low or high FD settings. The underlying mechanism is straight-forward: An invader species that differs only minutely from existing species does not bring new feeding relations into the food web, hence having a smaller impact on the already-equilibrated network of resident species.

Network structure

How does the change in species number affect the network structure, specifically the connectance and the trophic incoherence?

Figure II.4.4a shows the change in connectance of the three food webs, i.e. the fraction of realised links in the extracted interaction network (Section II.3.5.3). Compared to the previously studied observables, the connectance shows a noticeably more intricate structure which includes features with a higher frequency in trait space. In the low and medium FD scenario, regions of strong connectance loss coincide with those of large species loss. However, the opposite is the case for the high FD case, which shows an increase in connectance throughout almost the whole domain, despite a similar initial connectance.

Unlike the previous patterns, the reasons for the connectance change are somewhat more elusive. In general, if the invader coexists with the resident species, an increase in connectance requires a sufficient increase in the mean degree $\langle k \rangle = {}^{L}/_{N}$ of the network. Alternatively, if the invader causes a smaller network size, an increase in connectance would denote that the lost species had a smaller than average degree (and analogously for a decrease in connectance).

Again, the trophic level of displaced species (Figure II.4.4b) can be insightful in this context, as it allows to estimate the number of connections a displaced species had: Most links are from species on TL ≈ 3 feeding on species on TL ≈ 2; the primary producers typically have few connections as they feed predominantly on the external resource and predators at TL ≈ 4 and above again have fewer links because the number of prey becomes increasingly low. The low degree of primary producers explains the pattern of $\Delta C < 0$ for the $f = 1$ region. A notable exception with $\Delta C > 0$ is the region close to the external resource in the high FD case, where the connectance increase is caused by one species with fewer links than the invader being displaced. The invader species being closer to the primary resource also leads to additional species feeding on it, thus increasing its degree.

In turn, the high degree of species feeding on the primary producers leads to the increase in connectance in the regions around the $f = 10$ and $f = 100$ lines. Note that this requires that the invader actually survives (unlike in parts of the low FD scenario). In case of extinctions, the connectance change is additionally influenced depending on the trophic level of the displaced species: If species loss is on low trophic levels, the connectance will still increase (e.g. in the strongly positive regime for high FD), while losses at the well-connected TL ≈ 3 typically lead to a connectance decrease.

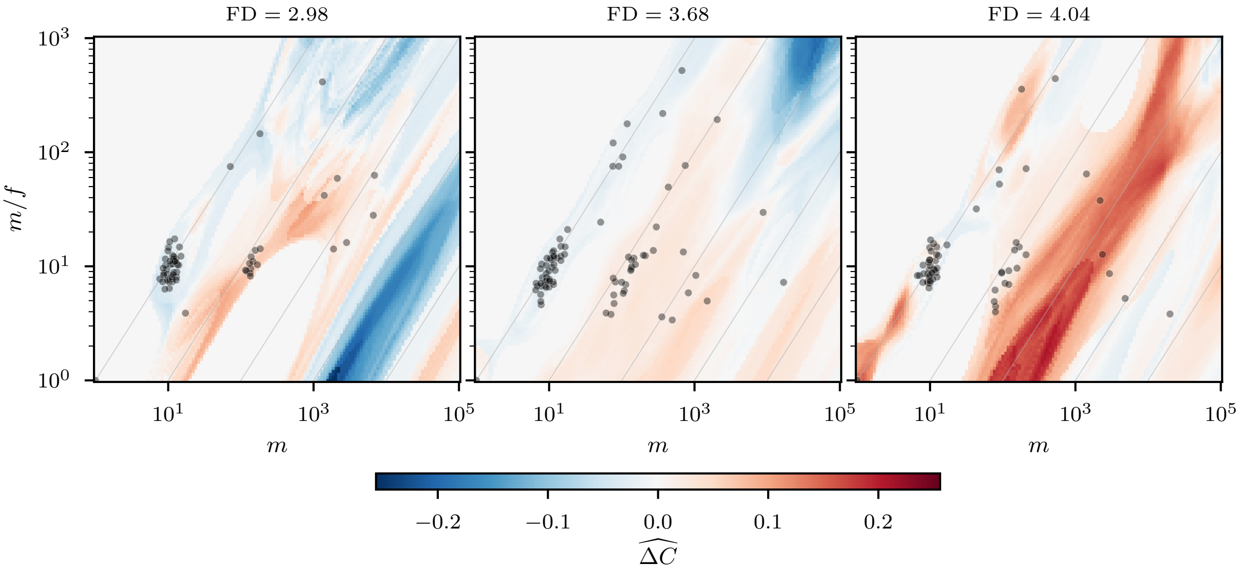

(a) Relative change in connectance $\widehat{\Delta C} = (C(t_1) - C(t_0))\,/C(t_0)$ induced by the addition of the invader species. The absolute initial connectance values $C(t_0)$ are 0.160, 0.196, and 0.156, respectively.

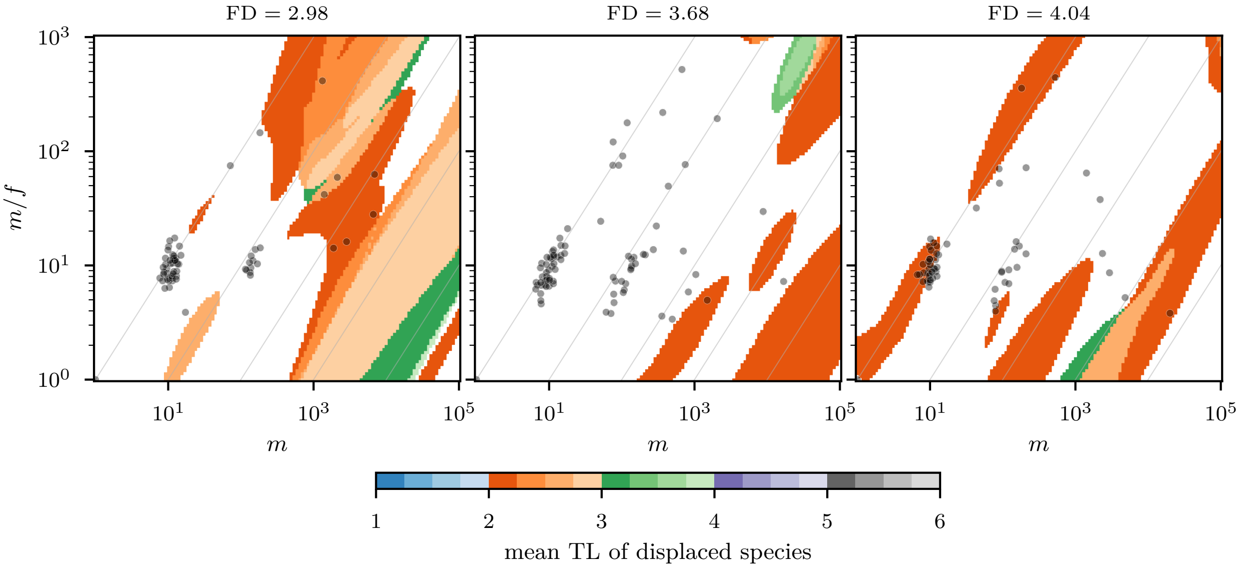

(b) Mean TL of displaced species (not including the invading species). White areas are those where no resident species became extinct.

Figure II.4.4 Inspecting the trophic level of displaced species can help explain the change in connectance (see main text).

The above walk-through addresses the large-scale patterns, but not the initially mentioned small-scale structures in the connectance. To better understand these, the thresholding algorithm has to be put into focus again: On top of the proximate effects from the change of links in the food web, the network generating algorithm imposes some additional changes in C: Recall that the network is extracted by including only those edges of the functional response matrix g with an interaction strength larger than 75 % of the mean interaction strength for a particular predator (see Section II.3.5.3). Subsequently, small changes in the species' interactions may cause many links to move beyond or below the threshold, strongly changing the number of links in the extracted network from which C is computed. While this depends on the distribution of link strengths, clusters of species typically show similar link strengths to other species and hence may lead to comparatively abrupt changes in the number of links, thereby forming the small-scale patterns in Figure II.4.4a.

How does the trophic incoherence TI change upon species invasion?

One expectation would be that connectance and increased trophic incoherence are somewhat correlated. Specifically, in scenarios where additional links are added in central parts of the food web, trophic incoherence can be expected to increase, as additional links will more likely be on a wider range of trophic levels. Comparing Figure II.4.5a to Figure II.4.4a, this appears to be the case for many but not all regions: For instance, regions of strong connectance loss are typically also regions of reduced TI. This can be explained by looking at the composition of species that got displaced (Figure II.4.4b), which typically show displacements with TL ≈ 3, i.e.: species that frequently feed on multiple trophic levels and thus contribute to a higher incoherence.

Some areas show similar overall positive or negative values in both C and TI, but additional effects seem to be present there as well. Specifically in the high FD case, there is a dominant increase in connectance, while TI appears hardly affected. One reason for TI not following C is that only the latter is affected by the thresholding procedure (Equation (II.3.12)), which can cause rapid jumps in such an invasion experiment. Thus, TI not only includes more information in terms of the interaction network's edge weights, but is additionally not prone to fluctuations caused by the thresholding. These considerations are relevant in the mentioned high FD case, where high m invaders lead to an increase in connectance but a decrease in trophic incoherence: The connectance increase suggests that the invader brings many new links to the food web, predominantly to the $m \approx 100$ cluster. At the same time, these links appear to make the trophic level differences more coherent, indicating that the new links have trophic level differences closer to 1 than other links in the network and that these links are strong enough to sufficiently contribute to TI.

As a brief summary, the presented network measures can help understanding the effect of the invader on the network structure. The trophic level can be insightful not only for describing the feeding relations of the invader, but also in explaining its effect on the rest of the network. However, care has to be taken to not con-

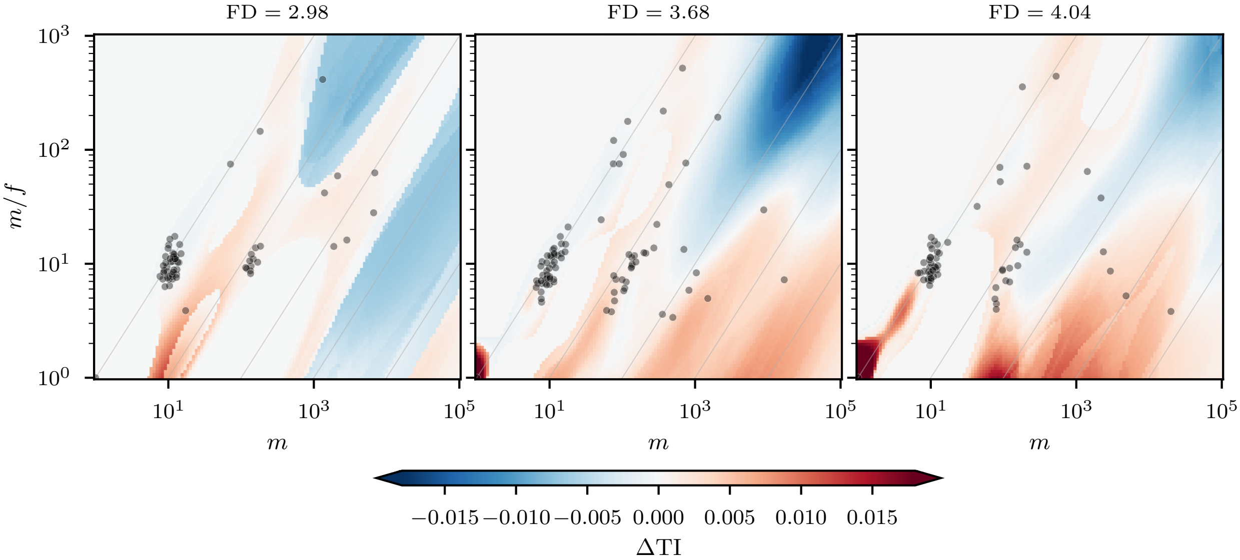

(a) Change in trophic incoherence $\Delta\mathrm{TI} = \mathrm{TI}(t_1) - \mathrm{TI}(t_0)$. Values before the introduction of the invader were $\mathrm{TI} \approx 0.036$, 0.062, and 0.047, respectively.

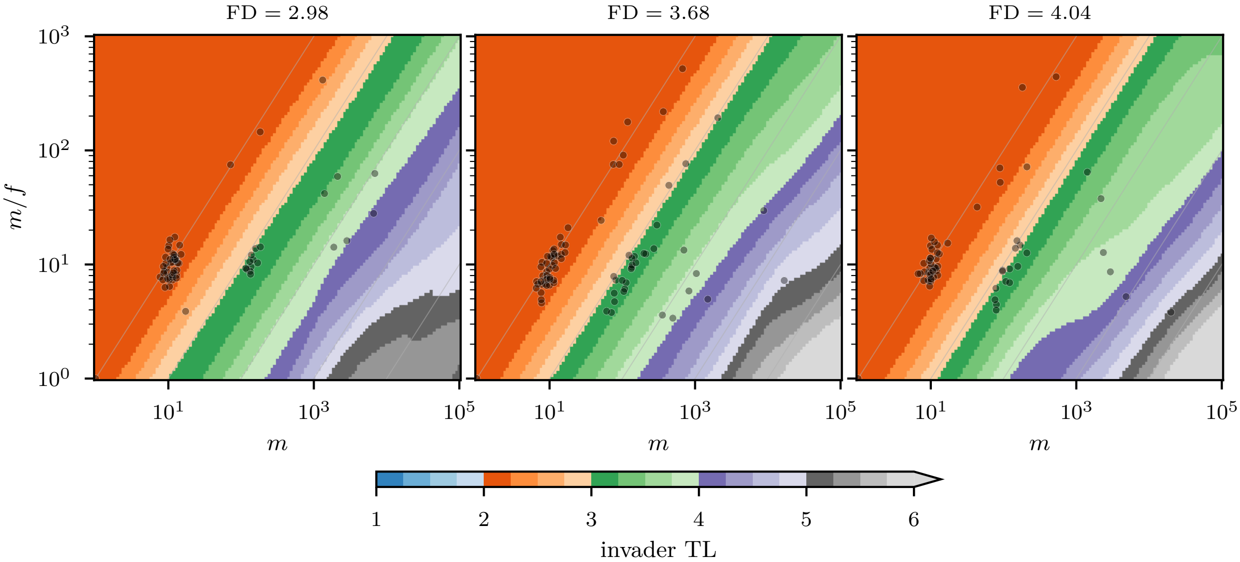

(b) The trophic level of the invader species. As the trophic level changes depending on food web composition and biomass densities, the value shown here is the mean value over the invader's life time. For the same reason, the region $f < 1$ shows values $\mathrm{TL} \approx 2$: the invader feeds primarily on the external resource, albeit with an increasingly small rate which quickly leads to its displacement by the resident species.

Figure II.4.5 Change in trophic incoherence TI and trophic level TL of the invader.

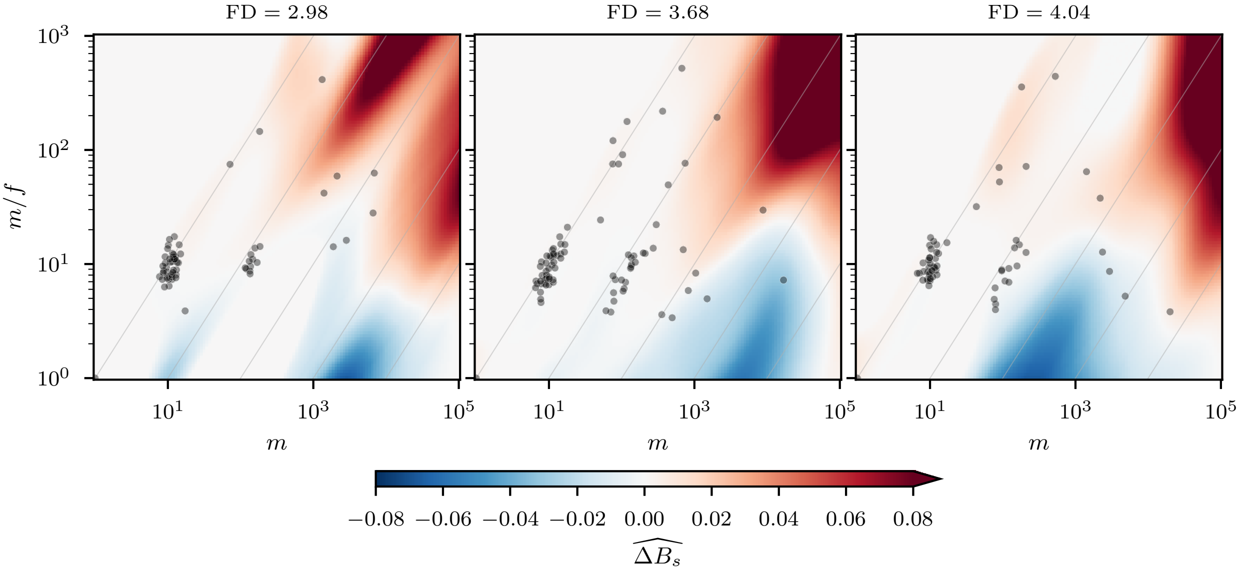

Figure II.4.6 Induced relative change in species biomass, $\widehat{\Delta B_\mathrm{S}} = (B_\mathrm{S}(t_1) - B_\mathrm{S}(t_0)) / B_\mathrm{S}(t_0)$, with t_0 and t_1 referring to the beginning and end of the invasion experiment, respectively. Note that maximum values go up to approximately +22 % but colour values are cropped in order to make negative biomass change visible.

flate actual structural changes with artefacts from the computation of the network measure as in the case of the connectance.

Biomass change

Apart from ΔN_S and the network measures, it is worthwhile to look at biomass-related measures of the system as these are indicators for changes in ecosystem functioning.

Figure II.4.6 shows the relative change in species biomass $\widehat{\Delta B_\mathrm{S}}$, which is overall positive in the high m/f regime and negative in the low m/f regime. The structure in trait space seems to be linked to the position of resident species in trait space, which often are at the boundary between a region of little change and the onset of positive or negative $\widehat{\Delta B_\mathrm{S}}$. In some cases, the resident species form 'incisions' into the regions of large change, e.g. the species with high m in the FD = 2.98 scenario, suggesting that their presence strongly influences the response of the food web and in turn the equilibrium biomass densities, which sum to B_S.

Notably, the change in species number seems not to correlate strongly with the change in biomass: In the low FD scenario, most extinctions occurred for invaders with $f < 10$, while the strongest biomass change is observed for $f > 10$. Similarly, the medium FD case shows a region with biomass loss for low $^m/_f$ that is in large parts a coexistence region.

Supporting observables to look at in this case are the change in biomass flow from the resource and between species (cf. Figure B.8), which can be used to characterise the relative strengths between the TL ≈ 2 cluster (which dominates consumption on the resource, Φ_0^{out}) and the species on higher trophic levels (which determine intra-guild consumption, Φ^{ig}). The regions with high biomass increase typically have an equivalent region with a decrease in Φ^{ig} (and vice versa for the increase in biomass). No matching regions are observable in Φ_0^{out}, which is dominated by an alternating pattern of increases and decreases, strongly dependent on the invader's f trait.

How can this seemingly systematic change observed in species biomass (increase for high $^m/_f$ invader, decrease for low $^m/_f$) be understood?

In principle, questions like these are approachable by the framework of *trophic cascades* [**??**], which attempts to describe the effects that changes to one trophic level in a food web have on the rest of the food web. As a simple example, consider a food chain of resource and producer. The introduction of herbivores, feeding on the producers, would exert *top-down control* on the rest of the chain: reduce the producer population, thus reduce predation on the resource and increase resource abundance. Of course, this example assumes an oversimplified ecology: a (linear) food chain, discrete trophic levels, as well as a linear, non-saturating functional response.

Wollrab et al. expanded this idea for more diverse topologies [*Wollrab* et al. 2012]. Specifically, they allowed the existence of a second branch, motivated by the observation that food webs often exhibit multiple energy pathways, e.g. originating from groups of edible and defended plant species. The authors then analytically and numerically studied the response of these generalised food web topologies to nutrient enrichment and increased predator mortality. They found that the food web response depended primarily on the lengths of the two energy pathways, whether they were odd- or even-numbered, and whether a generalist high-level consumer existed that fed on both chains.

In the context of the ECST modes, the authors' generalisations appear to be applicable to the patterns seen in the change in biomass flow from the resource (see Appendix B), where the three studied food webs show alternating patterns of increase or decrease in $\widehat{\Phi_0^{\mathrm{out}}}$. These regions seem to depend primarily on f, alternating between positive and negative values depending on the trophic distance from the external resource. Generously simplifying and extrapolating the pattern seen in the three cases, it appears that an odd-numbered distance leads to an increase in flow while even-numbered distances lead to flow reductions. Of course this is only a single observation at this point and it is superimposed with other effects like species extinction; while reminiscent of the mechanisms of top-down control, a systematic study of this pattern is needed (and done later on) to further support a link to the framework by *Wollrab* et al. [2012].

While conceptually powerful and seemingly applicable to the change in biomass flow from the external resource, using the framework to understand responses of species' properties to the invasion was not possible. The main difficulty was to reliably map clusters in the evolved food webs to the functional groups and chains the authors presume. For one, species clusters in the ECST model cannot be delineated clearly from each other due to their diverse traits and feeding relations. Furthermore, even if grouping clustered species into functional groups, the species in one group would frequently feed on multiple other groups, resulting in topologies that are not compatible with those of the framework (at most two chains, single feeding links unless for a top predator that connected two chains). Disregarding the weaker links would be an option to arrive at simpler topologies, but it is unclear how structurally important these weaker links are for the response to species invasions.

With these conceptual frameworks seemingly inapplicable to the situation at hand, it is worth reflecting on the actual aim for studying this particular scenario. After all, the invasion experiment leading to a biomass increase in the high ${}^{m}\!/\!{}_{f}$ regime is an *observation* from the simulations: the numerical solution for the change in equilibrium biomass densities according to the dynamic equations (Equation (II.2.11)) of this system and depending on different forcings (invaders). On a purely mechanistic level, this alone answers the question of *why* this change occurs.

Hence, the understanding that can come from studying this scenario has to be on a more abstracted level: Which changes are imposed on the food web by addition of the invading species? And which properties of the invader are relevant for this change? To investigate these questions further, individual invasion scenarios for representative cases offer a more detailed perspective.

Figure II.4.7 shows simulation results from two separate invasions into the evolved food web with FD = 4.04. As in previous figures, the individual species are located in the $(m, {}^{m}\!/\!{}_{f})$ trait space. In addition, feeding links between species are drawn according to the aforementioned thresholding algorithm (**??**). Node and edge colours are used to represent the absolute change in biomass density and biomass flow, respectively.

The particular invader traits were chosen in regions of trait space that showed a pronounced biomass increase or decrease. To not include secondary effects from species extinctions, only traits from coexistence regions were used. Furthermore, relative feeding distances ${}^{m}\!/\!{}_{f} \geq 3$ were used to reduce cannibalistic links and let invaders be in the same trait value range that is attainable via the regular evolution mechanism of the ECST model.

Let us first consider the case of the high ${}^{m}\!/\!{}_{f}$ invader (left subplot). The biomass change for the invader is very pronounced in this case and it holds approximately 20 % of total species biomass. Together with the comparably small losses and gains in other species this means that the additional biomass that led to the increase in B_S is primarily located in the invader species. Owing to its high m, the invader's number density is still very low – in the regular iteration of the model, it would be

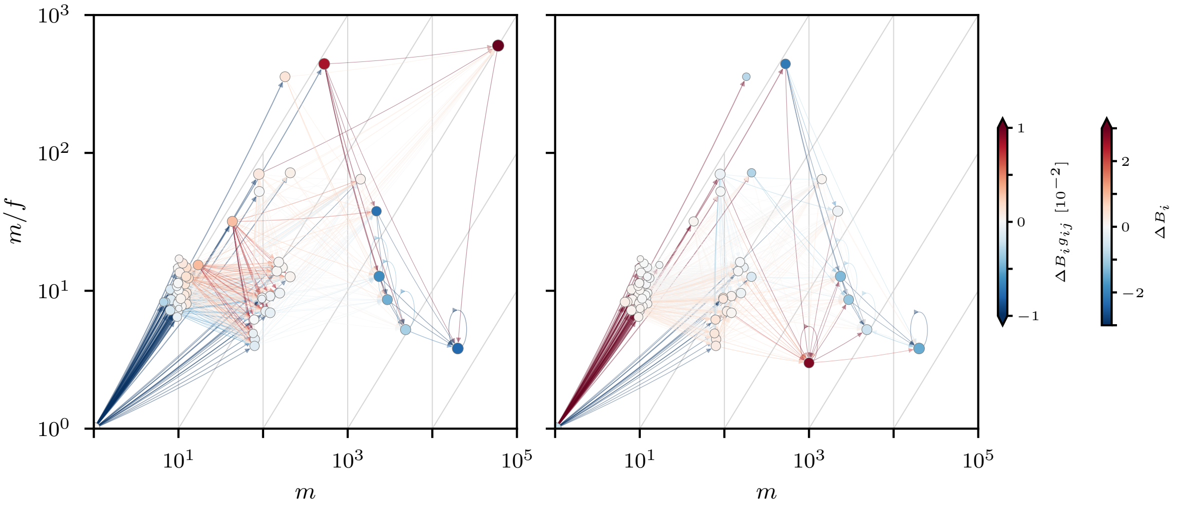

Figure II.4.7 Response to invasion visualised using the thresholded network representation. The invader species are those at $(6 \cdot 10^4, 600)$ **(left)** and $(1000, 3)$ **(right)**.
Unlike previously shown networks, species positions correspond directly to trait values. The node colour shows the absolute change in individual species biomass density $\Delta B_i = B_i(t_1) - B_i(t_0)$, node size hints at the absolute value of the species' B_i. Edge colour denotes change in biomass outflow over each link, $\Delta\Phi_{ij} = \Delta(B_i g_{ij})$.
The ΔB_i values of the invaders are approximately $+14$ (far exceeding the colourbar range) and $+3$, respectively. The total species biomass before the introduction of the invaders was $B_\mathrm{s}(t_0) \approx 68$. Note that the change in biomass flow from the resource is only large compared to the intra-guild flows; the total relative change is $\widehat{\Delta\Phi_0^\mathrm{out}} \approx -0.2\,\%$ and $+2.5\,\%$, respectively.
For a fully-connected network representation of the same scenarios, see Figure B.9.

only barely above the extinction threshold. Strong biomass loss occurs mostly in the five other species with high f; losses in the rest of the food web are rather small. In the $m \approx 10$ cluster, a shift occurs, where higher m species grow in biomass and lower m biomass is reduced. The changes in biomass outflows from species is of the same sign as their change in biomass – more abundant species lead to a larger biomass flow as a direct consequence of the foraging terms of the system equations. Interestingly, the $m \approx 100$ cluster still acquires a lot of biomass directly from the resource, despite its relatively high distance from it.

While the effects on the biomass densities and flows are diverse, they can indirectly all be attributed to an overall increased competition in the food web (cf. Figure B.10) triggered by the invader. The increased competition in turn leads to reduced functional responses[50], and reductions in biomass flow; however, the res-

---

[50] Recall the competition term $\sum_l c_{il} B_l$ in the denominator of the functional response, Equation (II.2.4), where c_{il} denotes the overlap of the respective predator's feeding kernels.

ulting changes for the equilibrium biomass densities are again dependent on all other species, thus eluding an intuitive explanation.

Which properties of the invader make it trigger such a strong change in competition that affects multiple species clusters? Basically, this is due its high body mass m leading to allometric advantages through reduced respiration rate ($x_i \propto m_i^{-1/4}$), reduced handling time ($h_i \propto m_i^{-3/4}$), and increased individual attack rate ($a_i \propto m_i^{3/4}$). Reduced respiration means it has lower living costs compared to low m species and the changes in handling time and attack rate increase the functional response for all its predatory interactions. In combination with it feeding not on high-level species but primarily on the $m \approx 100$ cluster in the lower levels of the web, it has a more immediate access to resources, allowing it to grow to sufficiently large abundances to persist. As seen before, these properties of the invader have repercussions on the macroscopic scale, namely that the system holds more biomass and has reduced losses.

The allometric mechanisms act in the opposite direction for the other invader (right subplot in Figure II.4.7): Its lower body mass make it have a higher respiration rate and handling time and a lower attack rate when compared to competing higher m predators. The invader still accrues a large amount of biomass (but much less than for the high m invader); in combination with the losses in the high m species, the overall biomass change is negative.

Additionally, the invader causes increased competition on the $m \approx= 100$ cluster, reducing the functional responses of higher-level predators. In effect, biomass flow from the resource to species with $f \approx 1$ strongly increases, while the flow to the indirectly-feeding $m \approx 100$ cluster slightly decreases due to the increased competition with the $m \approx 10$ cluster. Ultimately, this results in a diversion of biomass flow via the low m cluster instead of via the next-higher cluster, which explains both the increases in Φ_0^{out} and Φ^{ig}.

As shown in this section, the effects of the invader species can be diverse. It appears that in these scenarios – in absence of extinction events – the observed macroscopic changes are triggered primarily by a changes in competition which then cascade through the food web. Yet, understanding changes for individual species as attempted with Figure II.4.7 remains difficult, even (or especially) in such a detailed visualisation.

Summary

The above invasion experiments are an example of how diverse the reaction of key observables to species invasion may be and that these depend strongly on the specific food web that is being invaded. Also, they illustrate how difficult it can be to isolate underlying causes for a particular change in food web structure of biomass distribution.

Nevertheless, a few summarising observations of patterns and presumably acting mechanisms can be made:

- Survival and coexistence appears to be less frequent for low-level invaders that deviate too much from existing species' traits. This may be caused by low-level species being in an optimised configuration, thus making it more difficult for invaders with non-optimal traits to persist.
- In more functionally diverse food webs, an invader appears to have a lower overall effect on species composition (more coexistence regions, less extinctions), caused by the food web already having accommodated to more diverse feeding relations. In effect, a new species being introduced acts as a smaller perturbation of the resident species than in less diverse food webs.
- Allometric mechanisms seem to play a dominant role in the macroscopic changes to food webs, especially for invaders with high body mass.

While the above studies are initial indicators for these patterns in invasion response, they remain inherently qualitative due to the small number of compared food webs and the complex interplay between species depending on the initial state. The following section will attempt a more systematic study of the system response to invaders.

II.4.1.4 Systematic analysis

For a more systematic study of these invasion effects, the above experiments were carried out starting from a larger set of evolved food webs. For each of these initial conditions, the effect of adding the invader species was assessed by computing an effective observable from the whole space of trait combinations $(m, {}^{m}\!/_{f})$ and putting it into relation to a property of the initial food web (its functional diversity FD, connectance C, and trophic incoherence TI). The resolution of trait combinations was coarsened to 16 values along the ${}^{m}\!/_{f}$ dimension and 26 values along the m dimension, corresponding to 5 sampling points per decade in trait space. Given the size of the more prominent features (see previous analyses), the lower resolution reduced computational load and was still sufficiently high to capture the general effect of adding a species of a particular trait combination to an evolved food web. While sampling always occurred for exactly the same trait combination, the variance between food webs should average out any effects that depended on trait values much smaller than the sampling resolution.

A number of observables were computed to quantify the effect of the invader species depending on a particular initial state. These quantities were computed from the all samples in the space of $(m, {}^{m}\!/_{f})$ trait combinations, for instance: the proportion of samples in which the new species was able to coexist with resident species. Other observables were the mean change in biomass flow, or the proportion of states where large extinction cascades appeared or the system's biomass changed beyond some threshold.

Initial food webs were selected from a pool of 2048 evolved food webs[51]. Given the distribution of food web properties (Figure B.14), uniform sampling would have excluded many of the rare initial states with low or high values of these properties and would have made it more difficult to see dependencies on those properties. To address this issue, initial states were selected via binned sampling: The properties of interest of the 2048 available food webs (FD, C, TI) were computed and then sorted into bins of a certain width. From each of these bins, a random selection of food webs was chosen as initial conditions. In effect, the distributions of food web properties seen in this analysis are distorted towards a uniform distribution; therefore, these distributions should not be used to make statements about the frequency of such a state actually appearing in the regular iteration of the ECST model.

Typically, about 200 initial states were selected for a simulation run. For each initial state, simulations for $16 \cdot 26 = 416$ trait combinations were performed according to the procedure explained above, leading to upwards of 80 k simulations. Due to my implementation of the model and evaluation routines being optimised for long simulations rather than these kinds of invasion experiments, the bottleneck for this

---

[51] For a detailed description of how these were generated, see Section II.4.2.1. At this point, it is only relevant that these food webs were simulated much longer than the spin-up phase and that they were sufficiently diversified.

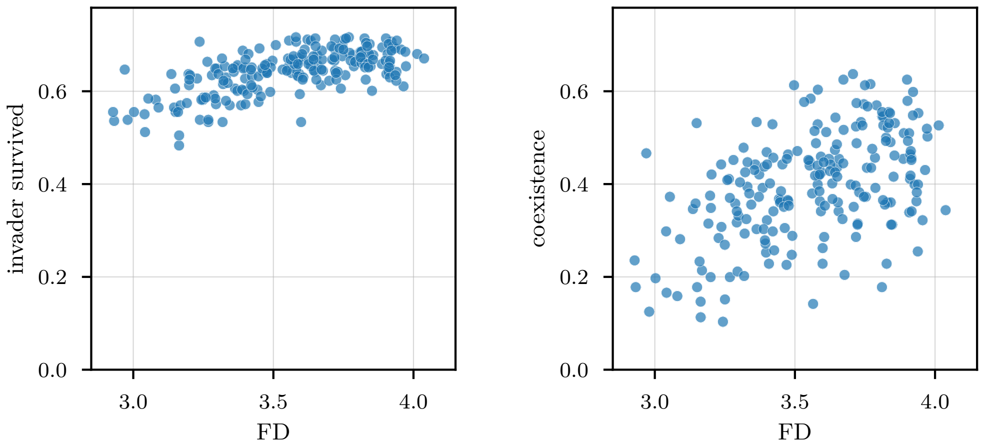

Figure II.4.8 Fraction of trait combinations for which the invading species was able to survive **(left)** or coexist with resident species **(right)**.

approach was in the number of simulations, becoming computationally unfeasible for more than 100 k simulations. With many observables showing a large variance caused by the different initial states, more simulation data would be desirable; but were not attainable in the scope of this thesis.

To not suggest a false sense of quantifiability, I will forego regression modelling or statistical significance testing (which is often problematic when modelling ecological systems, see [*White* et al. 2013]) and will focus on the qualitative differences observable from this analysis. This naturally restricts statements to the strongest effects – however, this is in line with the aim to better understand the *general* behaviour of the model.

In the following, I will first study the relation between invasion response and functional diversity, and then turn to the relation towards the network measures. With an apparent correlation between TI and C of the initial states (see Figure B.14), the analyses of these two measures is combined.

Dependencies to functional diversity

The relationship between functional diversity FD and change in food web properties caused by the invasion was studied by applying the above method to a set of food webs with FD values roughly in the range $3 \ldots 4$. Through binned sampling (bin width FD $= 0.1$, up to 24 samples per bin), 209 food webs were selected as initial states.

Figure II.4.8 shows the results from these invasion experiments depending on the functional diversity of the respective initial state. Both measures show the fraction of trait combinations that fulfilled a certain condition: the survival of the invader species or the coexistence with resident species, respectively. While both

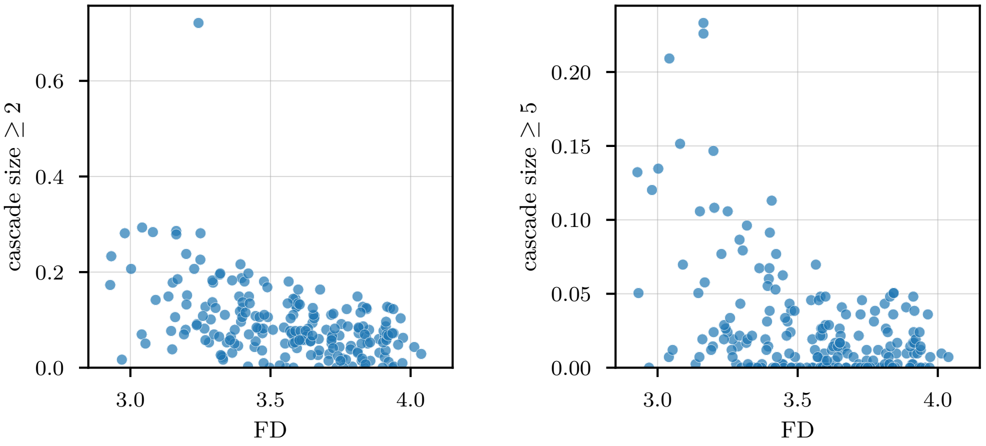

Figure II.4.9 Fraction of trait combinations in which an extinction cascade of a certain size was triggered. The cascade size is measured towards the food web size before the invader was introduced, hence a cascade size of two means that *at least* two resident species became extinct as a result of the invasion; if the invader survived, three resident species were displaced. See Figure B.15 for the number of species before invasion.

observables exhibit a large variance, they indicate a trend towards increased invader survival and coexistence for higher functional diversity, which is in accordance with the qualitative observations made above (Figure II.4.3). As mentioned there, the increase in coexistence can be explained by the wider spread of resident species in trait space for higher functional diversity, leading to more regions where the addition of the invader does not fundamentally change the feeding relations of the food web but merely causes a new distribution of equilibrium biomass densities. Thus, the invader may displace a resident species primarily through direct competition. The results shown in Figure II.4.8 support this explanation.

However, given the large variance in the coexistence, the specific state of the food web still appears to be very influential – the functional diversity measure FD alone does not capture how well an invader species can be integrated into an evolved food web. It seems that there are many food web configurations that lead to extinctions cascades. This is despite the fact that the invasion experiments use a reduced extinction threshold, meaning that all resident species are at least two orders of magnitude away from the extinction threshold. Subsequently, extinction cascades are unlikely to be the result of the biomass transients after the introduction of the invader, but reflect substantial changes in equilibrium densities caused by the addition of the invader.

Additional evidence for such a mechanism comes from the investigation of the trait combinations for which extinction cascades occur, see Figure II.4.9. While again showing a high variance, the food webs with low FD are those where larger regions in trait space exhibit small or larger extinction cascades. In turn, for

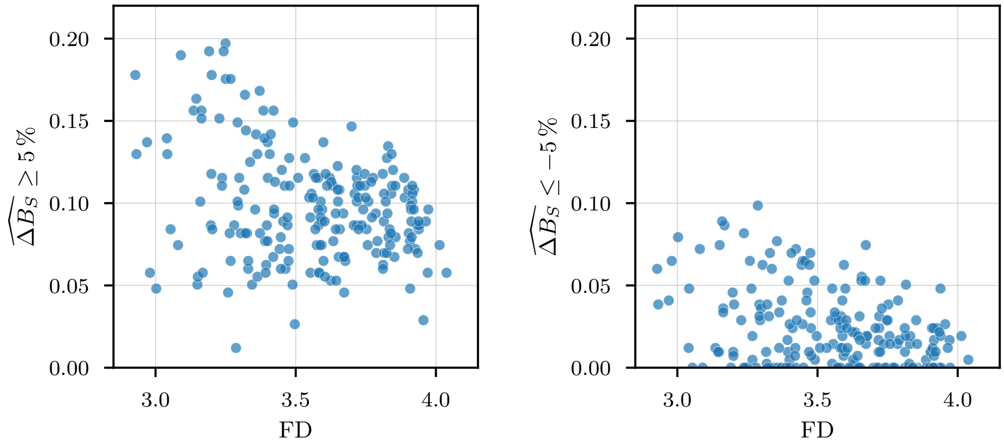

Figure II.4.10 Fraction of trait space with strong positive **(left)** or negative **(right)** relative changes in species biomass.

FD > 3.4, there are increasingly many initial states in which no cascades occurred. This is in accordance with the explanation given above, which relates the effective strength of the perturbation caused by the invader as based on the actual amount of changes in the interaction structure that it may introduce; for higher FD, the changes are smaller, hence acting as a weaker perturbation of the food web and leading to a typically lower occurrence of large extinction events.

Finally, turning to the biomass response of the system, Figure II.4.10 displays the size of regions where species biomass increased or decreased. As argued before, a strong biomass increase is to be expected in the high $^{m}/_{f}$ area and a decrease in the low $^{m}/_{f}$ area. By counting samples where the biomass change exceeds a certain threshold, these measures are an attempt to capture the extent of only those regions with a considerable effect on the species biomass densities and disregard those with small changes.

However, as before, the variance between observations is very high and shows no strong dependency on FD. There appear to be more regions with a large biomass increase for lower values of FD than for higher values, but these effects are too small for further consideration.

What can be said, comparing the two subplots, is that areas with a strong biomass increase are generally larger than those with a decrease – this of course depends on the choice of threshold, but is also the case for lower thresholds (not shown). Furthermore, these regions of biomass increase are not only larger, but also lead to a positive biomass change when averaging over the whole trait space (see Figure B.16), allowing the statement that an invader *on average* rather leads to an increase in species biomass than to a loss.

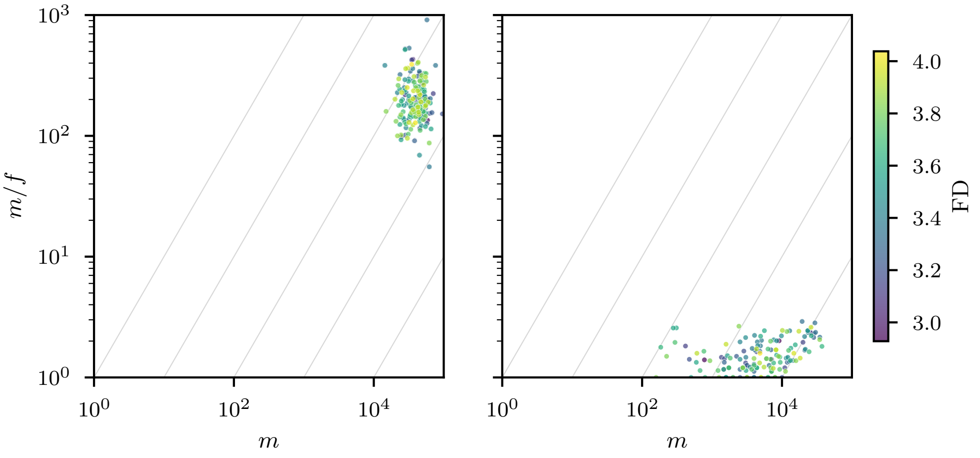

Figure II.4.11 Center of mass of the trait combination where the largest positive **(left)** or negative **(right)** relative species biomass changes $\widehat{\Delta B}_{\mathrm{S}}$ occurred – the same condition as used in Figure II.4.10. The marker colour denotes the corresponding initial FD value. The centre of mass was computed from log-transformed, equally weighted positions in trait space including only those states that showed $\geq 5\,\%$ positive or negative relative biomass change. Food webs that have no trait combination fulfilling such a condition do not appear in this plot.

To further corroborate that strong positive biomass change primarily occurs for invaders with high m and $^m/_f$ and strong negative change is typically in the low $^m/_f$ regions, Figure II.4.11 shows the centre-of-mass (COM) of the areas of strongest increase and decrease. The strongest changes are in fact found in the expected regions with $f \gtrapprox 100$. For the biomass losses, the COM is frequently localised at the boundary of the sampled trait space, indicating that only those 'cannibalistic' traits with $^m/_f = 1$ led to strong biomass losses. The right subplot also shows far fewer points because there are many food webs that do not exhibit a sufficiently strong biomass loss.

These COM plots can not only show the position of some effect in trait space, but also denote the FD of the corresponding initial state, hence allowing to detect whether effects shift to different parts in trait space depending on the properties of the invaded food web. However, there appears to be no obvious correlation between FD and the COM of strong biomass changes for these observables. In fact, the same is true for a wide range of other tested observables[52]: While the position and variance of the COM grows or shrinks, no systematic shift in trait space depending on FD was found. The strongest variances were observed for the position of large cascades, which is not particularly surprising given the previous observations.

---

[52] For a list of studied observables in this context, see Appendix B.

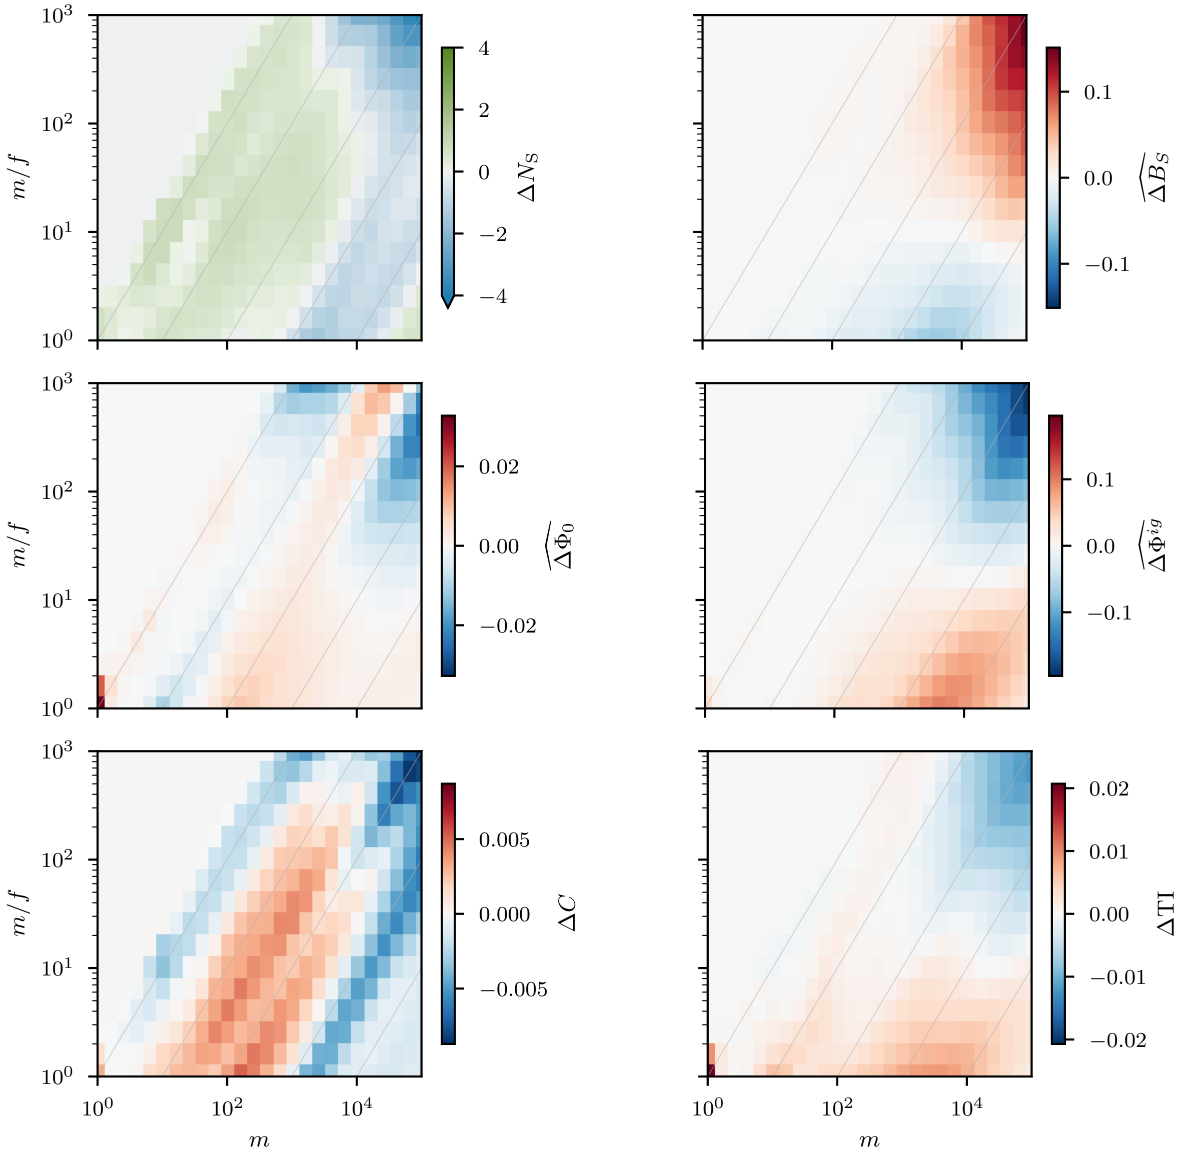

Figure II.4.12 Invasion experiment observables averaged over all 209 instantiations.

With none of the above-mentioned observables showing strong correlations to functional diversity or systematic shifts in trait space, all 209 food webs can be averaged to arrive at a visualisation of the more general effect of species invasion, regardless of functional diversity; see II.4.12. These averaged quantities show that many of the qualitatively studied responses to species invasion are in fact systematic responses that are independent of the exact initial conditions. They can be summarised as follows:

1. Invader survival and coexistence is only possible for low f, with a clearly showing dip for non-optimal feeding on the external resource.

2. Regions with extinction events align with invasion at the top-predator level, $f \in [10^3, 10^4]$, and in the high $(m, {}^{m}\!/\!{}_{f})$ regime. As demonstrated , food webs are more prone to extinctions for invaders of such traits because they are in direct competition with existing top-predators, the extinction of the latter can trigger a larger reorganisation of interaction strengths, hence often entailing extinctions.

3. Overall species biomass changes are minimal throughout most of trait space, unless for high m or high ${}^{m}\!/\!{}_{f}$, as discussed. The intra-guild biomass flow $\widehat{\Delta\Phi^{\mathrm{ig}}}$ correlates negatively to the change in biomass.

4. The presumed alternating pattern in $\Delta\Phi_0^{\mathrm{out}}$ is more clearly visible in the averaged heatmaps shown here. This further supports the notion that topologically motivated mechanisms like top-down control may be applicable if pertaining only to the external resource or other global observables like the total respiration rate, where exactly the same changes can be seen (not shown). In essence, the patterns are discernible here because all species-specific changes are averaged out and projected onto the resource input to the system.

5. The connectance of the discretised network structure increases for invaders on central levels and decreases in case of extinction cascades or invaders feeding directly on the external resource. As discussed, this is in large parts a result from the procedure by which the network is extracted.

6. Change in trophic incoherence correlates negatively with biomass change, caused in part by a change in dominance of ‘more incoherent’ connections. Trophic incoherence also slightly increases for invaders on low trophic levels, where the invader feeds on species in multiple clusters.

The previously observed responses to species invasion not only appearing in individual food webs but, in a broader sense, also in these averaged representations further supports the explanations given in the context of the qualitative analysis. Furthermore, it seems that the general properties of the invader species, like whether they act as a top-predator, are more relevant for the average food web response than the functional diversity. While functional diversity appears to play a role for coexistence, other observables hardly depend on it or show a large variance.

Dependencies on network structure

While the functional diversity of the evolved ECST model food webs seems not to correlate strongly to invasion response, the network measures of connectance C and trophic incoherence TI may show more pronounced effects. As seen before, regions in trait space that exhibited strong invasion effects on system biomass coincide with those of changes to trophic incoherence. If biomass *change* correlates to trophic incoherence *change* – mediated by interaction strengths shifting to or away from less coherent connections – the food web's trophic incoherence before the invasion likely plays an important role for biomass change as well.

In addition, studies like [*Johnson* et al. 2014] found trophic incoherence to be a better statistical predictor for linear food web stability of real food webs than other structural measures. The authors suggest that food webs may grow in size and complexity without becoming less stable as long as they are sufficiently coherent. Based on the linear stability analysis they conducted, they argue that food webs with increased incoherence require stronger regulation mechanisms (like competition, mutualism, cannibalism) to remain stable.

Regarding the connectance, as briefly motivated earlier[53], previous studies suggest that a food web's connectance may be important for determining its stability towards invasion: *Romanuk and Zhou* et al. [2017] use a similar setup of an invasion experiment as carried out here to investigate extinction cascades: They generate food webs[54] to attain networks of varying connectance $C \in [0.05, 0.30]$; subsequently, they introduce a new species with random properties and study species extinctions, defining *robustness* as the proportion of surviving resident species. They find that among the food webs they studied, those with high C are more robust but show larger extinction cascades in case there are extinctions; vice-versa, low C food webs were observed to be less robust but exhibit smaller extinction cascades.

However, due to the inherent difficulties in the procedure to attain a sensible connectance value from the ECST model food webs, especially when comparing these measures from times before and after the invader was added, I will focus on studying the trophic incoherence. Unlike the connectance, which relies on a threshold-based discretisation of interactions from the functional response matrix in this model, the trophic incoherence captures link strengths and thus does not

---

[53] See beginning of this section on species invasion, Section II.4.1.

[54] Unlike here, the process they use is based on randomly generated adjacency matrices with 30 species, followed by a filtering procedure that removes ecologically implausible networks. They then assign random biomass densities and use a niche model with a saturating functional response to arrive at equilibrium biomass densities; the invader is introduced into this resulting configuration.

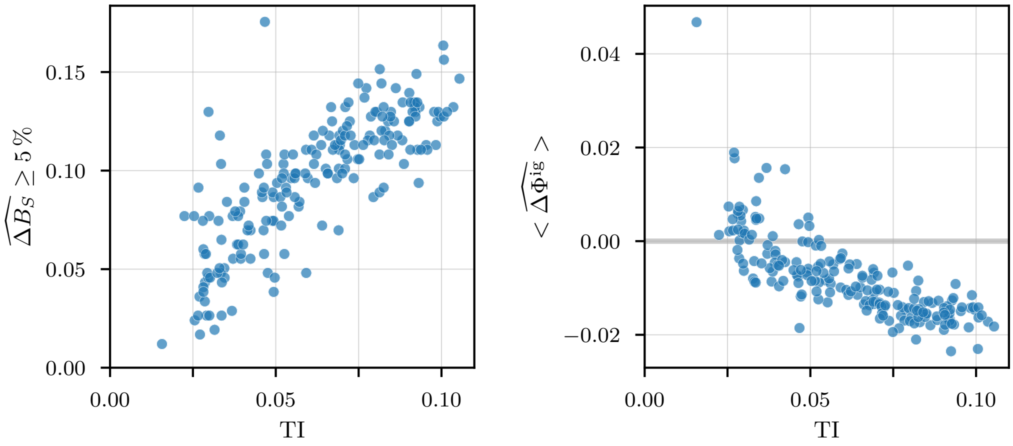

Figure II.4.13 The fraction of sampled trait combinations that led to a relative biomass change of more than 5 % **(left)**, and the mean relative intra-guild flow change **(right)**, plotted over the trophic incoherence of the initial state.

only alleviate any need for link thresholding but also takes into account more of the available information.

Another motivation behind this was that C and TI appear to be correlated. Thus, the set of food webs chosen to study the relationship towards trophic incoherence represents a similar distribution of connectance values (cf. Figures B.11 and B.14). The reason behind this correlation lies in the properties of links in certain parts of the food web both affecting the incoherence measure and the connectance. For instance, a new species with TL ≈ 2 will have a low link density and low incoherence contribution, thus reducing both values; a species on higher trophic levels will more likely have a high link density and thus also a trophic distance distribution with a high variance, increasing both values. The measures differ from each other when it comes to top predators, which can typically not form as many links as a species on TL ≈ 3, but the links may in some cases be very pronounced, thus contributing to a larger part into the trophic incoherence than in the connectance. Given the correlation between the two measures and the systematic problems with the connectance, focussing on the trophic incoherence seems like the more robust approach.

Same as with the functional diversity investigated above, initial states from the available range of trophic incoherence values were selected via binned sampling (bin width TI $= 0.01$, ≤ 24 samples per bin), leading to 196 evolved food webs as starting point for invasion experiments, totalling around 81 k individual simulations.

Figure II.4.13 shows two biomass-based measures of ecosystem functioning plotted against the trophic incoherence of the invaded food web. There appears to be a strong positive correlation between trophic incoherence and the biomass change of

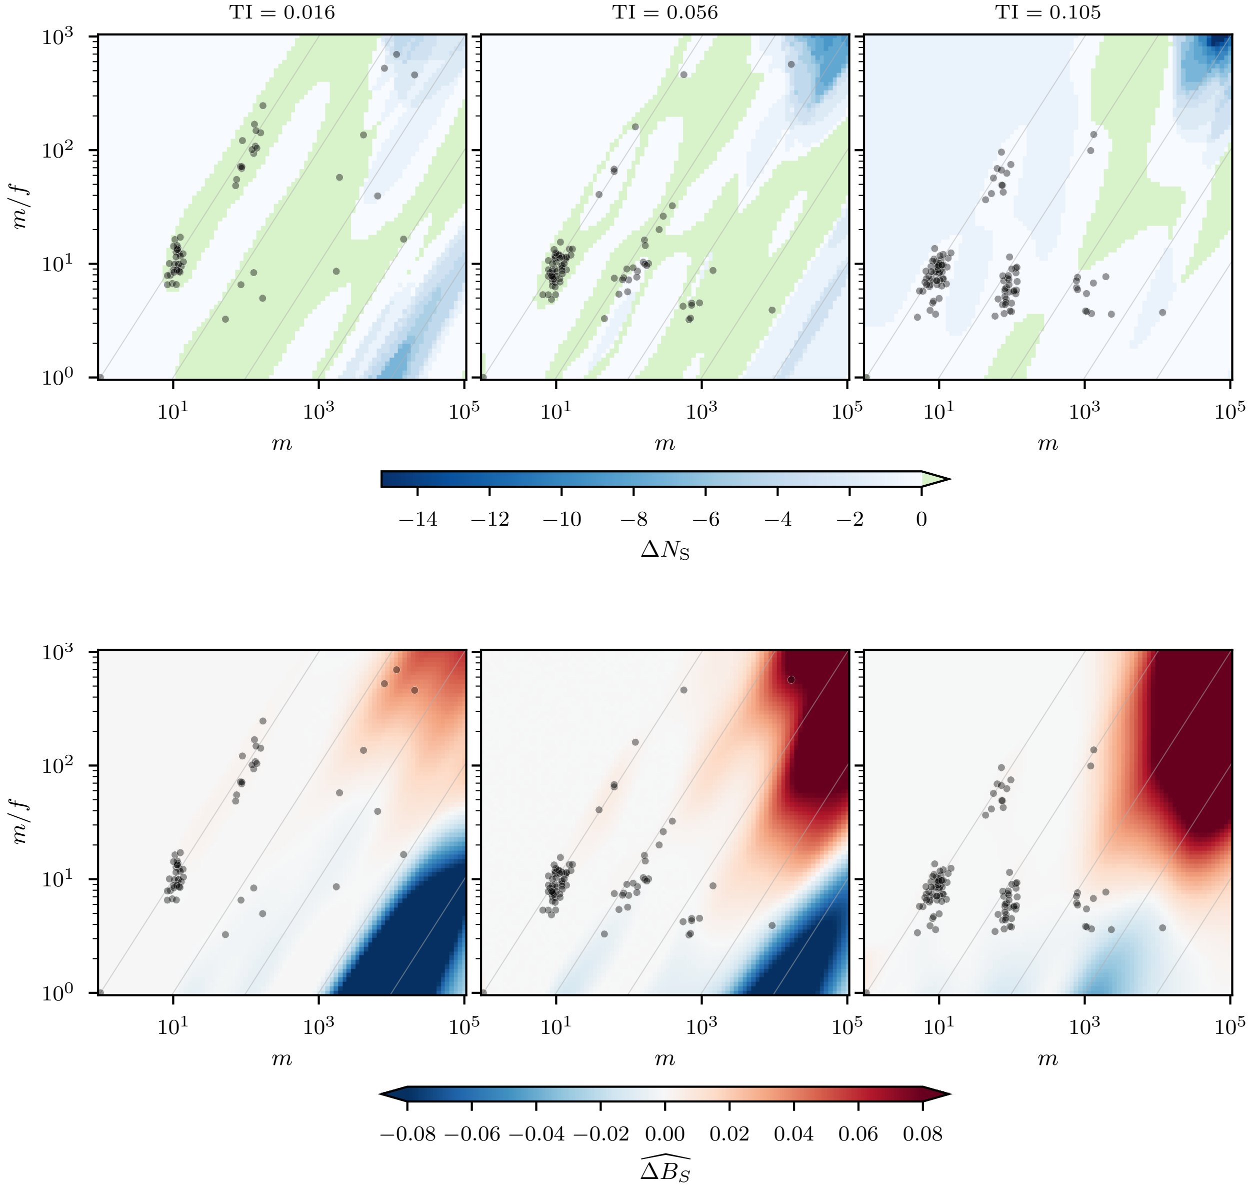

Figure II.4.14 Changes in food food web size and species biomass for three selected initial states. The respective number of species before invasion are $N_\mathrm{S} = 55$, 85, and 107.

the system and a negative correlation with the intra-guild biomass flow Φ^{ig}. At the same time, no correlation can be observed between the susceptibility to cascades depending on trophic incoherence (cf. Figure B.12).

This correlation is caused by two effects that appear to be related to the trophic incoherence of the food web: As Figure II.4.14 exemplifies, the area with negative biomass changes in the regime of low feeding distances is reduced while the area with positive biomass changes in the regime with high m and high feeding distance grows.

The mechanism behind the increase in species biomass is presumably some form of top-down control; yet, how exactly this may occur and how it relates to trophic incoherence is unclear at this point. As such, why invaders in the respective regions of trait space have these effects on the food web needs further investigation.

II.4.1.5 Summary

In this section, I studied the response of the ECST model to a new species being introduced into evolved food web states, an ‘invasion’. While some general patterns can be observed and allow making statements about the average response, it appears that an invaders’ trait values in combination with the traits of the resident species are the most relevant aspects determining the food web response.

II.4.2 Response to global perturbations

After having studied food web response to the microscopic perturbation of introducing a new species, the focus of this section is on a macroscopic perturbation, specifically: the reduction of the external resource's carrying capacity K. The carrying capacity is a quantity that substantially affects the extent of the structures that can emerge in the food webs as it controls the maximum abundance of the resource and thus the energy available to the whole food web. Subsequently, a sudden strong reduction will cause the extinction of many species. Due to trophic chaining and the corresponding conversion losses at each connection, species on higher trophic levels can be expected to be more strongly affected by the reduced energy flow.

Perturbations as studied here can be associated with the sudden loss of vital resources like light or nutrients in real-world ecosystems. While the model is formulated in terms of energy, the perturbation could also represent an effect of habitat loss that leads to a reduced availability of the external resource or other indirect effects.

In general, these perturbations aim to represent possible anthropogenic effects on ecosystems. However, in the strongly abstracted and simplified representation that is the ECST model, any such association can at best be of a qualitative nature and give hints at potentially important mechanisms in these kinds of systems. The aim for this section is hence primarily to illuminate to which extent specific food web properties may play a role in the response to a reduction in available energy.

II.4.2.1 Methods

The study of the response to a global perturbation is methodologically similar to that of species invasion: Sufficiently evolved food webs are loaded, the perturbation is introduced, and the system is iterated until reaching a stable state.

For the simulations studied in this section, the food webs were first loaded and iterated for 100 steps to ensure that they are in equilibrium; with the loaded food web themselves also being 100 steps away from their last mutation event, the initial states can be assumed to be equilibrated. At $t = 100$ steps, a perturbation was introduced, specifically: the external resource's carrying capacity $K = K_0$ was reduced by a factor $z \in [0, 1]$ to $K' = zK$. The system was then iterated until $t = 4\,\mathrm{k}$ steps.

Note that, as before, the mutation mechanism is completely disabled here. In this section, the focus is on investigating the direct response of the population dynamics on short time scales, depending on food web properties. Unlike in the invasion experiments, these simulations use the default value of the extinction threshold.

Generation of sufficiently evolved initial food web states

For the initial states, the same pool of evolved food webs is used as during the invasion experiment. Without the need to sample from the $(m, {}^{m}\!/\!{}_{f})$ trait space, the number of simulations is drastically smaller in this setting, making it computationally feasible to use the whole set of 2048 evolved food webs for studying perturbation response.

The general properties of these 2048 evolved food webs are shown in Figure B.14. They were generated in a multi-step process: First, an initial set of 32 instantiations of the ECST model was iterated over 20 M steps with different PRNG seeds. Then, resulting food webs were loaded and iterated again for 500 k steps and 16 different seeds, sufficiently diversifying them; this was repeated for 4 additional seeds, yielding a set of 2048 evolved food webs. The 500 k steps iteration time (or: 5000 mutation events) were more than enough for diversification given that the ECST model exhibits rapid changes in structural observables and other food web measures on much time scales of 50. . . 100 k steps. This was verified by inspecting the time series of key observables and the ensuing increase in variance between simulations starting from the same initial state but continuing with a different sequence of random numbers.

With the resulting final states of the procedure having nothing in common but the number of iteration steps, they can be considered to represent a random sample of food web states that may evolve throughout the iteration. Therefore, these kinds of food web sets are referred to as 'random' samples.

Generation of biased samples of initial states

For a second scenario, perturbation effects on initial states from the 'random' sample are compared to those on a biased set of initial states. To introduce a bias into the set, a conditional selection process is applied during the iteration, instead of randomly selecting states and thus approximating an unbiased sample of the available states. Depending on the selection condition, the bias may be more or less pronounced.

To generate the biased sample, the snapshotting feature of the model implementation is used. These snapshots can be conditioned on certain observables of the system and store snapshots at any point during the iteration, like the current value of trophic incoherence. The `SnapshotManager` then continuously tracks the relevant observable and will keep snapshots of those handful of states with the largest (or smallest) values of that observable.

One such biased sample that is studied below is based on the size of extinction cascades caused by the addition of a new species. In the case of extinction cascades, the effect of adding a new species can only be evaluated at a later time, because it takes time until potential secondary extinctions have occurred and the system has fully equilibrated. However, with the more interesting state of the system being *before* the cascade (representing a possibly volatile state of the system – more on this later), the snapshots are also made of the state before the species causing the cascade is added to the system. It is then kept as a candidate snapshot, which is only added to the `SnapshotManager` once the size of the cascade is clear.

This method of selecting states may introduce correlations to other samples created from the same simulations, specifically: correlations towards shared initial states (in simulations with different seeds) or to states in temporal proximity to each other. To avoid correlation to shared initial states, snapshots were collected no sooner than 200 k steps, allowing enough time for food webs to diversify before snapshots are created. Also, to not have snapshots from the same simulation in temporal proximity to each other, only a single snapshot is selected per simulation.

How about the case where the respective final states of the iterations are used for the 'random' sample and snapshots from the same simulations are used for a biased sample? This is indeed a case where correlations may be introduced if the snapshot is made close to the end of the simulation. However, for the conditioning on extinction cascades, this is hardly a problem, because snapshots from before a cascade and the respective state after the cascades would be different in at least one important aspect: the species extinctions, which often entail a restructuring of large parts of the food web.

II.4.2.2 Perturbation of a random food web states

Figure II.4.15 shows two observables from a perturbation experiment on the unbiased set of 2048 food webs (the 'random' sample). It includes simulations carried out with three different perturbation factors $z \in 0.1, 0.4, 0.9$, with a lower value denoting a stronger perturbation.

The following observations can be made: First, weak perturbations of $z = 0.9$ leave both the functional diversity and the (corrected) intra-guild flow mostly unchanged. While some extinctions seem to take place, necessarily leading to a lower FD, these seem to allow increases in Φ^{ig} for a few cases (which is not unexpected after the observations made during species invasion). A correction with z is important when comparing Φ^{ig} as it can be expected that all biomass flows in the system are reduced by that factor; with the correction applied, the observable quantifies changes that go beyond the mere reduction in biomass flow.

Second, the changes in FD show a clustered structure, which essentially reflects the trophic structure of the perturbed food web. Due to its definition, single species can contribute a great deal to the FD value if they feed on species that no other species feed on; this is primarily the case for apex predator species, of which there are typically only very few. With these species being most susceptible to becoming extinct as a result of the perturbation, FD changing in jumps is not surprising. The absolute values of FD are helpful in assessing the remaining structure of the perturbed food webs. For instance, the strongest perturbation often leads to FD $<$ 1.5, suggesting that only a cluster of species feeding directly on the external resource

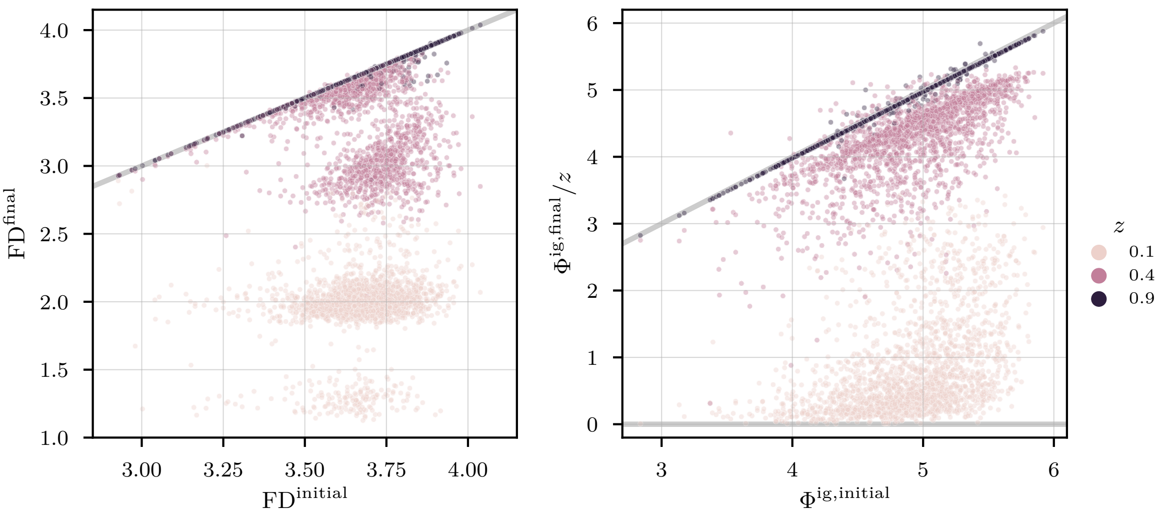

Figure II.4.15 The effect of a perturbation in carrying capacity K on the functional diversity **(left)** and the intra-guild flow **(right)** for 2048 evolved food webs. The grey line denotes equal values before and after the perturbation.
Here, perturbations with three different perturbation factors $z = {}^{K'}/_{K}$ are shown, where K' is the carrying capacity after perturbation. In the right panel, the intra-guild flow value after the perturbation is corrected by this factor to make the different outcomes more comparable. Note, that the distribution of initial properties is not uniform.

are able to survive. The perturbed webs with FD ≈ 2 suggest the existence of another trophic layer, such that these webs have their top predators at TL ≈ 3. Corresponding to the collapse of higher trophic structures, the intra-guild flow for stronger perturbations is very much reduced, meaning that even if trophic interactions are maintained between species, they are much weaker than before.

Third, stronger perturbations seem to eradicate most of the correlation to the initial state. While systems with higher initial Φ^{ig} seem to be able to retain higher flow values, this is not the case for FD which seems to be dominated by the trophic structures attainable for a certain energy availability. This again highlights a distinction between FD and measures which directly represent ecosystem functioning: FD may depend strongly on few species without quantifying how well-embedded they are in the food web while high Φ^{ig} values represent an established energy pathway through the system that is maintained not by a single species but by the whole food web.

Having studied the correlation between the final and the initial values of the same observable, the next step is to look for potential dependencies between the change in one observable compared to some property of the initial state. To that end, the relative change in species biomass $\widehat{\Delta B_{\mathrm{S}}}$ was plotted against the initial states' functional diversity, connectance, and trophic incoherence (not shown). However, no correlation was found between the response and the studied properties – or whatever effect there may be is hidden by the large variance.

This observation raises the question, whether there is *any* importance of the initial state regarding the perturbation response at all – for this unbiased sample and the studied observables, this does *not* appear to be the case. At the very least, this observation prompts for a different approach to this line of investigation, in the hope to gain another perspective on this question.

II.4.2.3 Perturbation of potentially volatile states

The above observations suggest that most of the structural information of the initial states seems to be lost upon perturbation. At this point, there are the following two possible explanations for this behaviour: Either the studied food webs' structure really is dominated by the available energy and not by the previous state; or, the studied observables do not accurately capture the food webs' ability to retain structure upon perturbation and other food web measures are needed. Testing the former explanation can inform whether it is worthwhile looking for measures that may describe this retaining capability more accurately.

To that end, the procedure described in Section II.4.2.1 is used to generate two sets of initial states, one that is unbiased (like the 'random' set used above) and one that is biased. The perturbation response of these two sets of initial states are then studied and compared; if they show different behaviour, it may indicate which properties of the initial state are relevant for eliciting a different perturbation response.

The initial states in the biased set will be states *before* a large extinction cascade occurred as the result of a new species being added. These states can be considered as potentially volatile, as they appear to be prone to a large number of extinctions – hence, the sample is denoted as 'volatile' sample.

In order for the volatile sample to contain enough values with large cascades, the procedure to generate initial states was repeated once more: Starting from the sample of 2048 evolved food webs, each one of them was again iterated for 800 k steps and with 4 PRNG seeds, thus yielding a set of 8192 evolved food webs. Throughout those simulations, the following conditions were used for snapshotting: Snapshots were only made for cascade sizes ≥ 3 and after the first 200 k steps had passed (to let initial states diversify). From the 8192 snapshots with that cascade size, only those with sizes $\geq$ 10 were included into the 'volatile', resulting in a sample size of 1249 volatile states. In the resulting set, the median cascade size is 12; there are about 25 % of states with cascade sizes ≥ 15, and about 10 % with sizes ≥ 20.

In how far these food webs really are in a more volatile state than the randomly selected ones is an open question at this point – after all, the extinctions that led to the selection of these states may have simply been caused by a new mutant species, where particular trait combinations can indeed trigger large extinction cascades (as

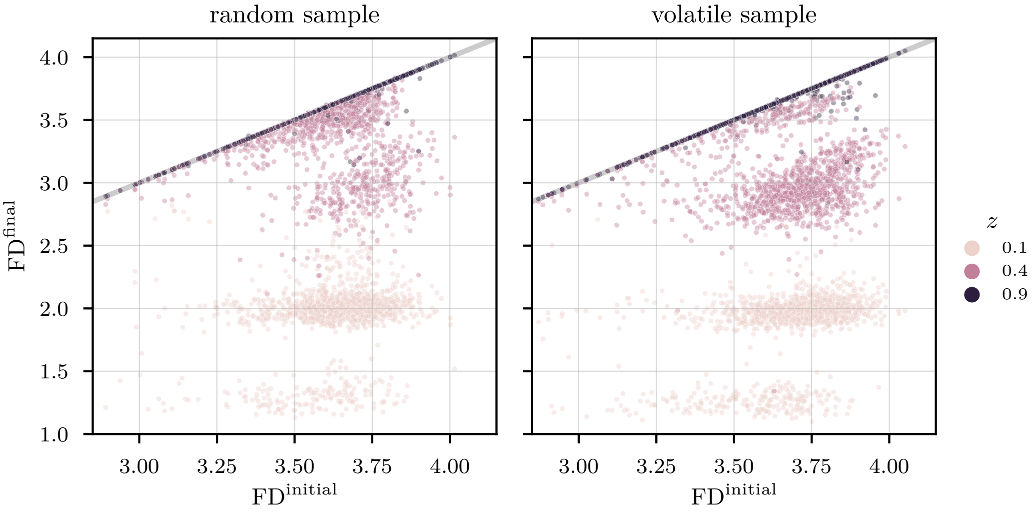

Figure II.4.16 Perturbation response in random and volatile sample.

seen earlier in the study of species invasion). However, by constructing a scenario that compares the essentially unbiased set of initial states to a set that is biased by this condition, such a hypothesis can be studied.

The 'random' sample for comparison consists of the respective final states of the 1249 simulations used for the volatile sample. Note that this is not the sample as used above, but a subset of the final states of the 8192 simulations used to create the initial states; this is mostly a technical limitation. See the methods section above for an argumentation why this is not a considerable issue.

A comparison of the distributions of basic observables of the two samples (cf. Figure B.17) shows some differences in the distributions: The food webs classified as volatile generally consist of more species and have a slightly larger amount of species biomass. Furthermore, the functional diversity distribution is shifted towards larger values. All these differences are comparably small, but appear to be systematic.

Simulations were performed for 19 different perturbation factors $z \in [0.1, 1.0]$ and both the random and the volatile sample. Figure II.4.16 shows the effect of the perturbation on the functional diversity of the random and the volatile sample for three selected perturbation strengths. Most differences can be seen for $z = 0.4$, where the volatile sample shows a larger loss in functional diversity. In both samples, initial states with high FD do not appear to be able to retain high values. Again, the changes appear to be dominated by the underlying possibilities for the trophic structure. For strong and weak perturbations, differences between the samples are comparably small.

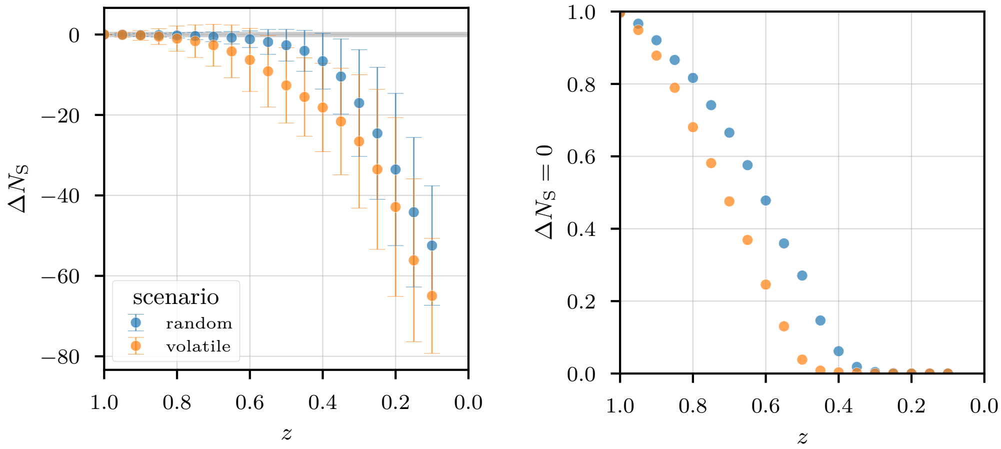

Figure II.4.17 Mean and standard deviation of change in N_S **(left)** and fraction of states where perturbation did not cause extinctions **(right)**, depending on the perturbation factor z. Note that low values of z indicate a strong perturbation and high values indicate a weak perturbation, which is why the x-axis is reversed in these plots.

While this approach was carried out for a wide range of other observables (e.g. Figure B.18) and is helpful to illustrate the induced changes in property distributions, studying the response solely via the distribution plots is difficult. To assess the response more systematically, scalar observables like the average change in N_S or biomass measures were evaluated over the whole ensemble of 1249 samples for each of the parameter combinations.

Figure II.4.17 displays results for two such observables and indeed shows differences between the scenarios: The volatile set not only had higher average species loss but also consistently fewer instances of food webs without any extinctions at all. As expected from the high variance in the distribution plots studied earlier, the standard deviation is large compared to the differences.

Turning to other observables, Figure II.4.18 shows some interesting effects in the biomass response: In general, the species biomass drops (on average) beyond what would be expected from the reduction in available energy alone. This observation indicates that the species loss induced by the perturbation leads to a reduced ability of the food web to retain biomass. While the difference between the two samples is generally small, it is larger for intermediate perturbation strengths; for very strong perturbations, $z \leq 0.2$, the differences begin to disappear, indicating that whatever distinctions the two samples might have no longer plays a role for the biomass response. Notably, the volatile sample shows a larger standard deviation, indicating a more variable response to perturbation, presumably caused by the larger changes in species number in that sample.

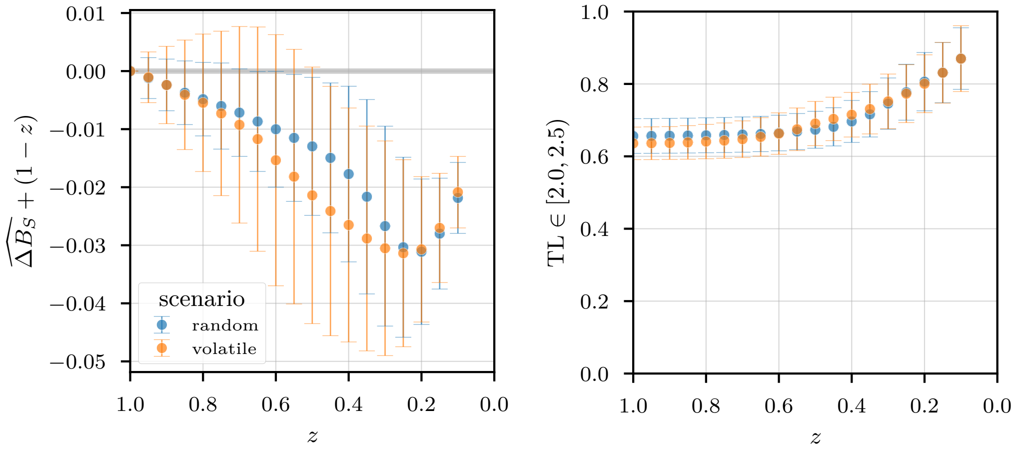

Figure II.4.18 Corrected mean and standard deviation of relative change in B_S **(left)** and fraction of states with TL ∈ [2.0, 2.5) **(right)**, depending on the perturbation factor z. Note that $z = 1$ corresponds to no perturbation having been applied.
The correction term $(1 - z)$ quantifies the expected relative biomass loss; by including it, this representation allows assessing how much the biomass changed *beyond* the expected change. For instance, for $z = 0.8$ a biomass loss of 20 % is caused directly by the perturbation, $\widehat{\Delta B_\mathrm{S}}$ being negative thus denotes that more biomass was lost than expected.

The right subplot in Figure II.4.18 shows the change in the number of species that feed primarily on the external resource (i.e., have TL just above 2). While the differences are not large between the two samples, the volatile sample appears to have slightly fewer low-level species before the perturbation (cf. point at $z = 1$). For increasing perturbation strengths, this relation is inverted, indicating that food webs in the volatile sample lose proportionally more high-level species than the random sample. For strong perturbations, these differences disappear; after such a strong perturbation, almost all species that were able to persist feed primarily on the external resource.

The results for the structural measures (FD, C, TI), are overall very similar: There is a systematic effect of all measures dropping off more quickly for the volatile sample than the random sample (cf. Figure B.19) for increase in perturbation strength. Overall, the variance is very large, and it is typically larger in the volatile sample than in the random sample.

To test the dependency on the threshold cascade size on these observations, further simulations were performed (not shown) where the cascade size of the volatile sample was smaller (3) or larger (15) than the one chosen above. In both cases, the qualitative observations are the same as above; the main difference is in the distribution of the initial state, which becomes more similar to the unbiased sample for smaller cascade sizes.

What makes the volatile sample react stronger to the perturbation and does this justify the classification as ‘volatile’?

Compared to the random sample, food webs in the volatile sample have higher N_S, B_S, and FD (see Figure B.17). The higher number of species is the dominant reason for the stronger reaction to a reduction in available energy: With more species sharing the same energy input, the average number density necessarily has to be lower, putting more species close to the extinction threshold. In effect, a reduction of available energy will more likely lead to species becoming extinct – which explains the early onset of extinctions in Figure II.4.17. With increasing number of primary extinctions as a direct result of the reduced carrying capacity, more secondary extinctions may be triggered during the equilibration phase of the food web.

While some differences between the initial states of the samples and their reaction can be observed, care should be taken to not overestimate the effect. Specifically, as visible from a distribution-based analysis (see Figure B.20) of the same quantities, it becomes clear that the response distributions have a large overlap, meaning that this event-based categorisation was in most cases not a good predictor for categorising food web states by ‘volatility’. Phrased differently: the food web states that were presumed to be prone to a strong reaction to a perturbation turned out to have in many cases very similar responses as randomly selected states. While this does not preclude a categorisation, the categories have to be defined somewhat softer, perhaps as ‘more volatile’ and ‘less volatile’.

One possible reason for the sampling condition not being particularly successful for this form of categorisation was already mentioned above: The extinction cascade may have simply been the response to a randomly selected mutant that happened to have a strong effect on equilibrium biomass densities and thus caused many extinctions. As seen in the study of species invasion, mutant traits appear to be significantly more relevant than macroscopic food web properties for the response of the food web, including the number of species. With this mechanism playing an important role, the sampling condition would assess the potential response to a perturbation only indirectly: via the distribution of species traits, which determines the probability of offspring species appearing in a certain region of trait space where a strong response may be elicited. However, this is a different type of perturbation than the reduction in carrying capacity: The categorisation happened following a mutation event, but what was studied in this section was the response to a reduction in primary biomass flow. Subsequently, the coupling between the specific response to a mutant and a general food web response has to be regarded as rather weak compared to the effect that mutants may have on the food web.

The motivation for comparing the two samples was to find out whether there are food web states that react systematically differently to perturbations than others. While this can in principle be confirmed, the effect appears to be weak for the chosen sampling condition.

There is another insight from this study which comes from the generation method for the volatile sample: Indirectly, this approach studied whether different kinds of perturbations are generalisable in their effect on the food web. The above results show that the response to new mutant species may *partly* inform the response to a change in carrying capacity, presumably mediated by the number of species in the food web and their number density distribution. However, it appears that the similarity is much weaker than the randomness introduced via the choice of parent species and the offspring traits.

II.4.2.4 Summary

The initial goal for the study of perturbation response was to assess whether food web properties can be found that make an evolved food web more or less robust against reductions in available energy. Overall, the above investigation found that strong perturbations lead to most food web measures becoming unrelated to the state before the perturbation and being dominated by the direct effect of reduced energy availability. For intermediate perturbation strengths, the biomass-based observables seem to be able to retain more information about the state before the perturbation than the structure-related measures. As shown above, the investigated structural observables of functional diversity, connectance, and trophic incoherence do not appear to be correlated to the perturbation response.

The attempt to construct a set of initial food web states that would exhibit a stronger response was only partially successful. While these ‘volatile’ states typically did elicit a stronger response than randomly selected initial conditions, the differences were not particularly strong, presumably due to the method chosen for acquiring the sample. However, the general approach of comparing a randomly selected set of food webs to a set that was created under a biased condition may be useful if a stronger or more focussed bias can be chosen, for instance by including additional qualifiers or using other events (e.g., large drops in FD or TI) as the trigger for snapshotting.

II.4.3 Recovery after perturbations

Ecosystem *resilience* denotes the capacity of a system to recover from some adverse influence, for instance by reorganising into a new structure that fulfils an equivalent function (refer to Section II.1.2.3 for an introduction).

Unlike the study of the direct effect of perturbations done in the previous sections, which focussed on the population dynamics and disregarded the evolution mechanism, investigating resilience requires that the modelled system actually has the ability to reorganise. For the ECST model, this requires the evolution mechanism to be activated which is the only means by which the model system can gain new species with different properties; without mutations, the system would lack the ability to reorganise itself and hence not be capable of recovery (and neither of any interesting dynamics, for that matter).

However, in the light of the findings regarding the perturbation response of the ECST model, it is unclear in how far the study of resilience would differ from directly studying the evolution of randomly evolved food web states. In other words: Given the lack of influence that food web properties seem to have on the perturbation response, is the ECST model at all suitable to study questions of resilience?

To approach this question, I will first apply the theoretical framework introduced by *Tamberg* et al. [2020] and summarised in Section II.1.2.3 to formulate what studying resilience using the ECST model would entail. I will then present some simulation results and reflect on the feasibility of using the ECST model for studying resilience.

II.4.3.1 Embedding into resilience framework

As described in Section II.1.2.3, the resilience framework proposed in [*Tamberg* et al. 2020] introduces a number of concepts to formalise the study of a system's resilience. Namely, the authors distinguish the following components: (i) the *system*, (ii) the *sustainant*, (iii) the *adverse influence*, and (iv) the *response options*. To apply this framework to the ECST model, these concepts need to be associated with mechanisms or structures of that model.

Some of these associations are easily made: Given that the ECST model already is a fairly compact model, the whole model can be regarded as the *system* of interest, namely: an evolutionary food web model of species and their trophic interactions, fed by an external resource.

Examples of *adverse influences* to such a system were already demonstrated in the previous sections: addition of a new ('invasive') species or a reduction of the energy available to the system. These previous findings can inform which kind of adverse influences would be interesting to study; specifically, the adverse influence would have to elicit a sufficient response, yet not bring the system into a state that is completely decoupled from the initial state (as could occur with an overly strong perturbation). As an example for a potentially unsuitable adverse influence, consider an invading top-predator, which typically led to an increase in species biomass but often did not elicit an extinction cascade – thus, if the system is able to fully absorb the perturbation in the first place, its recovery cannot be investigated.

There are many other options for adverse influences, which may also be in the form of direct manipulations of the interaction network or changes in global or local interaction parameters. The only requirement would be that the influence is temporary, allowing a comparison between a state before it was introduced and a well-defined time from which on the recovery can be assessed. In case of species-specific perturbations (invasion or dedicated removal), these can be considered as instantaneous. Changes to energy input or other food web parameters would only be applied over a certain time period and then be reset to their initial value.

The ECST model's *response options* to the adverse influence are given by its dynamic equation (the population dynamics) and the evolution mechanism (randomly selecting parent species and adding offspring species with mutated traits to the food web). For the system's resilience, the evolution mechanism is of larger importance than the population dynamics, because it models the ecosystem's ability to re-gain previously lost trophic interactions. The details of the evolution mechanism play an important role here: For example, parent species being selected randomly regardless of size or other factors make it more likely for low-level species to create offspring; depending on the adverse influence, this may strongly influence the recovery. A similar argumentation can be applied to the mutation rules, which transport only a small amount of information to the offspring species (because f and s traits are drawn from uniform distributions).

A more difficult association is that of the *sustainant*, i.e. the property of the system that is to be sustained or a state that the system is meant to recover to. As touched upon in Section II.1.2.3, this may be a normative choice, specifically because it requires to choose one functionality of the system as representing its to-be-sustained purpose. In the context of the ECST model, however, the choice is primarily of a practical nature: What is the function of the ECST model and which system observable best describes it? And depending on the choice, which would be the acceptable ranges or recovery times for those observables?

So far, when considering the functionality of the ECST model food web, it was associated either to the amount of biomass it retained, its biomass turnover, or some structural measure like its functional diversity. Specifically the functional diversity was formulated to capture the diversity of interactions in the food web, as discussed in Section II.3.5.2.

All of these measures would in principle be candidates for sustainants, but which of them are useful to capture functionality also depends on the adverse influence: For instance, if the adverse influence is a temporary reduction in available energy, the biomass-related measures would probably not act as a proxy for food web structure, mainly because high biomass flow is also possible with few low-level species (cf. Figure II.3.16). In contrast, the functional diversity would allow formulating recovery, because the reduction in available resources would lead to species extinctions and thus lead to a reduction in functional diversity; as soon as the adverse influence subsides, the mutation mechanism will lead to a recovery in this observable and the recovery time could be defined as the time it took to consistently cross a certain threshold value. In turn, for scenarios where there is little change in species composition, functional diversity will not be a useful choice of sustainant.

The above associations formally match the requirements of the resilience framework proposed by *Tamberg* et al. [2020]. However, as hinted to in the outset of this section, the observations made with the ECST model require further thoughts on whether this model is capable of a reasonable representation of resilience and what such an outcome would look like in such a case.

To illustrate why this is of importance, consider a temporary reduction in the carrying capacity as the adverse influence and the functional diversity as the sustainant. Now assume an extreme case where the adverse influence leads to all but one species becoming extinct. Subsequently – as all information about interactions are encoded in the species' traits alone –, the resulting state is equivalent to the initial state; there is no mechanism by which the system could develop in a qualitatively different fashion than said initial state[55].

While the adverse influence would surely not be chosen to be that strong, the point regarding the information content of any particular food web state is crucial: With all information about interactions encoded solely in the traits of existing species, species extinctions necessarily lead to information loss. This information loss occurs in trait space, where certain volumes in m space become no longer directly reachable by mutations[56]. Although still accessible via multiple mutations, the probability of those parts of trait space becoming consistently re-populated depends primarily on the trait distribution of the remaining species, as they determine the offspring body mass. Thus, any information that can be retained by the remaining species would be in the form of their trait distribution favouring certain kinds of feeding interactions (predators) over others. How much of the initial properties are relevant is, in addition, dependent on the coupling strength between species and their offspring; with the comparably large mutations of the ECST evolution mechanism, this coupling is rather weak.

For the situation at hand, the following research questions can be formulated: Does the remaining information suffice to noticeably affect food web recovery after the adverse influence subsided? If so, can macroscopic properties be extracted that quantify this capacity – or is the food web's whole micro-state relevant? Which

---

[55] Of course, the random number sequence will be different, thus leading to a different manifestation of the system. However, the observed dynamics will be qualitatively the same.

[56] This trait space perspective also puts the functional diversity measure into context: It is a measure that relates to the information content of the food web. However, it lacks a full representation, because it includes the body mass m trait only indirectly: via the relative feeding distance, $m/f \in [3, 1000]$.

parts of the model mechanism may be modified to allow retaining more information in such a scenario?

In the following, I will briefly study the first question, making specific choices for the sustainant and the adverse influence. However, it will soon become clear that very little information is retained, which is why further investigating these questions is not reasonable in the scope of the ECST model.

II.4.3.2 Recovery after a period of reduced energy availability

With the above considerations regarding the sustainant, using the functional diversity as a sustainant is a pragmatic choice, as it captures species richness; while it does not take into account their abundances, their existence alone is a representation of the depth of the food web. The adverse influence will be in form of a temporarily reduced carrying capacity. To quantify its resilience, the time until functional diversity recovers to a certain value is measured. The question to study thus becomes: Are there macroscopic properties of the ECST model that affect its resilience towards reductions in carrying capacity of the external resource?

To study this, I performed simulations starting from 161 evolved food webs, selected by binned sampling of their functional diversity value (bin width FD $= 0.1$, ≤ 18 samples per bin, using the pool of 2048 evolved food webs). By sampling along FD, a wide range of values can be covered.

These evolved food webs were loaded and at first propagated without the mutation mechanism enabled. The adverse influence set in at $t = 100$, at which time the carrying capacity was reduced to $K' = zK$ with $z \in \{0.2, 0.4, 0.6, 0.8\}$; it was reset to $K = 100$ at $t = 5000$, which gave the food webs sufficient time to equilibrate after the perturbation. Simultaneous to the reset of the carrying capacity at $t = 5000$, the mutations were enabled, such that the food webs would evolve as they normally do. The simulations continued until $t = 300\,\mathrm{k}$. Because the recovery depends strongly on the sequence of random numbers, the procedure was performed for 12 different PRNG seeds, leading to a total of 7728 simulations.

To quantify the recovery time, certain threshold values for FD were set and the time it took for the system to cross them was defined as the recovery time. If a food web's FD already was above the threshold, the recovery time was zero; if the food web did not recover beyond the threshold value during the simulation time, it was marked as 'not recovering'. To have a more robust measure, the recovery times were evaluated using the rolling mean of FD, averaging over 15 mutation events (1.5 k steps); this way, simulations where a short-lived species caused a temporary increase in FD did not directly cause a recovery time to be recorded but only if the increase in FD was sustained over a longer time.

Figure II.4.19 show the recovery time distributions for a threshold value of FD $=$ 3.4, plotted over the FD prior to the perturbation. Most food webs recover within 100 k steps and food webs with weaker perturbation recover more quickly. With weaker perturbations (increasing z), the number of food webs that did not exhibit

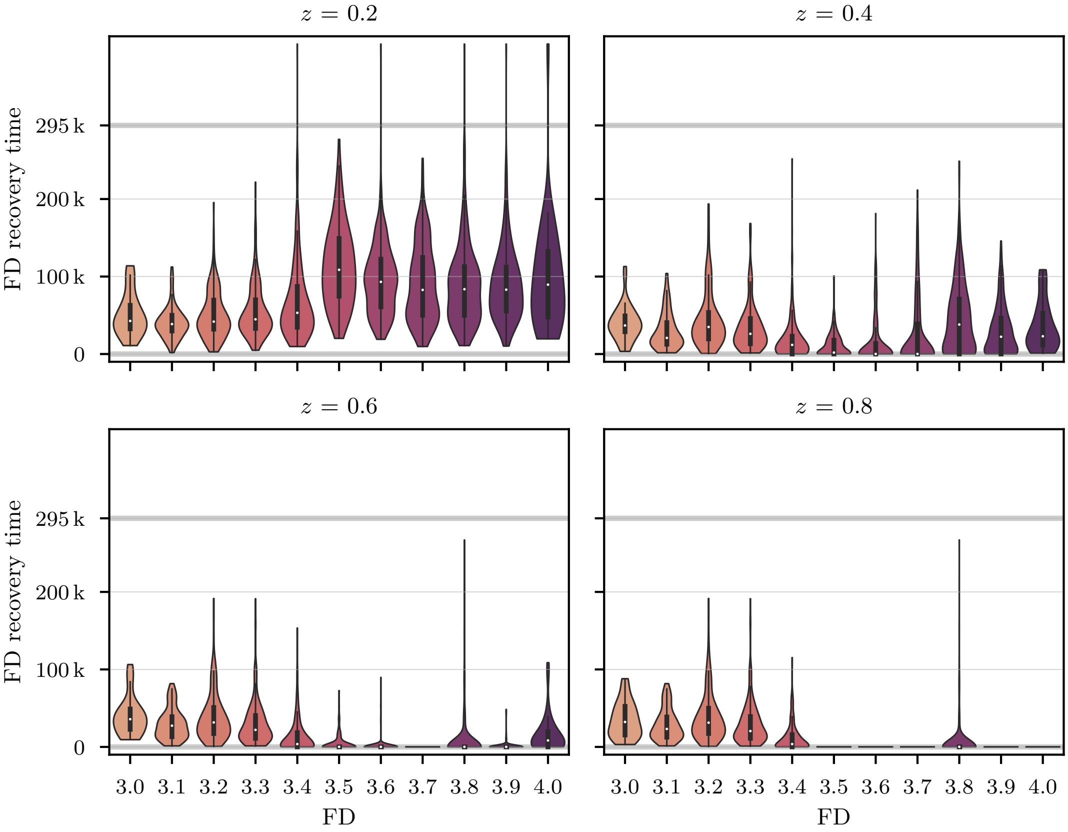

Figure II.4.19 Recovery time of food webs after perturbations: simulation times until the rolling-time-averaged functional diversity reached FD ≥ 3.4. The grouping on the x-axis is by the FD of the food web prior to perturbation, with values denoting the bin centre.

Note that food webs that did not recover during the time of the simulation were (for the purpose of this plot) assigned a value of 400, thus appearing in the unmarked parts of the y-axis. The violins are not extended beyond the data range and their areas are not normalised by counts. Each subplot represents 1932 simulations, roughly equally divided among the bins on the x-axis (due to the binned sampling of the initial states), but with slightly fewer counts in the first and last bin (due to fewer initial states with values in those data ranges).

a reduction in FD to values below the threshold increases, thus having a recovery time of zero.

For the food webs starting with FD < 3.4 the recovery times are always non-zero and are generally longer than for those with FD ≥ 3.4. This is not at all unexpected, given that the threshold was set to FD $= 3.4$. Even with weak perturbations, which often have little effect on the food webs (as seen earlier), the initial states of these food webs are below the threshold, thus necessarily leading to longer recovery times than for those that are above the threshold.

There is one unexpected observation in Figure II.4.19: For strong perturbations of $z = 0.2$, the food webs with FD > 3.4 show a systematically larger recovery time and wider distributions. However, this is a result of the model mechanisms and a form of selection bias in the pool of food webs: It is more likely that food webs with high FD also have a large size, making it more likely that their size after the perturbation is also large. A strong reduction in carrying capacity will primarily affect high TL species, and equally so in food webs with low and high FD values. As a result, the high FD food webs will be more likely to contain more species than food webs that had a low FD before perturbation. Now, the random selection of parent species comes in, which makes it more likely for low TL species to generate offspring – and thus more unlikely that FD increases. In effect, the food webs with high FD have a lower mutation speed, because low TL species will be selected more frequently than in food webs with fewer species.

The above observations are generally similar for lower or higher threshold values (not shown). For lower values, the recovery times are shorter and are frequently zero for the weaker perturbations; for higher values, distributions become much wider and food webs more often do not recover within the time of the simulation. The threshold value of FD $= 3.4$ was a compromise between these two extremes and a reasonable computational effort: Of course, longer simulations could be performed, but given the above results, it is unlikely that something qualitatively different would be observable. Also, effects become harder to discern for longer simulations, because there is a large amount of randomness being introduced by the mutations, which will cloud the effect of the initial state or the state after the perturbation.

These analyses were also repeated with the FD of the *perturbed* food web on the x-axis (see Figure B.21), again not exhibiting a strong dependency between recovery time and the respective FD value: While recovery time was slightly longer for food webs that had a lower FD after the perturbation, distributions were still very wide and this effect does not appear to be particularly strong. Again, the case of $z = 0.2$ behaves somewhat differently than those with weaker perturbations: Even for strong drops in FD caused by the perturbation, the recovery times are very similar to those cases with a smaller drop; which can again be explained by the above interplay of model mechanisms.

In addition, recovery times were correlated with other measures like TI and C, but showed no unexpected correlations apart from those that can be traced back to the effects described above. However, to rule out that there are any correlations, more extensive simulations and analyses would be needed.

The observations made during this analysis (of which only some are briefly presented here) match the presumptions made above that there is very little information retained across a perturbation, such that the state before and after a perturbation are largely decoupled. In addition, the mutation mechanism introduces further randomness such that longer simulation times further reduce correlations to the initial state – even without perturbations, food webs differ strongly after 300 k steps with different random number sequences.

Overall, studying resilience with the ECST model appears to be difficult, specifically because of the highly stochastic nature of the model and the correspondingly small amount of information that is retained by the food web structure. In hindsight, given the properties of the model, this is not particularly surprising: If all the information of the food web is stored in the species interactions and a perturbation removes many of these interactions, that information will be lost irrecoverably. Consequently, any following dynamic will show only very weak correlations across the perturbation.

The above analysis thus merely shows that studying food web resilience without a sufficiently complex model is not particularly insightful. In Section II.6.2, I will further discuss how this model could be adapted to retain more information and thus perhaps make such an investigation more reasonable.

In the framework proposed in [*Tamberg* et al. 2020] these aspects could be referred to as the recovery options of a system; for the ECST model, they are largely equivalent to the normal operation of the system. This illustrates one aspect of any food web model that *does* show a correlation between recovery time and initial state, namely that it retains sufficient information across the effect of the adverse influence. In other words, such a model would require that the adverse influence does only affect selective parts of the system and the recovery mechanisms can then use the remaining information to more quickly recover the previous state.

II.4.4 Summary

In this chapter, I investigated the response of the ECST model to species invasion, perturbations in carrying capacity, and its recovery times after such a perturbation.

The perturbation experiments where the carrying capacity was reduced studied the capacity of the food webs to retain structural information of its previous state across such a perturbation. It showed that even intermediate perturbations will remove almost any correlation between diversity measures. Furthermore, the comparison of volatile and random samples of initial states allowed to see that the species number can be a (weak) predictor for the size of extinction cascades. For these reasons, the investigation of food web resilience – one of the main goals for this thesis – did not yield any interesting insights, mostly due to the limited capacity of the ECST model for studying questions of this kind.

The main observations from the invasion experiments were that the invader species' traits determine strongly what the response of the food web will be. Interestingly, there are certain trait combinations which will, on average, elicit a similar response in the food web; these were discernible as repeated patterns in the invasion experiments. Attempts to explain this using the concept of generalised trophic cascades showed, that the interaction structure of the ECST model is too complicated to allow a reliable mapping to certain food web topologies, mainly due to the smooth feeding kernels and spread-out species clusters. This structure may be seen as the result of a self-organising process, which I will further study in the next chapter, Chapter II.5.

II.5

Food Webs as Self-Organised Systems

Statements referring to *self-organisation* and *self-organised criticality* (SOC) of food webs are not uncommon in the literature. For the ECST model, *Allhoff and Ritterskamp* et al. [2015] also briefly touch on this topic, mainly by arguing that the cascade size distribution resembles a (very steep) power law.

In the following, I will put the above statements into context and reflect on the extent to which the ECST model shows self-organisation and self-organised criticality. Informed by the observations of the previous chapters, I will try to isolate the relevant concepts and try to translate them to this particular model.

II.5.1 Does the ECST model show self-organisation?

For a system to be regarded as *self-organising*, it would need to exhibit structure formation as a result of the interactions between the system's entities.

In the case of the ECST model, the arising structure is that of a multi-layered food web: species of different sizes and feeding preferences dynamically arranging in a way that resembles multiple layers. Notably, these layers should not be understood as the nested layers emerging from Evolution Mechanics, but simply as interconnected species clusters arranging along a trophic structure.

The layered structure of the food web was already observed in many different forms, be it directly in the network representation (e.g. Figure II.3.12b), in the averaged species invasion patterns (Figure II.4.12), or in clustered response to perturbations (Figure II.4.16). How persistent this structure is can best be studied by performing species-level statistics on a large number of long simulations, thus averaging not only over multiple instantiations but also over life cycles within individual simulations. As a result of such an approach, Figure II.5.1 shows the life time distribution depending on species' mean trophic level. It is in the distribution of species with the longest life times that the trophic structure is best discernible: For trophic levels that are close to an integer value, more species are able to survive for more than 1 M steps; in between integer TL, there are much fewer long-living species, especially between the first two levels.

These kind of structures arising is – of course – ultimately a consequence of the definition of trophic interactions and the computation of trophic levels. The argumentation is as follows: (i) Feeding interactions in the ECST model are body-mass

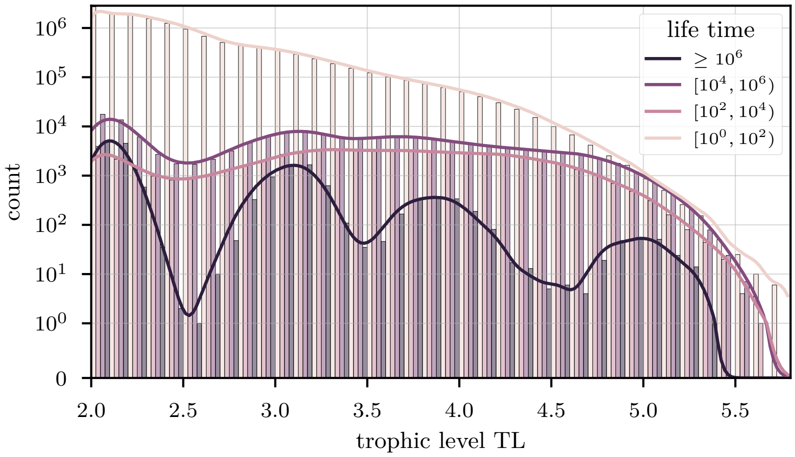

Figure II.5.1 Species life time distributions over their mean trophic level using data from 96 simulations over 15 M steps and a bin width of TL = 0.1. Kernel density estimates are shown to more easily distinguish the respective distributions.
This is a more detailed version of Figure II.3.18. The lightest shade corresponds to non-viable species and the remaining counts are from viable species.

structured due to the imposed relative feeding distance $^m/_f \in [3, 1000]$; (ii) successful prey species are more likely to generate offspring with similar body mass, which will only survive if they also have a similar feeding target, leading to clustering of prey species; (iii) in addition, predators feeding on clusters of prey can be assumed to be more successful than those feeding on isolated prey species, thus seeding another species cluster. However, while the ordering by body mass is imposed by the model mechanisms, the clustering of species can be seen as an emergent property: a result of the interplay of simple mechanisms acting between model entities and giving rise to macroscopically observable structures.

The centring around integer values is a simple consequence of the definition of the trophic level as *one plus* the flow-weighted average trophic level of all prey. The weighting by the respective predator's diet also explains why the third peak in Figure II.5.1 is slightly below TL = 4: Those species feed primarily on TL = 3, but can have a benefit if they additionally feed on species on TL = 2, thus reducing their distance to the external resource (and the corresponding conversion losses). Interestingly, this does not propagate to the fourth peak, which would then also be expected to be slightly shifted to lower values.

Studying structure using the trophic level distribution is a perspective focussed on the feeding interactions, which is a projection of the structures manifest in trait space. The trait space hold additional information, as can be seen in Figure II.5.2: For instance, it shows locations of trait clusters and how their feeding centre overlaps with other clusters, causing the shift away from integer trophic levels observed

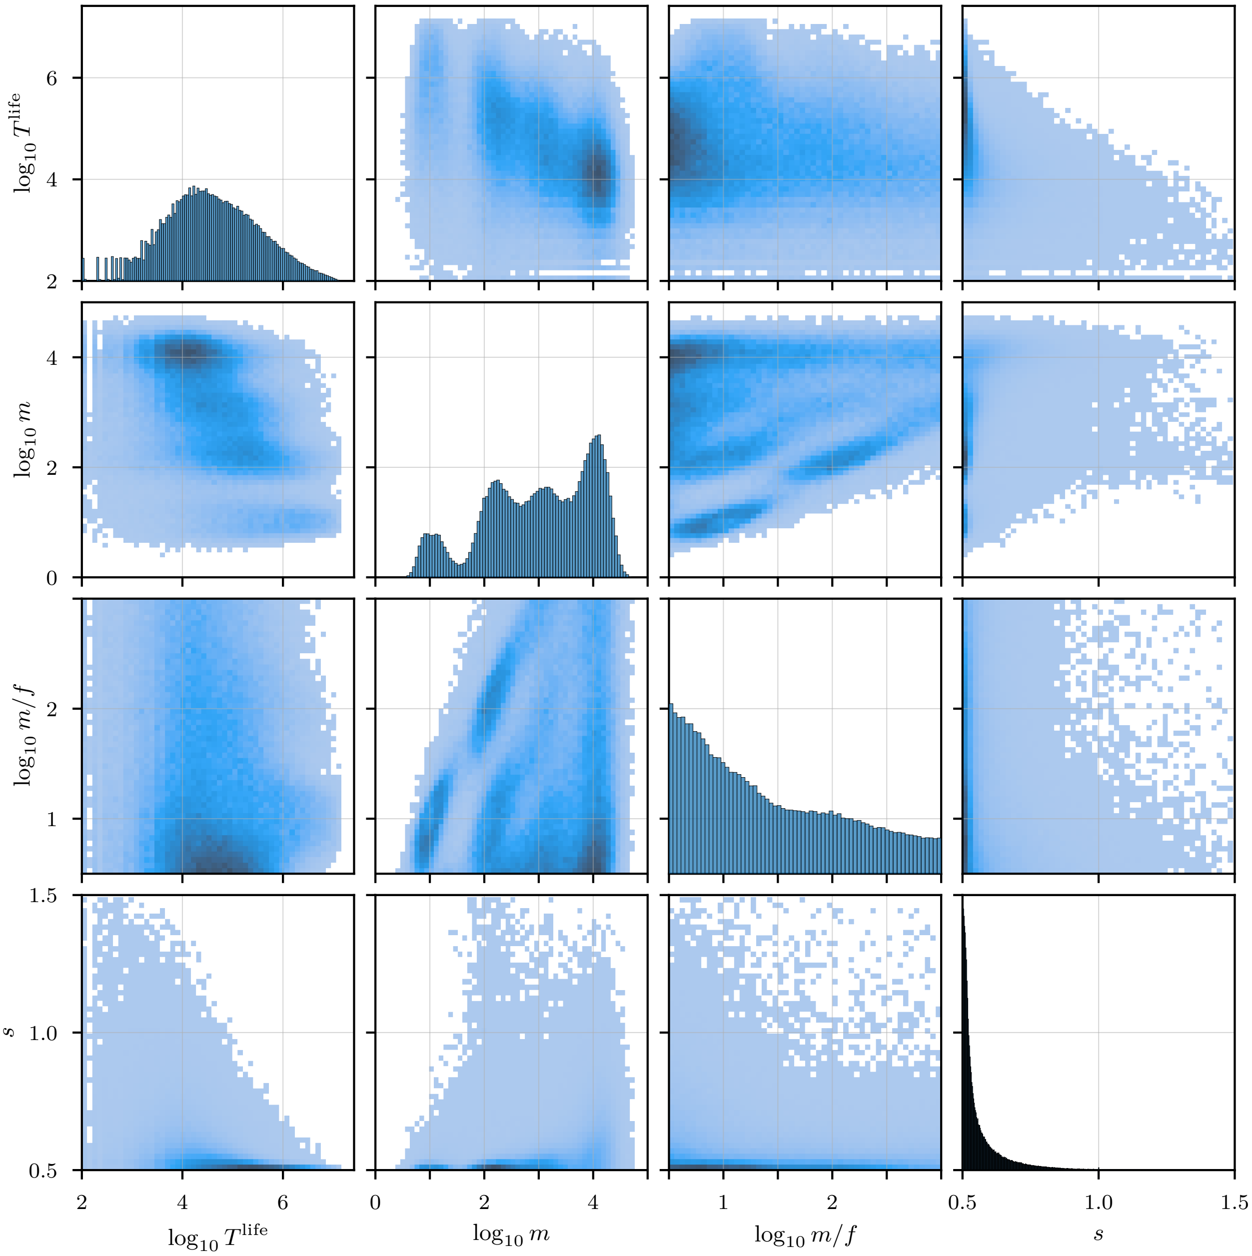

Figure II.5.2 Species life time and traits distributions as univariate (on the diagonal) and bivariate (all other subplots) histograms. Only viable species ($T^{\text{life}} \geq 100$) are shown. Data is from the same simulation as shown in Figure II.5.1.
Note that the univariate plots do not have a corresponding y-axis; the tick labels and horizontal grid lines relate *only* to the bivariate plots. For the univariate plots, the implicit y-axis is in counts, using a linear scale, and is shared among all the other plots on the diagonal.

above. Furthermore, it highlights how the diagonal structures in $(m, {}^{m}\!/\!_{f})$ are dominant only for low-m species, while high-m species have highly varying relative feeding distances ${}^{m}\!/\!_{f}$.

Turning to the species that survive the longest, they almost exclusively are specialists with $s \approx 0.5$, hence having adapted optimally to their prey. Most of these species have a low body mass, as already seen in Figure II.5.1, which is understandable as they will feed on the external resource, which is the most stable resource in the food web.

The univariate distributions in Figure II.5.2 also contribute a new perspective on which species constitute the continuous species turnover that is the namesake of the ECST model: Counting only viable species, there are more high-mass species than low-mass, despite the higher probability of low-mass species to be selected for generating offspring. This suggests that low-mass species optimise quickly and subsequent mutants are unlikely to replace resident species; as non-viable species, their offspring does not count into the shown distributions. Subsequently, it is primarily the higher-mass species that are viable, even with large niche widths. With the lower life time of the higher-mass species, it is primarily this group of entities that constitutes the species turnover in the ECST model.

Having established the trophic structure as one emergent property of this system: are there other emergent properties?

To investigate this question, varying the energy available to the food web can be insightful: As already seen in the study of the response to a reduced carrying capacity K of the external resource, the depth of the food web is primarily limited by the energy reaching the upper trophic levels and whether those species can remain viable. In the context of the whole model, this is additionally affected by the conversion efficiency λ (which affects the steepness of the trophic pyramid) and the extinction threshold $n^{\min}$ (which determines the range of viable number densities).

As expected, food webs with lower available energy are more shallow, see Figure II.5.3; with a rising carrying capacity, species with higher trophic levels become possible. Although lacking in counts compared to the histogram in Figure II.5.1, the long-living species again highlight which trophic levels are more stable than others. Again, the third peak is consistently slightly below $\mathrm{TL} = 4$ for all K. For $K > 100$, the peaks not aligning with integer values is also true for the trophic levels above. This indicates that for higher carrying capacities, trophic coherence seems to be less important for long-term stability and feeding relations are perhaps more general for high trophic levels.

Studying trait distributions for different carrying capacities also reveal something about the life time of species: It appears that the position of the peak in the life time distribution is not a boundary effect caused by the finite simulation time, but an emergent property that depends on the available energy.

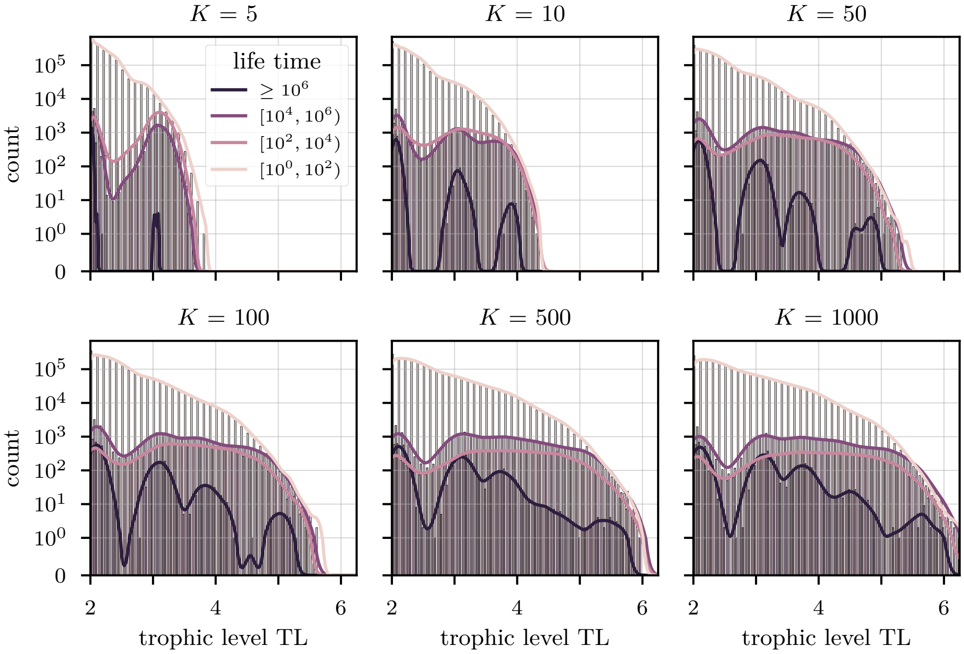

Figure II.5.3 Species life time distributions over their mean trophic level for different carrying capacities K. For each K, 24 simulations over 8 M steps were carried out; each subplot thus shows properties from roughly $2 \cdot 10^6$ mutation events.

Comparing the peak of life time distributions for $K = 100$ (Figure II.5.2) and for $K = 5$ and $K = 1000$ (Figures B.22 and B.23), a shift from $\log_{10} T^{\text{life}} \approx 3.5$ to 4.5 can be observed for increase in K. The peak position correlates with the life time observed for the species cluster with highest m. Taken together, it appears that the peak in the life time distribution is caused by the species with largest turnover, which are those with highest body mass. This high-turnover cluster appears to be more stable the more energy it has available – alternatively, the reason for its increased life time could its position in the food web, tolerating higher variance in its prey composition. With these underlying dynamics, it can be argued that the peak in the life time distribution (a consequence of the trait region with the highest species turnover rate) is an additional emergent property.

The changes in the trait distributions illustrate how the food web structure primarily varies with the available energy, a hallmark of a self-organised system. Importantly, the general structure arises regardless of the available energy; the energy merely controls how large it can grow. The same is true for other model parameters: food web structures can be seen to emerge for a large volume in parameter space[57]. While they vary in their exact manifestation depending on the parameters, the overall pattern (a multi-layered food web with species clusters) is preserved.

To conclude: Yes, the ECST model can be understood as a self-organised system. It generates a multi-layered food web with species clusters and fulfils all common characteristics of a self-organised system: (i) the structures emerge from a simple situation with only a single entity; (ii) they arise from the interactions of the system's entities, not by any outside agency or macroscopic process; (iii) they are driven by a random process which introduces new entities to the system and require no fine-tuning of parameters; and (iv) the diversity of structures depends on the available energy.

---

[57] Of course, scenarios can be constructed where no structure emerges. For instance, by supplying too little energy to the system, there will only be one level above the external resource. Alternatively, reducing mutation sizes can lead to a lock-in effect, effectively prohibiting the exploration of trait space and thus suppressing the evolutionary forcing. In both these cases, there will only be a single level of species feeding on the external resource.

II.5.2 Does the ECST model show self-organised criticality?

The previous section established that the ECST model can be understood as a self-organising system in which a food web structure emerges as the result of the interactions between species and (in many ways) independent of the details of the forcing. In this section, I discuss the extent to which the conceptual framework of *self-organised criticality* (SOC) may be applicable to this system, and where it may be insightful for the better understanding of the self-organising properties of food webs.

II.5.2.1 Self-organised criticality (SOC)

The concept of self-organised criticality was introduced by *Bak and Tang* et al. [1987] and refers to the property of a system to self-organise into a critical state, regardless of the initial state and without need for fine-tuning of parameters. Using a cellular automaton model that represents a heap of sand (known as *sand pile model* or *BTW model* after the authors), they demonstrated that microscopic interactions can give rise to a scale-invariant macroscopic property, akin to a critical point in statistical mechanics, and that the system is attracted to this critical state. They argued that such a mechanism 'might be *the* underlying concept for temporal and spatial scaling in a wide class of dissipative systems with extended degrees of freedom', and that it may offer an explanation for the frequent observations of scale-free patterns in nature.

The concept of SOC has since spread rapidly through many research fields and was applied to systems of various kinds. It was used to describe the dynamics of forest fires [*Malamud* 1998], punctuated equilibria in evolution [*Bak and Sneppen* 1993; *Paczuski* et al. 1996; *Sneppen* et al. 1995], but also extends to neuroscience [*Chialvo* 2010] and many other areas. In the field of ecosystem evolution, SOC has been observed in a number of simple evolutionary models of species interactions, thoroughly reviewed and discussed by *Drossel* [2001].

As an ancillary effect to its success, the vocabulary and understanding of SOC varies strongly between research fields and individual applications. *Watkins* et al. [2015] summarise the different ways in which SOC is perceived and used in the literature: In the most minimal formulation (following [*Bak and Tang* et al. 1987]), SOC refers to mechanisms that lead to a *self-tuned phase transition*. Others see SOC as the mechanism that causes fractals in nature – or even any form of power law. Finally, the 'visionary' interpretation is that all form of contingency in nature is caused by SOC. These different perceptions naturally lead to heated debates and misunderstandings regarding self-organised criticality [*Buchanan* 2015; *Watkins* et al. 2015].

Following *Roth* [2020], I will refer to the original formulation as *strict-sense SOC*: a self-organising, weakly-forced threshold system that is attracted to a scale-invariant critical state which can be described by power laws. In addition, *Roth* introduces *weak-sense critical* systems as those with a stronger forcing which still exhibit a robust critical state, albeit not necessarily scale-free.

II.5.2.2 SOC in food web models

As *Watkins* et al. [2015] note, many claims of systems exhibiting SOC are footed on the observation of power law distributions, suggesting that the cause of such a distribution is an underlying self-organised critical state. Yet, scale-free behaviour alone does not suffice as evidence for strict-sense SOC [*Watkins* et al. 2015], but requires a more careful consideration of the involved mechanisms.

In the field of ecosystem evolution, the sizes of extinction events extracted from the fossil record have been found to resemble power law distributions [*Raup* 1986]. This observation – despite being low in statistical power and also representable by an exponential function – lead to a number of studies that aimed to describe these findings using models of SOC systems [*Drossel* 2001]. As reviewed by *Drossel*, these models are conceptually close to threshold systems like the sand pile model. Some of the reviewed models (for instance, [*Paczuski* et al. 1996]) clearly show strict-sense SOC, while the mechanism that could give rise to SOC is less clear in other models. Based on the description as SOC systems, the models were able to illustrate how large extinction events may arise from the intrinsic dynamics of ecosystems and not necessarily require external causes.

However, these models do not include many of the mechanisms known to play an important role in ecosystem evolution. Instead, they make simplifying assumptions like associating fitness values with species which then determine the whole systems' dynamics. In particular, they do not allow for smooth adaptations to changes nor do they attempt to represent the underlying dynamics of the interactions within ecological interaction networks. Food web models like the ECST model differ in many ways from these simple evolution models and specifically include these two aspects of ecosystem dynamics: They do not prescribe fitness values, but a species' survival is determined by the explicitly modelled interactions within the system; in the ECST model, these are the trophic interactions.

With the increased level of ecological complexity represented by these models comes a departure from the conceptual proximity to threshold-based models. For instance, by no longer explicitly representing fitness – which indirectly constituted the threshold in those models –, some other part of the model would need to be associated with a threshold; it is not immediately clear if that part would behave in a comparable way with respect to SOC. In effect, it becomes more difficult to argue how observed scale-free behaviour is linked to SOC. To understand whether SOC plays a role in these more complex food web models, model mechanisms need to be dissected and it needs to be evaluated how they may give rise to SOC.

As an example, the niche-based food web models introduced and studied by [*Guill and Drossel* 2008] and [*Allhoff and Ritterskamp* et al. 2015] show power law

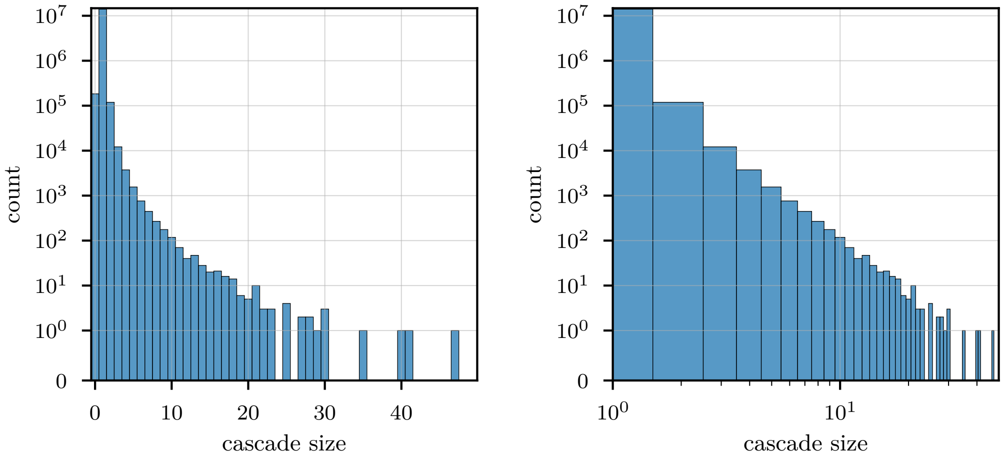

Figure II.5.4 Cascade size distribution from the same simulation as shown in Figure II.3.17 in lin-log **(left)** and log-log **(right)** representation. The distribution shows how many species became extinct as a result of the roughly $14.4 \cdot 10^6$ speciation events. The entry at zero denotes the introduction of viable species that led to zero extinctions; the peak at 1 consists mostly of non-viable species and their own extinction (cascade of size 1).

behaviour in the distribution of extinction cascade sizes. The authors of both studies (carefully) suggest SOC as a possible explanation for this observation, but omit to motivate the mechanisms by which SOC may arise there. However, given the conceptual differences between these food web models and the simple evolution models in which SOC was mechanistically plausible, a more detailed explanation would be needed for why SOC may also govern these more complex models.

II.5.2.3 SOC in the ECST model

As argued above, it is unclear by which mechanisms models like the ECST model may exhibit SOC. Below, I will attempt to isolate the parts of the model that are relevant for SOC, starting with the observations of power law behaviour.

Figure II.5.4 shows the distribution of extinction cascade sizes and indeed resembles a power law. This observation appears to be robust under changes to K (see Figure B.24), which does not appear to systematically affect the distribution. As also noted by *Allhoff and Ritterskamp* et al., the power law is very steep, with relatively few extinction cascades displacing more than 10 % of species in the food web and hardly any cascades being so large that more than a quarter of species become extinct. The authors conclude that ‘this is not the type of SOC required to explain the large extinction events in earth history, where up to 90 percent of all species went extinct’ but remain vague regarding other explanations for this observation.

However, SOC is not the only mechanism by which power law distributions may arise [*Buchanan* 2015]. For instance, as shown in [*Touboul and Destexhe* 2010] using avalanches in neuronal models, 'power law distributions may be a generic property of thresholded stochastic processes'.

There is another observation that suggests that it is not SOC that is causing the power law distribution in the cascade sizes: In the original argumentation by *Bak and Tang* et al., the lack of characteristic spatial scales is equated with the memorylessness on the temporal scale, i.e. an exponential life time distribution. The life time distributions shown in Figure II.3.17 exhibit a kind of distribution that is distinctively different from an exponential distribution, showing a peak at roughly 10^4 iteration steps. This suggests that a different process is governing the extinction of individual species.

Hence, a change in perspective may be useful: Given the considerably more complicated mechanisms of the ECST model and the possibility of alternative explanations for the observed distribution, which aspects of the model suggest that a power law distribution *would* be expected at all? In other words: Even if the system dynamics show SOC, can it be expected that this would be observable in the same unobscured way as in the sand pile model?

Comparison to sand pile model

To elevate this argumentation from the sole consideration of power-law-distributed observables, I will compare the ECST model mechanisms with those of the classical sand pile model [*Bak and Tang* et al. 1987]. This association may at first sight appear far-fetched, but there are surprisingly many properties of the sand pile model that find an analogy in the ECST model. Where a process lacks a suitable analogy, differences between the two become more tangible.

The sand pile model is an abstraction of grains of sand (the system's entities) in a spatial domain; by adding new grains (the forcing), a locally defined critical slope may be exceeded, causing toppling of the grains into the immediate neighbourhood. This effect may cause yet further toppling, thus leading to avalanches of different sizes.

In the ECST model, the entities are the individual species, which are localised not in a spatial domain but in a three-dimensional trait space (m, f, s), determining all species properties and interactions. The species are coupled via population dynamics, relaxing towards the equilibrium biomass densities of the current configuration. The mutation events can be understood as the forcing: new species are introduced into the system and may cause the extinction of other species once they drop below the extinction threshold. This in turn may cause other species to become displaced, thus leading to an extinction cascade.

Along these lines, most other parts can be associated with each other, which is done in Table II.5.1. In the following, I will go through some of these in more detail. Given the intertwined nature of the mechanisms, the argumentation soon encompasses most aspects of the model, illustrating why a holistic view of the model is relevant.

Table II.5.1 Comparison between the sand pile model and the ECST model, drawing possible analogies between model mechanisms in the context of self-organised criticality. Refer to the main text regarding the extent to which these analogies hold.

	Sand Pile Model	ECST Model
model entities	grains of sand	species
domain	spatial (typically 2D), uniformly discretised grid	trait space (3D), continuous
entity state	slope (or: height), discrete, small	biomass, continuous
system size	large (typically $\gg 1000$)	small (< 100 entities)
coupling[a]	via toppling (uniform, directed, local)	via population dynamics (structured, directed, non-local)[b]
forcing	adding a grain of sand to the domain (weak)	adding a species[c] to trait space via a mutation event (weak)
dissipation	grains removed at boundary	extinct species removed (incl. biomass)
threshold	critical slope	extinction threshold
relaxation	grains topple until all are sub-threshold	population dynamics: redistribution of biomass until equilibrated
absorbing phase	grain is added (no toppling)	mutant species coexists with residents (no extinctions)
active phase	grains topple, potentially causing further toppling (avalanche)	a species becomes extinct, potentially causing further extinctions (cascade)
time scales[d]	fully separated	approximately[e] separated
control parameter	average slope	FD? average number density?

[a] Denotes the possible interactions between entities.

[b] Species have more or less direct coupling via trophic interactions, but indirectly all affect each other.

[c] As a result of a speciation, i.e.: biomass is redistributed from the parent to the offspring, allowing to regard this as a conservative forcing.

[d] Refers to the time scales between forcing and relaxation.

[e] In almost all cases, relaxation is very fast; in few cases, it may overlap with the introduction of the next mutant.

System size and control parameter One of the most prominent differences between the two models is the system size, with the sand pile system typically comprising many thousands or millions of entities while ECST model systems rarely exceed 100 entities. Larger systems make it possible to observe the critical state via macroscopic observables that average over the microscopic state (akin to thermodynamic observables like pressure or temperature); in small systems, fluctuations may be too large to make such an observable discernible.

For the sand pile model, this macroscopic control parameter is the average slope. It is a robustly emerging property with relatively small fluctuations and its variance is further reduced for larger system sizes. Adding a grain to the system will temporarily increase the average slope; avalanches that redistribute grains within the domain will cause no changes, but avalanches that touch a boundary will lead to dissipation, hence creating a negative jump in the average slope. While the response of the sand pile to a new grain still depends on the exact micro-state of the system, a large average slope makes it more likely that large avalanches are triggered.

In contrast, the ECST model typically contains fewer than 100 species; at times, the system may comprise perhaps a few hundred species, depending on the available energy. While increasing the system size would in theory be possible by increasing the available energy, going beyond many thousand entities would quickly become unreasonable in the context of food web ecology.

Regardless of that choice, the ECST model and other food web models can be considered to be small systems. With few entities, it can be presumed that any macroscopic quantity that could characterise criticality will be subject to strong fluctuations or not manifest as a narrow-enough distribution to be detectable as the critical observable.

Can a suitable control parameter even be defined for the ECST model? Given the small system size, any such quantity will certainly not have the same characteristics as the average slope. Considering the observables used for studying the ECST model so far, there is no obvious choice for a suitable analogy to the average slope in the sand pile model; but there are two candidates which I want to briefly bring up here.

One candidate could be the average number density, which is a direct measure of how close the species are (on average) to the extinction threshold: $n^{\mathrm{min}} - \langle n \rangle$. Apart from the sign, this is structurally equivalent to the order parameter of the sand pile model and a seemingly appealing analogy: Naively, adding a species to a system where a common resource is shared would suggest that the average number density is reduced, thus approaching the threshold; species extinction in turn would lead to the opposite effect. However, the ECST model behaves differently due to its non-linear dynamics and the trophic structure: Adding a species may cause a configuration in which overall more biomass can be foraged from the external resource, leading to an increase in the average number density instead of the naively expected decrease. From observations of the behaviour of this observable (not shown), it appears as if it would neither capture the likelihood of a large cascade.

Another candidate is the functional diversity, which by definition grows with every added species and drops with every extinction event (as does the average slope). It also appears to be more robust with respect to the size of extinction cascades: As seen in the perturbation studies in Section II.4.2.3, the 'volatile' states in fact have a larger average FD (see Figure B.17) than the random sample – as they were selected because of the large cascades they exhibited, this suggests that FD behaves like the average slope in this respect. However, the change in FD depends on the (lost or added) contribution to the set of feeding kernels and is thus not a bijection to the number of extinct (i.e., dissipated) species, as is the case for the change in average slope[58].

As a combination of the two, the number of species may also be seen as a control parameter: It quantifies how many species have to split the available energy, while also being correlated to FD. However, as seen from the study of the food web observables, the species number is strongly fluctuating between instantiations (cf. Figure II.3.16), calling into question whether it may act as control parameter for individual simulations.

To summarise the above, the two systems strongly differ in their size and, correspondingly, in how well a macroscopic control parameter can be defined. In the ECST model, it is not clear which observable would act as such a control parameter and the presented candidates exhibit very wide distributions and only a vague resemblance to the well-defined emerging property of the sand pile model. This is not to say that the ECST model may not have such a control parameter at all, but only that it is not sufficiently discernible at its typical size.

Effective size and long-range coupling Above, the number of species was equated with the system size in the ECST model. However, it has to be taken into account that the state of all species in the food web is coupled via population dynamics. Introducing a (viable) mutant to the system thus leads to changes in equilibrium biomass densities across the whole system, even if there are no extinction events; these correlations between species equilibrium biomass densities in turn lead to correlations in the case of extinction events. This is unlike the sand pile model, where any non-local changes may be mediated solely via avalanches – and only in the event of toppling, not prior to that.

Subsequently, the species cannot be considered as completely independent entities in the same way as the sand grains are, which may be understood as a reduced effective system size. These correlations may be between distributed species in the food web (akin to the alternating pattern seen in the study of species invasion in Section II.4.1) or between species that have very similar traits, hence effectively behaving in the same way but with individually reduced number densities.

These considerations also highlight how the ECST model is not a threshold system in the same way as the sand pile model (or the Bak-Sneppen model [*Paczuski* et al. 1996]) is: the former includes long-range coupling even in the absence of threshold-crossing-events, while the latter system's long-range effects are restricted to those events. This puts into questions whether the analogy made between species and grains of sand is suitable at all.

---

[58] Of course, instead of counting cascade sizes via number of species, it could be defined as ΔFD instead, which would trivially lead to a bijection. The corresponding distribution was already shown in Figure II.3.19.
To not go deeper into the rabbit hole, this thought is not further elaborated here.

Boundary effects and domain structure In addition to the above points, the system size plays a role in determining how pronounced boundary effects will be; close to a boundary, they may be strong enough to distort or obscure a critical observable. Possible boundary effects are also a good context in which to reflect on the differences in the domain structure of the two models.

In the sand pile model, boundary effects considerably affect the size distribution of avalanches, such that (mostly) uninfluenced power law behaviour can only be seen for small avalanches on large grids. For smaller grids or larger avalanches, the finite resolution of the grid and the finite size of the domain lead to deviations from the power law in the complementary cumulative avalanche area distribution. In the latter case, the finite size leads to large avalanches frequently going beyond the domain and thus being counted as smaller ones, hence distorting the distribution for larger values. However, scale invariance in the sand pile model can be shown empirically by mapping the complementary cumulative probability distributions of differently-sized grids onto each other using a universal scaling function [*Kadanoff* et al. 1989; *Roth* 2020]. Hence, the sand pile model not only shows an emergent property to assess criticality, but its scale-invariance can additionally be shown by correcting for finite-size effects.

What would constitute boundary effects in the ECST model? Or, before that: what are the boundaries of the domain? In principle, the trait space is unbounded in m and f, but it is effectively restricted to regions where sufficient energy can be obtained: with the energy source at $m = 1$, much smaller f become infeasible; also, with the extinction threshold being formulated as a fixed number density $n^{\min}$, species with large m are not viable. The mutation rules further impose boundaries in $^{m}/_{f}$, requiring $m \gtrsim 3f$ and thus creating some form of ordering in trait space. Additionally, the niche width is limited to a fixed interval.

However, the non-spatial domain and boundaries are conceptually different to their spatial counterparts in the sand pile model. First, extinction cascades are associated only with a species count and cannot be associated with a volume in trait space. Second, as touched upon earlier, the cascades can include species from all parts of trait space, meaning that there is no locality as there is with avalanches, and hence no unevenly counted cascade sizes. Finally, while the addition of a new species – the potential trigger of an extinction – is affected by these boundaries, they play no role during the cascading of extinctions through the system; this process only depends on the population dynamics and the configuration of species in trait space.

On that backdrop, it seems that there is no mechanism by which there could be a direct effect of the boundaries on the structure of the extinction cascades. However, the effect of these boundaries pertains primarily to the influence they have during the absorption phase, where they determine the region in which a species may be introduced and hence: the interaction structure of the species.

Interaction structure and scale In the sand pile model, the interaction rule is exactly the same everywhere and it is only the micro-state that determines whether toppling will occur or not. Further, these interactions are purely local: any long-scale interaction is mediated via local topplings.

In the ECST model, this picture is less clear. On the one hand, interactions are defined in the same way for all species but the resource, Equation (II.2.11). On the other hand, the food web essentially has N_S^2 manifestations of individual species-species interactions at any time; while they are all defined by that same 'rule' (and many are effectively zero), the interaction *in principle* depends on all other species' traits and biomass densities: essentially the whole system's micro-state. This not only makes it effectively non-local, but introduces temporally changing heterogeneities that are shaped by the position and size of species clusters in trait space.

Hence, this is another point where an extinction cascade is not like an avalanche: Unlike an avalanche, which is propagated entirely via local interactions and covers a contiguous area, an extinction cascade may occur in any place in trait space and is affected by the whole system's configuration. The avalanches show long-range spatio-temporal correlations, depending on the system size; in contrast, the domain in which extinction cascades propagate is overall much smaller (or: strongly coupled) such that distinguishing into short- and long-range correlations might not be reasonable.

Furthermore, the heterogeneities depend on the chosen feeding distances and niche widths: As shown earlier in Figure II.5.2, viable species generally have niche widths close to the lower boundary. In addition, feeding distances of $^m/_f \approx 10$ give rise to the regular pattern in the body mass m of species clusters[59].

One conclusion from these observations about the domain and the boundaries is that the ECST model can be said to exhibit a scale in its domain: The arising trophic clusters are roughly equally-distanced and influenced by the possible mutation values for $^m/_f$ – which is a hallmark of the self-organisation of the model (see Figure II.5.2).

The effect of patterns in the domain may be illustrated using the *contagious disease model* (CDM) [*Roth* 2020], essentially a variant of the forest fire model where the time scales of clusters burning down (the avalanche) and new lightning strikes (the forcing) are no longer separated. This model may exhibit spatial pattern, which then lead to a distorted (i.e., no longer scale-free) cluster size distribution. Subsequently, the CDM is classified as a weak-sense SOC [*Roth* 2020].

However, it is yet another question whether a pattern in the trait distribution (the levelled structure) of the ECST model could have an effect on the cascade size distribution via correlations over the range of the domain. In the CDM, the connection between the spatial patterns and the cluster sizes is straight-forward; in

---

[59] While that feeding distance is not directly at the lower bound, it is certainly affected by it: a larger bound would lead to an increased distance between clusters in m.

the ECST model, however, further investigation is needed to find out how responses across the domain may be correlated.

With the observations surrounding Figure II.5.2, such as the species turnover occurring primarily in high TL species, cascade sizes might also be limited to one of the more active clusters, hence not being affected by the scale. Alternatively, extinctions may occur throughout the whole food web pyramid and be primarily affected by the trophic level distribution (more species on lower levels, hence more extinctions there) – in such a scenario, the trophic level distribution of individual cascades would not depend on the scale value per se, because a trophic interaction necessitates merely a matching feeding kernel.

Wrapping up these points, the interaction structure of the ECST model is considerably more complicated than in the sand pile model; this alone may warrant that whatever distribution results from this would not be scale-free but in some way distorted. In addition, the effect of the emergent scale of the system (distance between heavily populated species clusters) on the cascade size distribution is not obvious and requires further investigation, e.g. by studying the correlation patterns within the domain. Further research in that direction could also take into account how the parent selection algorithm affects this aspect.

Summary

Following a line of questioning whether the ECST model exhibits SOC, I reviewed a number of mechanisms of the model and compared them to the sand pile model. Rather than arguing how the observed power-law-distributed cascade sizes could be caused by SOC, I reflected on whether such a distribution would be observable at all, even in the case of SOC.

The above considerations exemplify how the parts of the ECST model mechanism are not only far more complicated than those of the sand pile model, but also not easily abstractable to that simpler scenario. Most of the analogies drawn between the two models turned out not to hold upon more thorough examination. Overall, food web models like the ECST model are self-organising threshold systems, but they still differ considerably from the comparably simple scenario that is the sand pile model or the Bak-Sneppen model [*Paczuski* et al. 1996]. This pertains in particular to the non-spatial domain with a system-wide coupling and the small effective system size which impedes definition and observation of useful control parameters.

As a consequence of this departure from the well-understood sand pile model, finding mechanistic arguments for why the studied food web model may show SOC become very difficult. Even presuming SOC in the ECST model, it would be unexpected if it would manifest as a perfect scale-free distribution in cascade sizes, mainly because the response of individual species in the food web is correlated either via trophic interactions or via a multiplicity in trait space. As suggested above, further studying the correlation pattern between extinctions could help to better understand this aspect of the interaction structure.

It is clear that a core property of food web models is to self-organise into exactly that structure: a layered network of trophic interactions. One aim of this discussion was to isolate in how far this may be accompanied by attraction towards a critical state. While the analysis showed that there are candidates that could be understood as such attractors, it is not clear what the robust state of the ECST model would be and whether these observables can be expected to capture this property, in particular given the small system size. In that sense, this discussion might have brought forward more questions than answers.

In closing, I propose that classifying the ECST model (or similar models) as 'SOC' or 'not SOC' is not in itself useful, but needs to be contextualised with how these effects manifest. Asking the corresponding questions en route to understand whether a system *may* exhibit SOC certainly offers valuable insights, regardless of the category it would be ascribed to; this might even be regarded as more important than the classification itself.

II.6

Discussion

In the previous chapters I studied the general food web evolution of the ECST model, its response to perturbations, and its capacity to self-organise. Here, I will discuss a selection of these aspects and give an outlook on possible future research along the lines started in this thesis.

In particular, I will discuss the repercussions of the assumptions made for the formulation of the ECST model mechanisms, which played a role in several parts of the results. In that context, I will also reflect on more general aspects of niche models. Finally, I will propose possibilities for how evolutionary models of ecological interaction can (and need to) be extended to better study questions pertaining to the resilience of food webs.

II.6.1 The evolution mechanism of the ECST model

The way evolution is represented in the ECST model (see Section II.2.2.2) is of course far abstracted from evolution in real ecosystems, where it occurs on the level of individuals or groups of individuals. The view taken in the ECST model and other niche models is on the level of the species and – correspondingly – a heuristic model: Instead of representing individuals which diversify over time and, at some point, become associated with different species, these speciation events are directly imposed on the system. This approach becomes necessary as soon as the underlying individuals are not explicitly modelled but abstracted into groups of sufficiently similar individuals (species).

The mechanism that represents these speciation events in the ECST model is characterised by the following properties: (i) parent species have the same probability to be selected, (ii) the offspring's body mass will be similar to that of the parent, and (iii) variation in other traits is high, leading to many non-viable species. Thus, it is an inheritance channel with comparably low fidelity. As already remarked in [*Allhoff and Ritterskamp* et al. 2015], these choices have an effect on the speed of exploration of trait space: With higher correlation between the traits of parent and offspring species, the variation is reduced, reducing the speed with which new traits can be 'discovered'.

Given the low inheritance fidelity in the ECST model, the introductions of new species in the ECST model are more similar to invasion events, but with a certain bias towards the prevalent body masses in the system. In contrast, the alternative mutation mechanism (briefly studied in Section II.3.4.1), where parent and offspring properties are very similar, exemplifies how a high-fidelity inheritance would act on such a system.

However, given the difficulties observed alongside the alternative mutation mechanism, it is a pragmatic choice to restrict the ECST model to a scenario with high rates of species invasions. While many of these species are not viable, some of them are; these viable species either cause a quick exploration of trait space which expands the food web, or they trigger an extinction cascade that restructures the whole food web. Overall, this is what causes the constant species turnover of the ECST model. In an essentially identical way, other niche models become less persistent if modelling invasion rather than speciation [*Romanuk and Binzer* et al. 2019].

Another aspect is the random probability of a parent to be selected to generate offspring, which introduces a bias towards the more populated low trophic levels. As seen in Figure II.3.18, however, these offspring are often non-viable, balancing out the selection bias. In effect, the relevant turnover occurs on the higher levels. In [*Allhoff* 2015, ch. 4.5], modifications to this mechanism are suggested, for instance by weighting the probability by the number density – however, this would make high-m species more unlikely to reproduce, further cementing the bias and reducing the occurrence of higher trophic levels. If selecting the parent weighted by m or by inverse generation time $m^{-1/4}$, the resulting food webs were observed to be qualitatively equivalent, but with a faster or slower time development, respectively.

These assumptions affect the range of questions that can be studied with the ECST model. In the context of this thesis, this pertains primarily to the resilience of food webs with different structural properties. As observed in Sections II.4.2 and II.4.3.2, perturbed food webs retain very little information of their previous structure. On top of that, with the mutation mechanism (representing the food web's recovery options) being as stochastic as it is, it is not unexpected that the study of food web resilience did not show discernible effects (or effects going beyond the sampling bias from the parent selection).

Overall, the ECST model includes few mechanisms that can persistently propagate information into its future. The only information it holds is in the species existing at a point in time, and that information is constantly perturbed with species invasions or lost upon extinction cascades. While these mechanisms give rise to the macroscopic, self-organising patterns observed over many instantiations and long time, the variance on short time scales is too large to use the ECST model as a basis to study food web resilience.

II.6.2 Towards more complex models of ecosystem evolution

One of the aims of the study of food web evolution throughout this thesis was to contribute to the understanding of what makes a food web resilient against adverse effects. The studies presented in this thesis showed that the ECST model system had only very limited response options and thus was not a suitable model system for properly investigating this particular topic, for instance due to the particular choices made for its evolution mechanism. However, these limitations illustrated a number of points that are informative about the conceptual requirements on a model system to study these questions. Together with Evolution Mechanics, this can assist in formulating models that are more suitable to study food web resilience.

I would argue that in order to better understand the evolution of ecological interaction networks, the complexity of these systems needs to be accounted for in some way or another. Naturally, reducing these systems to models that act basically on a single dimension (as done in niche models with the assumption of size-ordered feeding interactions) makes them easier to investigate and is a necessary step towards better understanding them. However, as seen throughout this thesis, the limitations of this approach appear when it comes to recovery mechanisms. Such mechanisms would arise from interactions with the environment or from evolutionary processes – but not from population dynamics alone. Because niche models typically use highly reduced evolutionary dynamics and a simplified representation of the environment, they are not able to provide a reasonable mechanism for food web recovery.

An interesting observation at this point is that the evolution mechanisms used in niche models seem to suffice to generate approximations of food web structures [*Allhoff and Ritterskamp* et al. 2015; *Fritsch* et al. 2021]. Furthermore, an influential study by *Eklöf* et al. [2013] suggests that the number trait dimensions needed to accurately describe a food web are surprisingly low: in many cases, fewer than five traits were enough to explain most of the interactions in the studied food webs.

However, a distinction should be made between the structural description of a food web and the mechanisms that play a role in its reorganisation. The above limitations seem to come into play only when investigating the dynamics of the model after an adverse influence. While few dimensions may suffice for a structural description, it is not clear why the same should be true for the reorganisation mechanism.

In the case of the ECST model, the recovery mechanism is basically equivalent to the general operation of the model and the situation after a perturbation is equivalent to a system that was initialised with a smaller size. However, it can be presumed that real ecosystems retain some information about previously existing interactions which then plays a role in the system's recovery – something that the ECST model is not capable of representing.

In the following, I will propose ways by which the environmental and evolutionary complexity of models of ecological interactions can be increased, with the main idea being that the higher complexity allows for more sensible recovery mechanisms. With 'more complex' I am *not* referring to a necessarily more *detailed* representation[60], but to models which include a larger amount of heritable information, more mechanisms that depend on these species properties, and a more important role of the environment in the interactions. As such, these extensions cover a larger range of possible interactions.

Having studied the ECST model in detail, I will align the extensions to that model where possible. Many of the points can also be seen detached from that particular model and also apply to niche or food web models in general. Thus, the two questions for this part of the discussion are: How could the ECST model be adapted to more persistently propagate structural information and, at the same time, retain the property of continuous species turnover? How could food web models in general be extended to take more of their evolutionary history into account and give rise to more sensible recovery mechanisms?

Answering these questions could be the foundation for further investigation of food web resilience with the help of extended niche models. As such, this section can be considered to be an outlook on future research in this field, informed by Evolution Mechanics[61] and the experiences made in the study of the ECST model. In closing, I will aggregate the open questions arising from these discussions, including the question of how increasingly complex models can be approached.

II.6.2.1 Expanded trait space and interactions

An easy way to increase the information content of a food web model is to extend the dimensionality of the trait space such that each species encodes more information. However, expanding the trait space is only sensible if the additional traits have a mechanistic meaning in the context of the model and the selection mechanism may act on them.

In niche-based food web models, traits encode primarily the feeding interactions; the central idea is to simplify these interactions by representing them all on the body mass dimension of trait space [*Fritsch* et al. 2021]. Thus, the information encoded by these traits is low: In the ECST model its (m, f, s); in models like studied in [*Ito* et al. 2009] an additional trait encodes the width in niche space as which a species is perceived when it is the prey.

---

[60] More detailed representations are chosen in models like the Madingley model [*Harfoot* et al. 2014; *Purves* et al. 2013], where one aim is to make quantitative predictions. Notably, that particular model does not include an evolution mechanism at all, presumably because the time scales are deemed too short.

[61] The processes suggested below can be regarded as applications of the general considerations about evolvability made in the Evolution Mechanics discussion (Section I.4.2.3). While they are not formulated on a detailed operational level, they form a bridge between Evolution Mechanics and applications in food web models.

Arguably, it appears to be difficult to extend these dynamics by further interactions, given how prominent the feeding relations are in niche models. The conceptual hurdle here is that there are direct physiological associations made for the (m, f, s) traits – assuming that these are the most dominant aspects determining feeding relations. While this does not exclude that feeding may be determined by other species traits, these models put the focus on size-ordered (encoded by body mass) trophic interactions.

One proposition to arrive at an expanded trait space would be to add more interaction types to the system. For each additional interaction, additional traits would be added to a species, increasing the dimensionality of trait space. If the interaction is similar to trophic interactions but perhaps with a focus on a certain resource type, the traits would again represent the source of some resource flow (in analogy to m) and the 'feeding kernel' along that dimension (in analogy to f and s).

Of course, these additional interactions would not need to be motivated by trophic interactions determined by body mass but other story lines are also conceivable: As one example, consider that prey may be in hiding at certain times of the day and predators might not be foraging constantly but only during some interval – this would be a circular interaction dimension where time intervals need to overlap in order for interactions to actually occur. Hence, such an interaction would modulate foraging success on the existing trophic interactions. As a result, it would allow for two prey species to coexist simply because they feed on a resource at different times of the day; in addition, prey species would have an evolutionary advantage if they have an expanded hiding period (but also a disadvantage if hiding for too long, e.g. by not being able to forage themselves).

As another example which adds a *parallel* interaction rather than a modification, consider mutualism between two species, a frequently-studied topic in food web ecology [*Cai* et al. 2020; *Olff* et al. 2009; *Rohr* et al. 2014]. The additional traits would need to encode whether two species engage in mutualism or not (or the degree to which they do). The effect of a mutualistic interaction may be that their respiration rates are reduced – but the mechanism could be designed such that these benefits would only be there for a certain ratio of biomass densities between the two, otherwise the mutual benefit turns into a disadvantage in the form of higher respiration rates for either or both of the species. Thus, the resulting interaction network would have two types of edges: trophic edges denoting resource flow and mutualistic edges which influence the species' dynamics. As suggested in [*Pilosof* et al. 2017], such parallel interaction dynamics can be understood as coupled ecological networks or multi-layer networks.

In general, additional interactions might play an important role in the stabilisation of ecosystems [*Grilli* et al. 2017]. These can also be achieved by introducing recycling dynamics or ecological feedbacks, allowing autocatalytic or inhibitory processes in the interaction network [*Veldhuis* et al. 2018].

These examples are ecologically motivated mechanisms that would expand the trait space and the range of different interactions. In essence, they introduce additional abstracted niches into the system, but they do so by adding dimensions to niche space, not by compressing the existing niche space (which could be achieved by smaller niche widths, for instance). This is a qualitative difference that substantially affects the evolutionary dynamics. For instance, in a trait space with two effective dimensions, an evolutionary branch may evolve to some other point in trait space without coming into conflict with another species – in the one-dimensional case, this is not possible.

II.6.2.2 High-fidelity inheritance, yet sufficient variation

The high stochasticity of the regular mutation mechanism of the ECST model was repeatedly stressed above, particularly because it reduced the amount of information that could persist beyond a certain number of mutation events. At the same time, variation needs to be sufficiently high in order to introduce species with novel trait combinations into the system. Both these concepts were discussed throughout Evolution Mechanics and, naturally, also apply to food web models.

To increase the inheritance fidelity, making offspring traits more similar to those of the parent proved to be successful in principle (see Section II.3.4.1), but introduced an operational challenge in the form of many species with very similar properties coexisting and driving up computation times. To maintain feasible computation times in that scenario, the number of nearly identical species would need to be reduced. This can either be achieved via an approximation (identifying similar species and merging them into a species with averaged traits), or by adapting the population dynamics such that a weaker form of intra-specific competition also acts on nearly identical species, displacing them more frequently [62].

While this alleviates some of the problems with low-fidelity inheritance, it needs to be taken care that there is sufficient variation introduced into the system, otherwise the systems become prone to run into an evolutionary dead end. This is also influenced by the mechanism with which the parent species is selected (discussed above), which can strongly influence the composition of the food web.

Furthermore, any additional interaction mechanisms also needs to be inherited with sufficiently high fidelity. The challenge in modelling a system with increasingly many traits is that every mutating trait comes with parameters defining sensible mutation strengths and bounds; especially if the traits encode qualitatively different mechanisms. As seen with the bounds of the niche width parameter in the ECST

---

[62] The latter method picks up a problematic aspect of the definition of the functional response in the ECST model (as remarked on by *Allhoff* in personal communication with me): With higher competition, the denominator of the functional response becomes larger and the overall consumption is reduced. With the story line of interference competition this is sensible, but it weakens the relative competition pressure and thus promotes the coexistence of many similar species.

model – and highlighted in [*Fritsch* et al. 2021] regarding similar approaches in other niche models –, the choices made regarding these bounds have strong effects on the observed outcome. They can be interpreted as some sort of fine-tuning, which is an unfeasible approach if aiming to extend models to be more complex.

Of course, these parameters are the result of a heuristic that aims to represent inheritance in food web models on the level of species instead of on the level of individuals. An ideal and more general inheritance model would aim to imitate what occurs in natural evolving systems, where mutations act on the genotype of individuals, potentially causing changes in phenotype; a species can then only be defined from a larger number of sufficiently similar individuals. The mutation mechanism in such a case would be comparably simple and take the shape of point mutations in the genotype occurring with a certain rate. The actual complexity lies (as so often) in the form-function relationship and the genotype-phenotype map: this is what determines whether a mutation causes some phenotypic changes and whether they lead to changes in selection pressures.

However, such an approach is not attainable with food web models, because too many temporal and spatial scales would need to be bridged, ranging from the single individual and their life cycle to the long-term evolutionary development. Imitating natural inheritance mechanisms would simply not be feasible to simulate. In other words, these models (and models of evolutionary systems in general) need to balance the detail with which inheritance mechanisms are represented with the conceptual cost of representing them and the computational cost of simulating them. This leads to the need for heuristic inheritance mechanisms.

Can fine-tuning be reduced even if heuristic inheritance mechanisms acting on the level of the species need to be used?

One possibility is inspired by the self-modifying nature of evolutionary systems: By incorporating the mutation parameters themselves into the genotype and allowing them to mutate upon inheritance, the parameters could coevolve with the rest of the system – at least in principle. It is likely that there are a number of operational difficulties in doing this; for instance, it would need to be prevented to have mutation rates evolve to zero or diverge altogether. In essence, such an approach shifts the point at which decisions regarding the mutation parameters need to be made. It may make some decisions easier if they can be made on the mutation parameters and not on the traits; however, deciding on bounds for trait values could be informed by knowledge about these systems, which is valuable information to include into a model. Thus, it remains speculative whether this idea would work in practice or if it would introduce too much unnecessary complexity by representing the self-modification as an additional process.

Another possibility involves the general design of the model and the selection pressures in the resulting systems. Ideally, the mechanisms are designed in a way that the boundary values are mere safeguards, but they do not play a substantial role in limiting evolution in trait space. Assumptions about these safeguard can be sufficiently general, because the mechanisms are balanced in a way that makes it disadvantageous for species to have trait values close to these bounds. An example of such a model design was briefly touched on in the context of the alternative mutation mechanism (Section II.3.4.1), where the trend towards small niche widths was countered by a rising cost.

This requires sufficient knowledge about the behaviour of the system and sensible story lines for including the respective mechanisms. Yet, it may still not be applicable to all situations. In fact, a certain ambiguity in the mechanism is required for unexpected behaviour to occur; if all mechanisms would push the system towards an optimal point in trait space, there would be no point of having all these evolutionary mechanisms in the first place. This – reminiscent of the discussions made throughout EM – is the intrinsic challenge when modelling a system with the aim of novel behaviour to occur.

In summary, there appears to be no general solution for a 'canonical' inheritance mechanism; it seems that these mechanisms need to take the shape of a heuristic when modelling food webs. I tried to isolate some of the points that may inform a model design in such a context; however, they will most likely still require educated guesses of parameters and sufficient knowledge about the system in itself.

II.6.2.3 Environment representation and niche construction

The interactions in the ECST model as well as the extensions proposed above have in common that they are exclusively between species. This neglects a crucial aspect of every ecosystem: the environment it is embedded in.

In the ECST model, the environment is represented solely by the external resource, replenishing with logistic growth and acting as energy input to the primary producers of the food web; the resource itself is not evolving and has no dynamics other than being replenished. Expanding the notion of what the environment constitutes in this context can expand niche space in two conceptual ways: by adding a compartmentalisation or by providing additional interactions with the environment which influence the topology of niche space.

Meta-community models Compartmentalisation was discussed in Section I.4.2.1 as an important mechanism to provide further niches and parallel evolutionary development. The same approach is applied for food web models under the keyword of *meta-community models*.

One approach is to use the community matrix approach from [*May* 1972] and expand it to a meta-community matrix [*Gravel* et al. 2016], which consists of local, strongly-interacting community matrices that are coupled via dispersal terms. These systems show increased stability due to the homogenising effect of dispersal between the individual communities and stability appears to be higher if the individual compartments are more heterogeneous.

While the above is primarily a linear stability analysis on a meta-community matrix, niche models like the ECST model can also be expanded to represent a fragmented habitat: Basically, for each *patch* in the fragmented environment, a single food web is simulated; these food webs are then loosely coupled via diffusion terms, forming a (spatial) network of ecological networks [*Allhoff and Weiel* et al. 2015; *Bolchoun* et al. 2017]. The authors find that biodiversity in the patches is generally promoted and so-called *rescue effects* come into play, buffering extinctions in one patch by the dispersal of species between the patches.

It has been recognised that these environment heterogeneities are highly relevant in the reorganisation of ecological networks in changing environments, e.g. as caused by anthropogenic climate change [*Ryser* et al. 2021; *Thompson and Gonzalez* 2017; *Thompson and Guzman* et al. 2020]. In addition, dispersal processes appear to interact with other processes in ecological networks like local adaptations, sometimes counteracting the positive effect of dispersal on species persistence [*Thompson and Fronhofer* 2019].

Overall, these meta-community models expand the space of possible interactions by representing an abstracted form of spatiality. This is a qualitative departure from the basically well-mixed ECST model, where species extinction led to a complete loss of the associated information; in meta-community models, the species may still exist in another compartment. Even if the whole system is subject to an adverse influence, the larger number of ecological networks makes it more likely for a species to persist in one of the patches. As a result, the recovery mechanism would be governed by the dispersal of a species into the perturbed habitat.

Niche construction Another approach mentioned before is that of providing additional interactions with the environment itself, i.e.: not regarding the environment as the mere spatial embedding of the ecosystem but as an interaction partner which exhibits its own dynamics. As mentioned alongside the diversification process of EM (see Section I.3.3.1), one such interaction would be *niche construction* which denotes species modifying their habitat in some form [*Erwin* 2008]. This may be directly to their own advantage by simplifying access to resources or by providing shelter, or indirectly via the effects on other species.

Niche construction is a mechanism by which interactions become facilitated while others may be inhibited. As such, the changes in the environment encode certain characteristics of the ecological network – these characteristics are encoded not only in the interacting species, but also in the environment.

To illustrate this, consider a scenario as studied with the global perturbations in the ECST model. If the effect of the reduction in carrying capacity is species extinction, the regular ECST model loses all the information about previously existing interactions. However, in a model that allows persistent changes to the environment, some of the information about the previously existing interactions is encoded in the changes made to the environment. As long as the changes persist, any interaction that is similar to a previously facilitated interaction will have an evolutionary advantage; in effect, information about the previous interaction structure is propagated via the environment.

For an environment to retain structural information of the ecological network, it needs to be a more complex structure than a single replenishable resource as used in many niche models. The difficulty in modelling this is that the environment needs to have a representation of niches which can be modified by the resident species: an approximation of the naturally given form-function relationship of an environment. This can be achieved by introducing multiple independent resources that can be foraged on, each with their own modifiable properties (like a replenishment rate). In essence, the resource space becomes multi-dimensional, thus providing not only more niches to forage on but can also encode more information.

However, encoding niches only in the external resource neglects the fact that the environment shapes all interactions that occur in an ecosystem. Consequently, in a more general implementation of a niche construction mechanism, species should be able to modify not only trophic interactions with the external resource but all kinds of interactions throughout the ecosystem.

I explored such a mechanism in the context of this thesis. While more research is needed to draw conclusions from that extended model – which is the reason why it is not included here –, I want to briefly sketch the main ideas of the mechanism as an illustration of how a complex environment can be implemented.

As all interactions are represented in the functional response matrix, an environment modification mechanism can also be attached at that point, for instance as a modification matrix which encodes element-wise modulations of the functional response. The elements of the modification matrix basically represent the environment. Each element of that matrix can be affected by species diverting some of their foraged resources not into reproduction (their own biomass), but into the construction of a facilitative or inhibitory modification. This modification will then persist in the environment and only slowly decay towards zero if no more energy is diverted to maintaining it.

One important aspect is that the modifications need not be of interactions that they directly benefit from; for instance, a top predator may benefit more strongly from more resources being introduced into the system via the primary producers than from preying more efficiently on one particular prey. Subsequently, modifications act as higher-order interactions; they could even form catalytic loops if an excitatory modification causes a species to further divert resources to a modification.

How would modifications arise in such a scenario? The easiest assumption would be that of randomly appearing modifications which are then inherited to offspring and hence subject to selection pressures: If a modification yields a benefit despite the cost of diverting energy to it, it will be more likely that it persists – and vice versa for disadvantageous modifications.

In order to not lose this information in the event of a strong perturbation, the characteristics of an interaction need be encoded in more places than in the participating species alone. This is a crucial point: the interaction needs to arise from the *combination* of environment properties and the involved species, not exclusively from the species traits. Again, modelling such a system requires to find a sensible way to encode the relevant form-function relationships. In this case, the aim would be to find an encoding for the interaction between two species that would allow for two different, but similar species to have a similar encoding if they are in the same environment: A classification of a position in the combined trait space of the involved species to a certain group of modifications. How this should be derived from the combined feeding-related traits is mostly arbitrary and thus difficult to implement.

An alternative and more feasible approach would be to let the feeding interactions arise not from feeding kernels but from abstracted species properties – like done in the Webworld model [*Drossel* et al. 2001]. There, a number of random integers is associated with each species and encodes these abstracted properties; the feeding interactions are then derived from these integers and the modification groups could be derived analogously. By using abstracted traits, it becomes much easier to define a high-dimensional niche space, and simplifies extensions with other mechanisms. At the same time, this is a departure from allometrically-motivated feeding interactions, thus making it more difficult to associate concrete ecological properties with the species. Depending on the research question at hand, these costs and benefits of the various modelling approaches need to be weighed against each other.

II.6.2.4 Open questions

How can the additional complexity be handled? The model mechanisms suggested above – ranging from additional interactions to more complex environments – can be presumed to be suitable candidates for making food web models retain more structural information. It can thus be expected that such an extended model would allow studying recovery upon partial loss of structural information. But even in a more general context that aims at understanding dynamics of evolving ecosystems, the coupling of ecological processes with evolutionary processes is becoming increasingly important; for instance, as argued in [*Govaert* et al. 2021], incorporating both evolutionary and community-level process can 'inform broader questions about the maintenance of diversity and the resilience of diverse communities to disturbances'. Thus, it seems that there is certainly a need for more complex models.

However, it is clear that these extensions bring with them a large amount of assumptions and model complexity. This prompts the question of how the additional complexity can be handled – or if the expanded models would become *too* complicated to handle and understand. After all, each additional mechanism brings with it a number of new parameters, trait values, and observables; the curse of dimensionality looms over all efforts to expand a model.

This question is one of the most pressing ones, as it limits in how far models remain feasible to study these questions. One general approach is, of course, to better investigate isolated aspects of the systems and then attempt to bring them together. *Govaert* et al. [2021] follow a similar argumentation and see the understanding of relevant couplings between ecological and evolutionary processes as a crucial step into that direction. However, this is an approach that does not scale well, given that each new model component, interaction, or unknown parameter regime may cause completely unexpected behaviour.

Ultimately, this question is immensely difficult and no ultimate answers can be expected. Yet, a better understanding of underlying processes may assist in focussing models on the relevant processes.

In which scope should ecosystem resilience be studied? A second open question I want to address here is about the relevant scope for studying food web resilience.

So far, the focus in this discussion was on the behaviour of rather isolated ecosystems and how to potentially extend them to better describe recovery options, while anthropogenic effects were reduced to their direct interactions with ecosystems. Specifically this latter aspect might be overly reductive in the context of the Anthropocene (see Section I.2.8.2), where many of the scenarios that are presumed to be isolated are becoming increasingly interconnected with human societal and economic dynamics. In particular, ecosystem dynamics might be affected more strongly by the internal dynamics of human societies than by evolutionary mechanisms.

Hence, the open questions for studying ecosystem resilience are: In which scenarios is a reduction to eco-evolutionary dynamics still suitable? In which scenarios would it be imperative to choose a wider context, e.g. explicitly including socio-economic dynamics of human civilisation?

In light of the Anthropocene, these questions are certainly not new – the bigger picture is aptly summarised by *Donges and Winkelmann* et al.:

> The Anthropocene qualitatively differs from previous eras in Earths history in three key characteristics: (i) There is planetary-scale human agency. (ii) There are social and economic networks of teleconnections spanning the globe. (iii) It is dominated by planetary-scale social-ecological feedbacks. Bolting together old concepts and methodologies cannot be an adequate approach to describing this new geological era. Instead, we need a new paradigm in Earth system science that is founded equally on a deep understanding of the physical and biological Earth system – and of the economic, social and cultural forces that are now an intrinsic part of it. It is time to close the loop and bring socially

> mediated dynamics explicitly into theory, analysis and models that let us study the whole Earth system. *Donges and Winkelmann* et al. 2017

In this larger context, the study of eco-evolutionary dynamics still plays a very important role, as it helps to elucidate crucial aspects of the Earth system. However, it is important to recognise that ecosystems can no longer be seen as isolated system but are strongly coupled to other dynamics of the Earth system.

II.7

Conclusion & Outlook

In this part of my thesis, I presented studies on the evolution of abstracted ecosystems, their response to perturbations, their capacity for resilience against such perturbations, and their self-organising behaviour. To that end, I implemented and investigated the niche-based food web model proposed in [*Allhoff and Ritterskamp* et al. 2015] and analysed numerical simulations of that model and variations of it.

The main observations can be summarised as follows:

First, the study of local perturbations (in the form of species invasion, Section II.4.1) and global perturbations (in the form of reduced energy availability, Section II.4.2) elucidated the response of the ECST model in absence of the evolutionary mechanism. Specifically, one aim was to study whether there are correlations between measures of ecosystem functioning and the perturbation response. The observed responses to species invasion were strongly dependent on the existing structure in the food web and the trait combination of the invading species – and less so on the macroscopic measures. Yet, as studied later, the self-organisation of these food webs allowed to make some observations regarding the average response to a certain species invasion. When perturbing the food web by reducing the available energy, stronger perturbations removed basically all correlations to biodiversity and ecosystem measures prior to the perturbation; given the model mechanisms, this was not unexpected. The overall number of species in the food web was found to be a weak predictor for perturbation response.

Second, with the evolution mechanism present, I showed that the initially posed question regarding the resilience of a food web could not be studied with the chosen model, mainly because it does not retain sufficient information and hence lacks a recovery mechanism (Section II.4.3). As such, this particular model is too simplified to allow any sensible study of food web resilience. To address these shortcomings, I proposed a number of possible model extensions that would allow for recovery options to emerge, primarily by encoding structural information in the model that would either be unaffected by a certain kind of perturbation or be more long-lived than the species themselves.

Third, the heuristic mechanisms used to represent evolution in niche-based food web models was found to cause self-organisation in the form of reproducible patterns of species clusters in trait space (Section II.5.1). Depending on the available energy, the range of the structures in trait space varied. These inherent structures also explained patterns in previous observations. Given the properties of the model, I argued in Section II.5.2 that self-organised criticality could not be expected to manifest in the form of power-law distributions, mainly because the size and interaction structure of the system would not support such expectations. In addition, I suggested that the valuable insights are not in the classification of a system as SOC but in the questions studied in the process of attempting such a classification.

Throughout these studies, most of the observations could be explained more or less directly by choices in the model formulation. For instance, the choices made for the evolution mechanism or the extinction threshold turned out to be the underlying cause for a certain observation; this can be problematic if the choice is difficult to motivate from observations.

While I observed these challenges in the context of the chosen model and other models with which I attempted to study these questions, many of the difficulties are more general in that they pertain to finding appropriate model representations. Hence, they can be assumed to also play a role in other modelling-based studies of food web ecology or ecosystem evolution.

Despite the large body of work aiming to understand the processes that govern ecosystems, e.g. in the context of the complexity-stability debate, there appear to be few truly generalisable statements regarding the underlying mechanisms. Even when attempting to inform food web or biodiversity research by empirical data, many shortcomings remain in the form of lack of knowledge about mechanisms, trait-function relationships, or their interactions in general [*Hortal* et al. 2015]. While there are new approaches that aim to address these shortcomings [*Strydom* et al. 2021], these challenges appear to be difficult to surmount.

At this interface, Evolution Mechanics could provide a high-level description of the processes relevant in the evolution of ecosystems, the resulting interactions, and their behaviour. Arguably, many research questions may not benefit from more general theories, for instance because bridging conceptual scales becomes difficult. This may also be the case for the more quantitatively oriented modelling of anthropogenic effects or couplings to socio-economic systems. However, approaches that aim for a more conceptual understanding of abstracted systems may find benefits in the ideas of Evolution Mechanics.

Perhaps the research field of ecosystem evolution can be described as being at a fairly early point on the path towards a general understanding of these system – maybe akin to the field of biology prior to Darwin: There is a large number of observations and an increasing understanding of the processes relevant in specific scenarios, yet there is a lack of more widely applicable, fundamental concepts.

Given the complex interplay of processes and entities, the large differences in time and spatial scales, and the inherently evolutionary nature of ecosystems, this may not come as a surprise. In fact, it should be put into question if such a representation is attainable at all for these systems; it cannot be dismissed that this may not be the case. Either way, the path that lies ahead seems to hold many intriguing questions about the nature of ecosystems and their part in the Earth system at large.

Epilogue

In closing, I would like to briefly recall the initially formulated questions that set the wider context of this thesis: What are the fundamental processes that describe the Unfolding of Life? And which processes shape the emergence of ecosystems and how do they react to changes in their environment?

Surely, these questions are far from answered and will continue to inspire many interesting thoughts. I hope that this thesis was able to shed some light on these fascinating aspects of our world.

Bibliography

Ackermann, M. & *Doebeli*, M. (2004). Evolution of niche width and adaptive diversification. *Evolution*, *58*(12), 2599–2612. DOI: 10.1111/j.0014-3820.2004.tb01614.x

Allhoff, K. T. (2015). *Evolutionary food web models in fragmented landscapes* (Doctoral dissertation, Technische Universität Darmstadt).

Allhoff, K. T. & *Drossel*, B. (2016). Biodiversity and ecosystem functioning in evolving food webs. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *371*(1694), 20150281. DOI: 10.1098/rstb.2015.0281

Allhoff, K. T., *Ritterskamp*, D., *Rall*, B. C., *Drossel*, B. & *Guill*, C. (2015). Evolutionary food web model based on body masses gives realistic networks with permanent species turnover. *Scientific Reports*, *5*(1). DOI: 10.1038/srep10955

Allhoff, K. T., *Weiel*, E. M., *Rogge*, T. & *Drossel*, B. (2015). On the interplay of speciation and dispersal: An evolutionary food web model in space. *Journal of Theoretical Biology*, *366*, 46–56. DOI: 10.1016/j.jtbi.2014.11.006

Axelrod, R. M. & *Hamilton*, W. D. (1981). The evolution of cooperation. *Science*, *211*(4489), 1390–1396. DOI: 10.1126/science.7466396

Bak-Coleman, J. B., *Alfano*, M., *Barfuss*, W., *Bergstrom*, C. T., *Centeno*, M. A., *Couzin*, I. D., … *Weber*, E. U. (2021). Stewardship of global collective behavior. *Proceedings of the National Academy of Sciences*, *118*(27), e2025764118. DOI: 10.1073/pnas.2025764118

Bak, P. & *Sneppen*, K. (1993). Punctuated equilibrium and criticality in a simple model of evolution. *Physical Review Letters*, *71*(24), 4083–4086. DOI: 10.1103/physrevlett.71.4083

Bak, P., *Tang*, C. & *Wiesenfeld*, K. (1987). Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters*, *59*(4), 381–384. DOI: 10.1103/physrevlett.59.381

Bar-On, Y. M., *Phillips*, R. & *Milo*, R. (2018). The biomass distribution on earth. *Proceedings of the National Academy of Sciences*, *115*(25), 6506–6511. DOI: 10.1073/pnas.1711842115

Barfuss, W., *Donges*, J. F., *Wiedermann*, M. & *Lucht*, W. (2017). Sustainable use of renewable resources in a stylized social-ecological network model under heterogeneous resource distribution. *Earth System Dynamics*, *8*(2), 255–264. DOI: 10.5194/esd-8-255-2017

Bell, E. A., *Boehnke*, P., *Harrison*, T. M. & *Mao*, W. L. (2015). Potentially biogenic carbon preserved in a 4.1 billion-year-old zircon. *Proceedings of the National Academy of Sciences*, *112*(47), 14518–14521. DOI: 10.1073/pnas.1517557112

Ben-Kiki, O., *Evans*, C. & *Ingy döt Net.* (2009). YAML Ain't Markup Language (Version 1.2).

Betts, H. C., *Puttick*, M. N., *Clark*, J. W., *Williams*, T. A., *Donoghue*, P. C. J. & *Pisani*, D. (2018). Integrated genomic and fossil evidence illuminates life's early evolution and eukaryote origin. *Nature Ecology & Evolution*, *2*(10), 1556–1562. DOI: 10.1038/s41559-018-0644-x

Birch, J. (2017). The inclusive fitness controversy: Finding a way forward. *Royal Society Open Science*, *4*(7), 170335. DOI: 10.1098/rsos.170335

Birch, J. & *Okasha*, S. (2014). Kin selection and its critics. *BioScience*, *65*(1), 22–32. DOI: 10.1093/biosci/biu196

Bolchoun, L., *Drossel*, B. & *Allhoff*, K. T. (2017). Spatial topologies affect local food web structure and diversity in evolutionary metacommunities. *Scientific Reports*, *7*(1). DOI: 10.1038/s41598-017-01921-y

Boost Developers. (2021). Boost C++ Libraries (Version v1.75+). Retrieved from https://www.boost.org

Bourke, A. F. G. (2009). *Principles of social evolution.* Oxford University Press, USA.

Braitenberg, V. (1986). *Künstliche Wesen – Verhalten kybernetischer Vehikel.* Braunschweig / Wiesbaden: Friedr. Vieweg & Sohn.

Brännström, Å., *Johansson*, J., *Loeuille*, N., *Kristensen*, N., *Troost*, T. A., *Lambers*, R. H. R. & *Dieckmann*, U. (2012). Modelling the ecology and evolution of communities: A review of past achievements, current efforts, and future promises. *Evol. Ecol. Res. 14*, 601–625.

Brovkin, V., *Brook*, E., *Williams*, J. W., *Bathiany*, S., *Lenton*, T. M., *Barton*, M., … *Yu*, Z. (2021). Past abrupt changes, tipping points and cascading impacts in the earth system. *Nature Geoscience*, *14*(8), 550–558. DOI: 10.1038/s41561-021-00790-5

Brown, J. H., *Gillooly*, J. F., *Allen*, A. P., *Savage*, V. M. & *West*, G. B. (2004). Toward a metabolic theory of ecology. *Ecology*, *85*(7), 1771–1789. DOI: 10.1890/03-9000

Brueckner, J. & *Martin*, W. F. (2020). Bacterial genes outnumber archaeal genes in eukaryotic genomes. *Genome Biology and Evolution*, *12*(4), 282–292. DOI: 10.1093/gbe/evaa047

Buchanan, M. (2015). SOC revisited. *Nature Physics*, *11*(6), 442–442. DOI: 10.1038/nphys3354

Budd, G. E. (2003). The cambrian fossil record and the origin of the phyla. *Integrative and Comparative Biology*, *43*(1), 157–165. DOI: 10.1093/icb/43.1.157

Butterfield, N. J. (2015). Early evolution of the eukaryota. *Palaeontology*, *58*(1), 5–17. DOI: 10.1111/pala.12139

Cai, W., *Snyder*, J., *Hastings*, A. & *D'Souza*, R. M. (2020). Mutualistic networks emerging from adaptive niche-based interactions. *Nature Communications*, *11*(1). DOI: 10.1038/s41467-020-19154-5

Chialvo, D. R. (2010). Emergent complex neural dynamics. *Nature Physics*, *6*(10), 744–750. DOI: 10.1038/nphys1803

Collette, A. (2013). *Python and HDF5.* O'Reilly Media.

Dakos, V., *Matthews*, B., *Hendry*, A. P., *Levine*, J., *Loeuille*, N., *Norberg*, J., ... *Meester*, L. D. (2019). Ecosystem tipping points in an evolving world. *Nature Ecology & Evolution*, *3*(3), 355–362. DOI: 10.1038/s41559-019-0797-2

Darwin, C. (1859). *On the origin of species by means of natural selection, or the preservation of favoured races in the struggle for life* (1st ed.). John Murray, London.

Darwin, C. & *Wallace*, A. (1858). On the tendency of species to form varieties; and on the perpetuation of varieties and species by natural means of selection. *Journal of the Proceedings of the Linnean Society of London. Zoology*, *3*(9), 45–62. DOI: 10.1111/j.1096-3642.1858.tb02500.x

Dawkins, R. (1982). *The extended phenotype.* Oxford: Freeman.

Dobzhansky, T. (1964). Biology, Molecular and Organismic. *American Zoologist*, *4*(4), 443–452. DOI: 10.1093/icb/4.4.443

Dodd, M. S., *Papineau*, D., *Grenne*, T., *Slack*, J. F., *Rittner*, M., *Pirajno*, F., ... *Little*, C. T. S. (2017). Evidence for early life in earth's oldest hydrothermal vent precipitates. *Nature*, *543*(7643), 60–64. DOI: 10.1038/nature21377

Donges, J. F., *Lucht*, W., *Müller-Hansen*, F. & *Steffen*, W. (2017). The technosphere in earth system analysis: A coevolutionary perspective. *The Anthropocene Review*, *4*(1), 23–33. DOI: 10.1177/2053019616676608

Donges, J. F., *Winkelmann*, R., *Lucht*, W., *Cornell*, S. E., *Dyke*, J. G., *Rockström*, J., ... *Schellnhuber*, H. J. (2017). Closing the loop: Reconnecting human dynamics to earth system science. *The Anthropocene Review*, *4*(2), 151–157. DOI: 10.1177/2053019617725537

Doolittle, W. F. (2020). Evolution: Two domains of life or three? *Current Biology*, *30*(4), R177–R179. DOI: 10.1016/j.cub.2020.01.010

Drossel, B. (2001). Biological evolution and statistical physics. *Advances in Physics*, *50*(2), 209–295. DOI: 10.1080/00018730110041365

Drossel, B., *Higgs*, P. G. & *McKane*, A. J. (2001). The influence of predator–prey population dynamics on the long-term evolution of food web structure. *Journal of Theoretical Biology*, *208*(1), 91–107. DOI: 10.1006/jtbi.2000.2203

Ehrenfeld, J. G. (2010). Ecosystem consequences of biological invasions. *Annual Review of Ecology, Evolution, and Systematics*, *41*(1), 59–80. DOI: 10.1146/annurev-ecolsys-102209-144650

Eigen, M. (1971). Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften*, *58*(10), 465–523. DOI: 10.1007/BF00623322

Eklöf, A., *Jacob*, U., *Kopp*, J., *Bosch*, J., *Castro-Urgal*, R., *Chacoff*, N. P., ... *Allesina*, S. (2013). The dimensionality of ecological networks. *Ecology Letters*, *16*(5), 577–583. DOI: 10.1111/ele.12081

Elhacham, E., *Ben-Uri*, L., *Grozovski*, J., *Bar-On*, Y. M. & *Milo*, R. (2020). Global human-made mass exceeds all living biomass. *Nature*, *588*(7838), 442–444. DOI: 10.1038/s41586-020-3010-5

Erwin, D. (2008). Macroevolution of ecosystem engineering, niche construction and diversity. *Trends in Ecology & Evolution*, *23*(6), 304–310. DOI: 10.1016/j.tree.2008.01.013

Foster, K. R., *Schluter*, J., *Coyte*, K. Z. & *Rakoff-Nahoum*, S. (2017). The evolution of the host microbiome as an ecosystem on a leash. *Nature*, *548*(7665), 43–51. DOI: 10.1038/nature23292

Fritsch, C., *Billiard*, S. & *Champagnat*, N. (2021). Identifying conversion efficiency as a key mechanism underlying food webs evolution : A step forward, or backward ?

Frost, L. S., *Leplae*, R., *Summers*, A. O. & *Toussaint*, A. (2005). Mobile genetic elements: the agents of open source evolution. *Nature Reviews Microbiology*, *3*(1740-1526 (Print)), 722–732. DOI: 10.1038/nrmicro1235

Gause, G. F. (1932). Experimental studies on the struggle for existence. *Journal of Experimental Biology*, *9*, 389–402.

Gause, G. F. & *Witt*, A. A. (1935). Behavior of mixed populations and the problem of natural selection. *The American Naturalist*, *69*(725), 596–609. DOI: 10.1086/280628

Geier, F., *Barfuss*, W., *Wiedermann*, M., *Kurths*, J. & *Donges*, J. F. (2019). The physics of governance networks: critical transitions in contagion dynamics on multilayer adaptive networks with application to the sustainable use of renewable resources. *arXiv e-prints*, arXiv:1906.08679.

Girardot, B., *Gauduchon*, M., *Ménard*, F. & *Poggiale*, J. C. (2020). Does evolution design robust food webs? *Proceedings of the Royal Society B: Biological Sciences*, *287*(1930), 20200747. DOI: 10.1098/rspb.2020.0747

Goldenfeld, N., *Biancalani*, T. & *Jafarpour*, F. (2017). Universal biology and the statistical mechanics of early life. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *375*(2109), 20160341. DOI: 10.1098/rsta.2016.0341

Govaert, L., *Altermatt*, F., *Meester*, L. D., *Leibold*, M. A., *McPeek*, M. A., *Pantel*, J. H. & *Urban*, M. C. (2021). Integrating fundamental processes to understand eco-evolutionary community dynamics and patterns. *Functional Ecology*. DOI: 10.1111/1365-2435.13880

Gravel, D., *Massol*, F. & *Leibold*, M. A. (2016). Stability and complexity in model meta-ecosystems. *Nature Communications*, *7*, 12457. DOI: 10.1038/ncomms12457

Gribaldo, S. & *Brochier-Armanet*, C. (2019). Evolutionary relationships between archaea and eukaryotes. *Nature Ecology & Evolution*, *4*(1), 20–21. DOI: 10.1038/s41559-019-1073-1

Grilli, J., *Barabás*, G., *Michalska-Smith*, M. J. & *Allesina*, S. (2017). Higher-order interactions stabilize dynamics in competitive network models. *Nature*. DOI: 10.1038/nature23273

Grosberg, R. K. & *Strathmann*, R. R. (2007). The evolution of multicellularity: A minor major transition? *Annual Review of Ecology, Evolution, and Systematics*, *38*(1), 621–654. DOI: 10.1146/annurev.ecolsys.36.102403.114735

Guill, C. & *Drossel*, B. (2008). Emergence of complexity in evolving niche-model food webs. *Journal of Theoretical Biology*, *251*(1), 108–120. DOI: 10.1016/j.jtbi.2007.11.017

Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology*, *7*(1), 1–16. DOI: 10.1016/0022-5193(64)90038-4

Hardin, G. (1960). The competitive exclusion principle. *Science*, *131*(3409), 1292–1297. DOI: 10.1126/science.131.3409.1292

Harfoot, M. B. J., *Newbold*, T., *Tittensor*, D. P., *Emmott*, S., *Hutton*, J., *Lyutsarev*, V., … *Purves*, D. W. (2014). Emergent global patterns of ecosystem structure and function from a mechanistic general ecosystem model. *PLoS Biology*, *12*(4), e1001841. DOI: 10.1371/journal.pbio.1001841

Harris, C. R., *Millman*, K. J., van der *Walt*, S. J., *Gommers*, R., *Virtanen*, P., *Cournapeau*, D., … *Oliphant*, T. E. (2020). Array programming with NumPy. *Nature*, *585*(7825), 357–362. DOI: 10.1038/s41586-020-2649-2

Heitzig, J., *Barfuss*, W. & *Donges*, J. (2018). A thought experiment on sustainable management of the earth system. *Sustainability*, *10*(6), 1947. DOI: 10.3390/su10061947

Herdeanu, B. (2021). *Emergence of Cooperation in Evolutionary Social Interaction Networks* (Doctoral dissertation, Heidelberg University).

Higgs, P. G. & *Lehman*, N. (2014). The RNA world: Molecular cooperation at the origins of life. *Nature Reviews Genetics*, *16*(1), 7–17. DOI: 10.1038/nrg3841

Holland, H. D. (2006). The oxygenation of the atmosphere and oceans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *361*(1470), 903–915. DOI: 10.1098/rstb.2006.1838

Hölldobler, B. & *Wilson*, E. O. (1990). *The Ants.* Springer.

Hordijk, W. (2013). Autocatalytic sets: From the origin of life to the economy. *BioScience*, *63*(11), 877–881. DOI: 10.1525/bio.2013.63.11.6

Hortal, J., de *Bello*, F., *Diniz-Filho*, J. A. F., *Lewinsohn*, T. M., *Lobo*, J. M. & *Ladle*, R. J. (2015). Seven shortfalls that beset large-scale knowledge of biodiversity. *Annual Review of Ecology, Evolution, and Systematics*, *46*(1), 523–549. DOI: 10.1146/annurev-ecolsys-112414-054400

Hoyer, S. & *Hamman*, J. (2017). Xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software*, *5*(1). DOI: 10.5334/jors.148

Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in Science Engineering*, *9*(3), 90–95. DOI: 10.1109/MCSE.2007.55

Huxley, J. (1942). *Evolution: The modern synthesis.* George Allen and Unwin.

Isson, T. T., *Love*, G. D., *Dupont*, C. L., *Reinhard*, C. T., *Zumberge*, A. J., *Asael*, D., … *Planavsky*, N. J. (2018). Tracking the rise of eukaryotes to ecological dominance with zinc isotopes. *Geobiology*, *16*(4), 341–352. DOI: 10.1111/gbi.12289

Ito, H. C., *Shimada*, M. & *Ikegami*, T. (2009). Coevolutionary dynamics of adaptive radiation for food-web development. *Population Ecology*, *51*(1), 65–81. DOI: 10.1007/s10144-008-0113-5

Johnson, S., *Domnguez-Garca*, V., *Donetti*, L. & *Muñoz*, M. A. (2014). Trophic coherence determines food-web stability. *Proceedings of the National Academy of Sciences*, *111*(50), 17923–17928. DOI: 10.1073/pnas.1409077111

Judson, O. P. (2017). The energy expansions of evolution. *Nature Ecology & Evolution*, *1*, 0138. DOI: 10.1038/s41559-017-0138

Kadanoff, L. P., *Nagel*, S. R., *Wu*, L. & *Zhou*, S.-m. (1989). Scaling and universality in avalanches. *Physical Review A*, *39*(12), 6524–6537. DOI: 10.1103/physreva.39.6524

Kauffman, S. A. (1986). Autocatalytic sets of proteins. *Journal of Theoretical Biology*, *119*(1), 1–24. DOI: 10.1016/S0022-5193(86)80047-9

Keeling, P. J. & *Palmer*, J. D. (2008). Horizontal gene transfer in eukaryotic evolution. *Nature Reviews Genetics*, *9*(8), 605–618. DOI: 10.1038/nrg2386

Kleidon, A. (2010). Life, hierarchy, and the thermodynamic machinery of planet earth. *Physics of Life Reviews*, *7*(4), 424–460. DOI: 10.1016/j.plrev.2010.10.002

Knoll, A. H. & *Nowak*, M. A. (2017). The timetable of evolution. *Science Advances*, *3*(5), e1603076. DOI: 10.1126/sciadv.1603076

Koonin, E. V. & *Novozhilov*, A. S. (2017). Origin and evolution of the universal genetic code. *Annual Review of Genetics*, *51*(1), 45–62. DOI: 10.1146/annurev-genet-120116-024713

Koonin, E. V. & *Wolf*, Y. I. (2012). Evolution of microbes and viruses: A paradigm shift in evolutionary biology? *Frontiers in Cellular and Infection Microbiology*, *2*. DOI: 10.3389/fcimb.2012.00119

Krakauer, D. C., *Bertschinger*, N., *Olbrich*, E., *Flack*, J. C. & *Ay*, N. (2020). The information theory of individuality. *Theory in Biosciences*, *139*(2), 209–223. DOI: 10.1007/s12064-020-00313-7

Lade, S. J., *Bodin*, Ö., *Donges*, J. F., *Enfors Kautsky*, E., *Galafassi*, D., *Olsson*, P. & *Schlüter*, M. (2017). Modelling social-ecological transformations: an adaptive network proposal. *arXiv e-prints*, arXiv:1704.06135.

Laland, K. N., *Uller*, T., *Feldman*, M. W., *Sterelny*, K., *Müller*, G. B., *Moczek*, A., … *Odling-Smee*, J. (2015). The extended evolutionary synthesis: Its structure, assumptions and predictions. *Proceedings of the Royal Society B: Biological Sciences*, *282*(1813), 20151019. DOI: 10.1098/rspb.2015.1019

Laland, K. N., *Uller*, T., *Feldman*, M., *Sterelny*, K., *Müller*, G. B., *Moczek*, A., … *Strassmann*, J. E. (2014). Does evolutionary theory need a rethink? *Nature*, *514*(7521), 161–164. DOI: 10.1038/514161a

Lammer, H., *Bredehöft*, J. H., *Coustenis*, A., *Khodachenko*, M. L., *Kaltenegger*, L., *Grasset*, O., ... *Rauer*, H. (2009). What makes a planet habitable? *The Astronomy and Astrophysics Review*, *17*(2), 181–249. DOI: 10.1007/s00159-009-0019-z

Landi, P., *Minoarivelo*, H. O., *Brännström*, Å., *Hui*, C. & *Dieckmann*, U. (2018). Complexity and stability of ecological networks: A review of the theory. *Population Ecology*, *60*(4), 319–345. DOI: 10.1007/s10144-018-0628-3

Lane, N. (2011). Energetics and genetics across the prokaryote-eukaryote divide. *Biology Direct*, *6*(1), 35. DOI: 10.1186/1745-6150-6-35

Lane, N. & *Martin*, W. F. (2010). The energetics of genome complexity. *Nature*, *467*(7318), 929–934. DOI: 10.1038/nature09486

Lenton, T. M. & *Watson*, A. (2011). *Revolutions that made the earth.* DOI: 10.1093/acprof:oso/9780199587049.001.0001

Loeuille, N. & *Loreau*, M. (2005). Evolutionary emergence of size-structured food webs. *Proceedings of the National Academy of Sciences*, *102*(16), 5761–5766. DOI: 10.1073/pnas.0408424102

Lotka, A. J. (1920). Analytical note on certain rhythmic relations in organic systems. *Proceedings of the National Academy of Sciences*, *6*(7), 410–415.

Lotka, A. J. (1925). *Elements of physical biology.* Williams and Wilkins Company, Baltimore.

Lyons, T. W., *Reinhard*, C. T. & *Planavsky*, N. J. (2014). The rise of oxygen in earth's early ocean and atmosphere. *Nature*, *506*(7488), 307–315. DOI: 10.1038/nature13068

MacArthur, R. H. (1972). *Geographical ecology: Patterns in the distribution of species.* Biology / Princeton University Press. Princeton University Press.

MacKay, R. S., *Johnson*, S. & *Sansom*, B. (2020). How directed is a directed network? *Royal Society Open Science*, *7*(9), 201138. DOI: 10.1098/rsos.201138

Malamud, B. D. (1998). Forest fires: An example of self-organized critical behavior. *Science*, *281*(5384), 1840–1842. DOI: 10.1126/science.281.5384.1840

Malthus, T. R. (1798). *An essay on the principle of population.* J. Johnson, London.

Martin, W. F. (2017). Symbiogenesis, gradualism, and mitochondrial energy in eukaryote origin. *Periodicum Biologorum*, *119*(3), 141–158. DOI: 10.18054/pb.v119i3.5694

Martin, W. F., *Baross*, J., *Kelley*, D. & *Russell*, M. J. (2008). Hydrothermal vents and the origin of life. *Nature Reviews Microbiology*, *6*(11), 805–814. DOI: 10.1038/nrmicro1991

Martin, W. F., *Garg*, S. & *Zimorski*, V. (2015). Endosymbiotic theories for eukaryote origin. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *370*(1678), 20140330. DOI: 10.1098/rstb.2014.0330

Martin, W. F. & *Russell*, M. J. (2003). On the origins of cells: A hypothesis for the evolutionary transitions from abiotic geochemistry to chemoautotrophic prokaryotes, and from prokaryotes to nucleated cells. *Philosophical Transactions of the*

Royal Society of London, Series B: Biological Sciences, *358*(1429), 59–85. DOI: 10.1098/rstb.2002.1183

May, R. M. (1972). Will a large complex system be stable? *Nature*, *238*(5364), 413–414. DOI: 10.1038/238413a0

Mereschkowsky, C. (1910). Theorie der zwei plasmaarten als grundlage der symbiogenesis, einer neuen lehre von der entstehung der organismen. *Biologisches Centralblatt*, *30*(9–11), 278–367.

Moran, N. A. & *Sloan*, D. B. (2015). The hologenome concept: Helpful or hollow? *PLOS Biology*, *13*(12), e1002311. DOI: 10.1371/journal.pbio.1002311

Müller, P. M., *Heitzig*, J., *Kurths*, J., *Lüdge*, K. & *Wiedermann*, M. (2021). Anticipation-induced social tipping: Can the environment be stabilised by social dynamics? *The European Physical Journal Special Topics.* DOI: 10.1140/epjs/s11734-021-00011-5

Nghe, P., *Hordijk*, W., *Kauffman*, S. A., *Walker*, S. I., *Schmidt*, F. J., *Kemble*, H., ... *Lehman*, N. (2015). Prebiotic network evolution: Six key parameters. *Molecular BioSystems*, *11*, 3206–3217. DOI: 10.1039/C5MB00593K

Nowak, M. A., *McAvoy*, A., *Allen*, B. & *Wilson*, E. O. (2017). The general form of hamilton's rule makes no predictions and cannot be tested empirically. *Proceedings of the National Academy of Sciences.* DOI: 10.1073/pnas.1701805114

Nowak, M. A., *Tarnita*, C. E. & *Wilson*, E. O. (2010). The evolution of eusociality. *Nature*, *466*(7310), 1057–1062. DOI: 10.1038/nature09205

Olejarz, J., *Iwasa*, Y., *Knoll*, A. H. & *Nowak*, M. A. (2021). The great oxygenation event as a consequence of ecological dynamics modulated by planetary change. *Nature Communications*, *12*(1). DOI: 10.1038/s41467-021-23286-7

Olff, H., *Alonso*, D., *Berg*, M. P., *Eriksson*, B. K., *Loreau*, M., *Piersma*, T. & *Rooney*, N. (2009). Parallel ecological networks in ecosystems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1524), 1755–1779. DOI: 10.1098/rstb.2008.0222

Paczuski, M., *Maslov*, S. & *Bak*, P. (1996). Avalanche dynamics in evolution, growth, and depinning models. *Physical Review E*, *53*, 414–443. DOI: 10.1103/PhysRevE.53.414

Payne, J. L., *Bachan*, A., *Heim*, N. A., *Hull*, P. M. & *Knope*, M. L. (2020). The evolution of complex life and the stabilization of the earth system. *Interface Focus*, *10*(4), 20190106. DOI: 10.1098/rsfs.2019.0106

Pennisi, E. (2018). The power of many. *Science*, *360*(6396), 1388–1391. DOI: 10.1126/science.360.6396.1388

Petchey, O. L., *O'Gorman*, E. J. & *Flynn*, D. F. B. (2009). A functional guide to functional diversity measures. In *Biodiversity, ecosystem functioning, and human wellbeing* (pp. 49–59). DOI: 10.1093/acprof:oso/9780199547951.003.0004

Pigliucci, M. (2007). Do we need an Extended Evolutionary Synthesis? *Evolution*, *61*(12), 2743–2749. DOI: 10.1111/j.1558-5646.2007.00246.x

Pilosof, S., *Porter*, M. A., *Pascual*, M. & *Kéfi*, S. (2017). The multilayer nature of ecological networks. *Nature Ecology & Evolution*, *1*, 0101 EP -.

Pimm, S. L. & *Lawton*, J. H. (1978). On feeding on more than one trophic level. *Nature*, *275*(5680), 542–544. DOI: 10.1038/275542a0

Podani, J. (2017). Different from trees, more than metaphors: Branching silhouettescorals, cacti, and the oaks. *Systematic Biology*, *66*(5), 737–753. DOI: 10.1093/sysbio/syx039

Poole, A. M. (2009). Horizontal gene transfer and the earliest stages of the evolution of life. *Research in Microbiology*, *160*(7), 473–480. DOI: 10.1016/j.resmic.2009.07.009

Purves, D., *Scharlemann*, J. P. W., *Harfoot*, M., *Newbold*, T., *Tittensor*, D. P., *Hutton*, J. & *Emmott*, S. (2013). *Time to model all life on earth.* Nature Publishing Group.

Queller, D. C. (1997). Cooperators since life began. *The Quarterly Review of Biology*, *72*(2), 184–188. DOI: 10.1086/419766

Queller, D. C. (2000). Relatedness and the fraternal major transitions. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *355*(1403), 1647–1655. DOI: 10.1098/rstb.2000.0727

Ratcliff, W. C., *Denison*, R. F., *Borrello*, M. & *Travisano*, M. (2012). Experimental evolution of multicellularity. *Proceedings of the National Academy of Sciences*, *109*(5), 1595–1600. DOI: 10.1073/pnas.1115323109

Raup, D. M. (1986). Biological extinction in earth history. *Science*, *231*(4745), 1528–1533. DOI: 10.1126/science.11542058

Riedel, L., *Herdeanu*, B., *Mack*, H., *Sevinchan*, Y. & *Weninger*, J. (2020). Utopia: A Comprehensive and Collaborative Modeling Framework for Complex and Evolving Systems. *Journal of Open Source Software*, *5*(53), 2165. DOI: 10.21105/joss.02165

Ripple, W. J., *Wolf*, C., *Newsome*, T. M., *Barnard*, P. & *Moomaw*, W. R. (2019). World scientists' warning of a climate emergency. *BioScience.* DOI: 10.1093/biosci/biz088

Rohr, R. P., *Saavedra*, S. & *Bascompte*, J. (2014). On the structural stability of mutualistic systems. *Science*, *345*(6195), 1253497. DOI: 10.1126/science.1253497

Romanuk, T. N., *Binzer*, A., *Loeuille*, N., *Carscallen*, W. M. A. & *Martinez*, N. D. (2019). Simulated evolution assembles more realistic food webs with more functionally similar species than invasion. *Scientific Reports*, *9*(1). DOI: 10.1038/s41598-019-54443-0

Romanuk, T. N., *Zhou*, Y., *Valdovinos*, F. & *Martinez*, N. (2017). Robustness trade-offs in model food webs. In *Networks of invasion: A synthesis of concepts* (pp. 263–291). DOI: 10.1016/bs.aecr.2016.11.001

Roth, K. (2020). *Chaotic, complex, and evolving environmental systems.* Lecture Notes, v0.6.

Russell, M. J. & *Hall*, A. J. (1997). The emergence of life from iron monosulphide bubbles at a submarine hydrothermal redox and pH front. *Journal of the Geological Society*, *154*(3), 377–402. DOI: 10.1144/gsjgs.154.3.0377

Ryser, R., *Hirt*, M. R., *Häussler*, J., *Gravel*, D. & *Brose*, U. (2021). Landscape heterogeneity buffers biodiversity of simulated meta-food-webs under global change through rescue and drainage effects. *Nature Communications*, *12*(1). DOI: 10.1038/s41467-021-24877-0

Sagan, L. (1967). On the origin of mitosing cells. *Journal of Theoretical Biology*, *14*(3), 225–IN6. DOI: 10.1016/0022-5193(67)90079-3

Sanderson, C. & *Curtin*, R. (2016). Armadillo: A template-based C++ library for linear algebra. *The Journal of Open Source Software*, *1*(2), 26. DOI: 10.21105/joss.00026

Schlebusch, C. M., *Malmström*, H., *Günther*, T., *Sjödin*, P., *Coutinho*, A., *Edlund*, H., … *Jakobsson*, M. (2017). Southern african ancient genomes estimate modern human divergence to 350,000 to 260,000 years ago. *Science*, *358*(6363), 652–655. DOI: 10.1126/science.aao6266

Schopf, J. W. (2006). Fossil evidence of archaean life. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *361*(1470), 869–885. DOI: 10.1098/rstb.2006.1834

Sevinchan, Y. (2020). Paramspace (Version v2.4.1). DOI: 10.5281/zenodo.3826471

Sevinchan, Y., *Herdeanu*, B., *Mack*, H., *Riedel*, L. & *Roth*, K. (2020). Boosting group-level synergies by using a shared modeling framework. In V. V. *Krzhizhanovskaya*, G. *Závodszky*, M. H. *Lees*, J. J. *Dongarra*, P. M. A. *Sloot*, S. *Brissos* & J. *Teixeira* (Eds.), *Computational Science – ICCS 2020* (pp. 442–456). DOI: 10.1007/978-3-030-50436-6_32

Sevinchan, Y., *Herdeanu*, B. & *Traub*, J. (2020). dantro: A Python package for handling, transforming, and visualizing hierarchically structured data. *Journal of Open Source Software*, *5*(52), 2316. DOI: 10.21105/joss.02316

Simon, H. A. (1962). The Architecture of Complexity. *Proceedings of the American Philosophical Society*, *106*(6), 467–482.

Sneppen, K., *Bak*, P., *Flyvbjerg*, H. & *Jensen*, M. H. (1995). Evolution as a self-organized critical phenomenon. *Proceedings of the National Academy of Sciences*, *92*(11), 5209–5213. DOI: 10.1073/pnas.92.11.5209

Solé, R. V. (2016). The major synthetic evolutionary transitions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *371*(1701), 20160175. DOI: 10.1098/rstb.2016.0175

Solé, R. V., *Valverde*, S., *Casals*, M. R., *Kauffman*, S. A., *Farmer*, D. & *Eldredge*, N. (2013). The evolutionary ecology of technological innovations. *Complexity*, *18*(4), 15–27. DOI: 10.1002/cplx.21436

Staps, M., van *Gestel*, J. & *Tarnita*, C. E. (2019). Emergence of diverse life cycles and life histories at the origin of multicellularity. *Nature Ecology & Evolution*, *3*(8), 1197–1205. DOI: 10.1038/s41559-019-0940-0

Steffen, W., *Richardson*, K., *Rockstrom*, J., *Cornell*, S. E., *Fetzer*, I., *Bennett*, E. M., … *Sorlin*, S. (2015). Planetary boundaries: Guiding human development on a changing planet. *Science*, *347*(6223), 1259855–1259855. DOI: 10.1126/science.1259855

Steffen, W., *Rockström*, J., *Richardson*, K., *Lenton*, T. M., *Folke*, C., *Liverman*, D., … *Schellnhuber*, H. J. (2018). Trajectories of the earth system in the anthropocene. *Proceedings of the National Academy of Sciences*, *115*(33), 8252–8259. DOI: 10.1073/pnas.1810141115

Steiner, F. M., *Schlick-Steiner*, B. C., *Moder*, K., *Stauffer*, C., *Arthofer*, W., *Buschinger*, A., … *Crozier*, R. H. (2007). Abandoning aggression but maintaining self-nonself discrimination as a first stage in ant supercolony formation. *Current Biology*, *17*(21), 1903–1907. DOI: 10.1016/j.cub.2007.09.061

Strydom, T., *Catchen*, M. D., *Banville*, F., *Caron*, D., *Dansereau*, G., *Desjardins-Proulx*, P., … *Poisot*, T. (2021). A roadmap toward predicting species interaction networks (across space and time). DOI: 10.32942/osf.io/eu7k3

Szathmáry, E. (2015). Toward major evolutionary transitions theory 2.0. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(33), 10104–10111. DOI: 10.1073/pnas.1421398112

Szathmáry, E. & *Maynard Smith*, J. (1995). The major evolutionary transitions. *Nature*, *374*(6519), 227–232. DOI: 10.1038/374227a0

Takeuchi, N. & *Hogeweg*, P. (2012). Evolutionary dynamics of RNA-like replicator systems: A bioinformatic approach to the origin of life. *Physics of Life Reviews*, *9*(3), 219–263. DOI: 10.1016/j.plrev.2012.06.001

Tamberg, L. A., *Heitzig*, J. & *Donges*, J. F. (2020). A modeler's guide to studying the resilience of social-technical-environmental systems.

Taylor, T. (2014). Evolution in virtual worlds. In M. *Grimshaw* (Ed.), (Preprint). Oxford University Press.

Taylor, T. (2015). Requirements for open-ended evolution in natural and artificial systems. *CoRR*, *abs/1507.07403.*

Taylor, T. & *Dorin*, A. (2020). *Rise of the self-replicators.* DOI: 10.1007/978-3-030-48234-3

Theis, K. R., *Dheilly*, N. M., *Klassen*, J. L., *Brucker*, R. M., *Baines*, J. F., *Bosch*, T. C. G., … *Bordenstein*, S. R. (2016). Getting the hologenome concept right: An eco-evolutionary framework for hosts and their microbiomes. *mSystems*, *1*(2). DOI: 10.1128/msystems.00028-16

Thompson, P. L. & *Fronhofer*, E. A. (2019). The conflict between adaptation and dispersal for maintaining biodiversity in changing environments. *Proceedings of the National Academy of Sciences*, *116*(42), 21061–21067. DOI: 10.1073/pnas.1911796116

Thompson, P. L. & *Gonzalez*, A. (2017). Dispersal governs the reorganization of ecological networks under environmental change. *Nature Ecology & Evolution*, *1*(6), 0162. DOI: 10.1038/s41559-017-0162

Thompson, P. L., *Guzman*, L. M., *Meester*, L. D., *Horváth*, Z., *Ptacnik*, R., *Vanschoenwinkel*, B., … *Chase*, J. M. (2020). A process-based metacommunity framework linking local and regional scale community ecology. *Ecology Letters.* DOI: 10.1111/ele.13568

Touboul, J. & *Destexhe*, A. (2010). Can power-law scaling and neuronal avalanches arise from stochastic dynamics? *PLoS ONE*, *5*(2), e8982. DOI: 10.1371/journal.pone.0008982

Trifonov, E. N. (2011). Vocabulary of definitions of life suggests a definition. *Journal of Biomolecular Structure and Dynamics*, *29*(2), 259–266. DOI: 10.1080/073911011010524992

van der Walt, S., *Colbert*, S. C. & *Varoquaux*, G. (2011). The numpy array: A structure for efficient numerical computation. *Computing in Science Engineering*, *13*(2), 22–30. DOI: 10.1109/MCSE.2011.37

Vandenkoornhuyse, P., *Quaiser*, A., *Duhamel*, M., *Van*, A. L. & *Dufresne*, A. (2015). The importance of the microbiome of the plant holobiont. *New Phytologist*, *206*(4), 1196–1206. DOI: 10.1111/nph.13312

Vasas, V., *Fernando*, C., *Santos*, M., *Kauffman*, S. A. & *Szathmáry*, E. (2012). Evolution before genes. *Biology Direct*, *7*(1), 1–14. DOI: 10.1186/1745-6150-7-1

Veldhuis, M. P., *Berg*, M. P., *Loreau*, M. & *Olff*, H. (2018). Ecological autocatalysis: A central principle in ecosystem organization? *Ecological Monographs*, *88*(3), 304–319. DOI: 10.1002/ecm.1292

Verhulst, P. F. (1838). *Correspondance mathématique et physique* (J. G. *Garnier* & A. *Quetelet*, Eds.). Impr. d'H. Vandekerckhove.

Vetsigian, K., *Woese*, C. & *Goldenfeld*, N. (2006). Collective evolution and the genetic code. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(28), 10696–701. DOI: 10.1073/pnas.0603780103

Volterra, V. (1926). Fluctuations in the abundance of a species considered mathematically. *Nature*, *118*(2972), 558–560. DOI: 10.1038/118558a0

von *Neumann*, J. (1966). *Theory of self-reproducing automata* (A. W. *Burks*, Ed.). University of Illinois Press.

Waters, C. N., *Zalasiewicz*, J., *Summerhayes*, C., *Barnosky*, A. D., *Poirier*, C., *Gauszka*, A., ... *Wolfe*, A. P. (2016). The anthropocene is functionally and stratigraphically distinct from the holocene. *Science*, *351*(6269). DOI: 10.1126/science.aad2622

Watkins, N. W., *Pruessner*, G., *Chapman*, S. C., *Crosby*, N. B. & *Jensen*, H. J. (2015). 25 years of self-organized criticality: Concepts and controversies. *Space Science Reviews*, *198*(1-4), 3–44. DOI: 10.1007/s11214-015-0155-x

Watson, R. A., *Mills*, R., *Buckley*, C. L., *Kouvaris*, K., *Jackson*, A., *Powers*, S. T., ... *Power*, D. (2015). Evolutionary connectionism: Algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. *Evolutionary Biology*, 1–29. DOI: 10.1007/s11692-015-9358-z

Watson, R. A. & *Szathmáry*, E. (2016). How can evolution learn? *Trends in Ecology & Evolution*, *31*(2), 147–157. DOI: 10.1016/j.tree.2015.11.009

Weiss, M. C., *Preiner*, M., *Xavier*, J. C., *Zimorski*, V. & *Martin*, W. F. (2018). The last universal common ancestor between ancient earth chemistry and the onset of genetics. *PLOS Genetics*, *14*(8), e1007518. DOI: 10.1371/journal.pgen.1007518

Weiss, M. C., *Sousa*, F. L., *Mrnjavac*, N., *Neukirchen*, S., *Roettger*, M., *Nelson-Sathi*, S. & *Martin*, W. F. (2016). The physiology and habitat of the last universal common ancestor. *Nature Microbiology*, *1*(9). DOI: 10.1038/nmicrobiol.2016.116

West, S. A., *Fisher*, R. M., *Gardner*, A. & *Kiers*, E. T. (2015). Major evolutionary transitions in individuality. *Proceedings of the National Academy of Sciences*, *112*(33), 10112–10119. DOI: 10.1073/pnas.1421402112

White, J. W., *Rassweiler*, A., *Samhouri*, J. F., *Stier*, A. C. & *White*, C. (2013). Ecologists should not use statistical significance tests to interpret simulation model results. *Oikos*, *123*(4), 385–388. DOI: 10.1111/j.1600-0706.2013.01073.x

Williams, R. & *Martinez*, N. (2004). Limits to trophic levels and omnivory in complex food webs: Theory and data. *The American Naturalist*, *163*(3), 458–468. DOI: 10.1086/381964

Williams, T. A., *Cox*, C. J., *Foster*, P. G., *Szöllsi*, G. J. & *Embley*, T. M. (2019). Phylogenomics provides robust support for a two-domains tree of life. *Nature Ecology & Evolution*, *4*(1), 138–147. DOI: 10.1038/s41559-019-1040-x

Wilson, D. S. & *Wilson*, E. O. (2007). Rethinking the theoretical foundation of sociobiology. *Quarterly Review of Biology*, *82*(4), 327–348.

Wilson, E. O. (2012). *The social conquest of earth*. Liveright.

Wilson, E. O. & *Hölldobler*, B. (2005). The rise of the ants: A phylogenetic and ecological explanation. *Proceedings of the National Academy of Sciences*, *102*(21), 7411–7414. DOI: 10.1073/pnas.0502264102

Woese, C. R. (2002). On the evolution of cells. *Proceedings of the National Academy of Sciences of the United States of America*, *99*(13), 8742–8747. DOI: 10.1073/pnas.132266999

Woese, C. R. & *Fox*, G. E. (1977). Phylogenetic structure of the prokaryotic domain: The primary kingdoms. *Proceedings of the National Academy of Sciences*, *74*(11), 5088–5090. DOI: 10.1073/pnas.74.11.5088

Woese, C. R., *Kandler*, O. & *Wheelis*, M. L. (1990). Towards a natural system of organisms: Proposal for the domains archaea, bacteria, and eucarya. *Proceedings of the National Academy of Sciences*, *87*(12), 4576–4579. DOI: 10.1073/pnas.87.12.4576

Wollrab, S., *Diehl*, S. & *Roos*, A. M. D. (2012). Simple rules describe bottom-up and top-down control in food webs with alternative energy pathways. *Ecology Letters*, *15*(9), 935–946. DOI: 10.1111/j.1461-0248.2012.01823.x

Xavier, J. C., *Hordijk*, W., *Kauffman*, S., *Steel*, M. & *Martin*, W. F. (2020). Autocatalytic chemical networks at the origin of metabolism. *Proceedings of the Royal Society B: Biological Sciences*, *287*(1922), 20192377. DOI: 10.1098/rspb.2019.2377

Yodzis, P. & *Innes*, S. (1992). Body size and consumer-resource dynamics. *The American Naturalist*, *139*(6), 1151–1175. DOI: 10.1086/285380

Zalasiewicz, J., *Williams*, M., *Waters*, C. N., *Barnosky*, A. D., *Palmesino*, J., *Rönnskog*, A.-S., … *Wolfe*, A. P. (2017). Scale and diversity of the physical technosphere:

A geological perspective. *The Anthropocene Review*, *4*(1), 9–22. doi: 10.1177/2053019616677743

Zhang, L., *Andersen*, K. H., *Dieckmann*, U. & *Brännström*, Å. (2015). Four types of interference competition and their impacts on the ecology and evolution of size-structured populations and communities. *Journal of Theoretical Biology*, *380*, 280–290. doi: 10.1016/j.jtbi.2015.05.023

Publications

I was first author or co-author on the following peer-reviewed publications during the time of my doctoral studies:

Riedel, L., *Herdeanu*, B., *Mack*, H., *Sevinchan*, Y. & *Weninger*, J. (2020). Utopia: A Comprehensive and Collaborative Modeling Framework for Complex and Evolving Systems. *Journal of Open Source Software*, *5*(53), 2165. DOI: 10.21105/joss.02165

Sevinchan, Y., *Herdeanu*, B., *Mack*, H., *Riedel*, L. & *Roth*, K. (2020). Boosting group-level synergies by using a shared modeling framework. In V. V. *Krzhizhanovskaya*, G. *Závodszky*, M. H. *Lees*, J. J. *Dongarra*, P. M. A. *Sloot*, S. *Brissos* & J. *Teixeira* (Eds.), *Computational Science – ICCS 2020* (pp. 442–456). DOI: 10.1007/978-3-030-50436-6_32

Sevinchan, Y., *Herdeanu*, B. & *Traub*, J. (2020). dantro: A Python package for handling, transforming, and visualizing hierarchically structured data. *Journal of Open Source Software*, *5*(52), 2316. DOI: 10.21105/joss.02316

Appendix

A Supplementary Material on Evolution Mechanics

This part of the appendix contains supplementary material for Part I of this thesis.

A.1 Auto-catalytic sets

An *auto-catalytic set* (ACS) is defined as a set of species and their interactions which collectively catalyse the production of each of the set's components. While initially developed in the context of chemical reaction networks [*Eigen* 1971; *Kauffman* 1986], the concept is applicable to a much wider range of phenomena [*Hordijk* 2013], basically any system that can be abstracted to a reaction network.

Furthermore, ACS have been shown [*Nghe* et al. 2015; *Vasas* et al. 2012] to be capable of self-reproduction and hence behaving as evolutionary units – as discussed in Section I.2.2.3. This early form of evolution in chemical reaction networks may have played an important role in the *Origin of Life* and is the reason to include a brief description of the main ideas here.

Definitions

Species & the food set In this context, a *species* is any population of effectively identical entities. If applied to chemical reactions, each type of molecule would be considered a species.

The *food set* is a small set of species that are available in large abundance, such that they can be considered 'freely available' [*Hordijk* 2013] and far from becoming depleted.

Reaction In the simplest case, a *reaction* is the process of turning a species A into another species B, producing some waste energy W:

$$\mathrm{A} \underset{k_-}{\overset{k_+}{\rightleftharpoons}} \mathrm{B} + \mathrm{W} \ , \tag{1}$$

where the $\rightleftharpoons$ denote the reaction being an equilibrium reaction, with k_+ and k_- being the respective forward and backward reaction rates. The rate constants are determined by the kinetics of the reaction, typically including an energetic barrier that needs to be overcome in order for the reaction to take place.

Catalysts & inhibitors A *catalyst* is a species C, that changes the kinetics of a reaction such that the energetic barrier is reduced and the effective reaction rates are higher[63]. Analogously, an *inhibitor* I leads to reduced reaction rates.

The reduction in activation energy is achieved by the catalyst forming an intermediate species with the reactant and then proceeding to react to the product:

$$\mathrm{A} + \mathrm{C} \rightleftharpoons \mathrm{AC} \rightleftharpoons \mathrm{B} + \mathrm{C} + \mathrm{W}\ . \tag{2}$$

With C being part of both the product and reactant side of the reaction, it does not get consumed in this process. Ignoring C thus yields the same reaction as in 1 but with higher net reaction rates.

Auto-catalysis An *auto-catalytic reaction* is a reaction where the product catalyses its own synthesis, meaning: at least one of the products is also a reactant. The simplest example is the following reaction (basically Equation (2) with A and B referring to the same kind of species):

$$\mathrm{A} + \mathrm{C} \rightleftharpoons 2\mathrm{C} + \mathrm{W}\ . \tag{3}$$

Auto-catalytic sets A set of species and their reactions is called an *auto-catalytic set* (ACS) if they *collectively* catalyse each other's synthesis: (i) every reaction in the set is catalysed by one or more other species from the set, and (ii) every species in the set can be produced using only the reactions from the set and input from the food set. These conditions capture catalytic closure of the set and its self-sustainability, respectively.

Hordijk [2013] further developed the above approach under the term *reflexively auto-catalytic and food-generated sets* (RAF sets, RAFs), contributing a mathematically rigorous formulation.

Auto-catalytic loop Auto-catalytic loops refer to a set of species that form a closed, circular path of catalytic relationships: each species depends on the previous species in the cycle to catalyse its own synthesis.

---

[63] Note that both the forward *and* the backward reaction rate are equally influenced by the presence of a catalyst. Subsequently, the equilibrium constant k_+/k_- is the same for catalysed and uncatalysed reactions.

B Supplementary Material on Ecosystem Evolution

This part of the appendix contains supplementary material for Part II of this thesis.

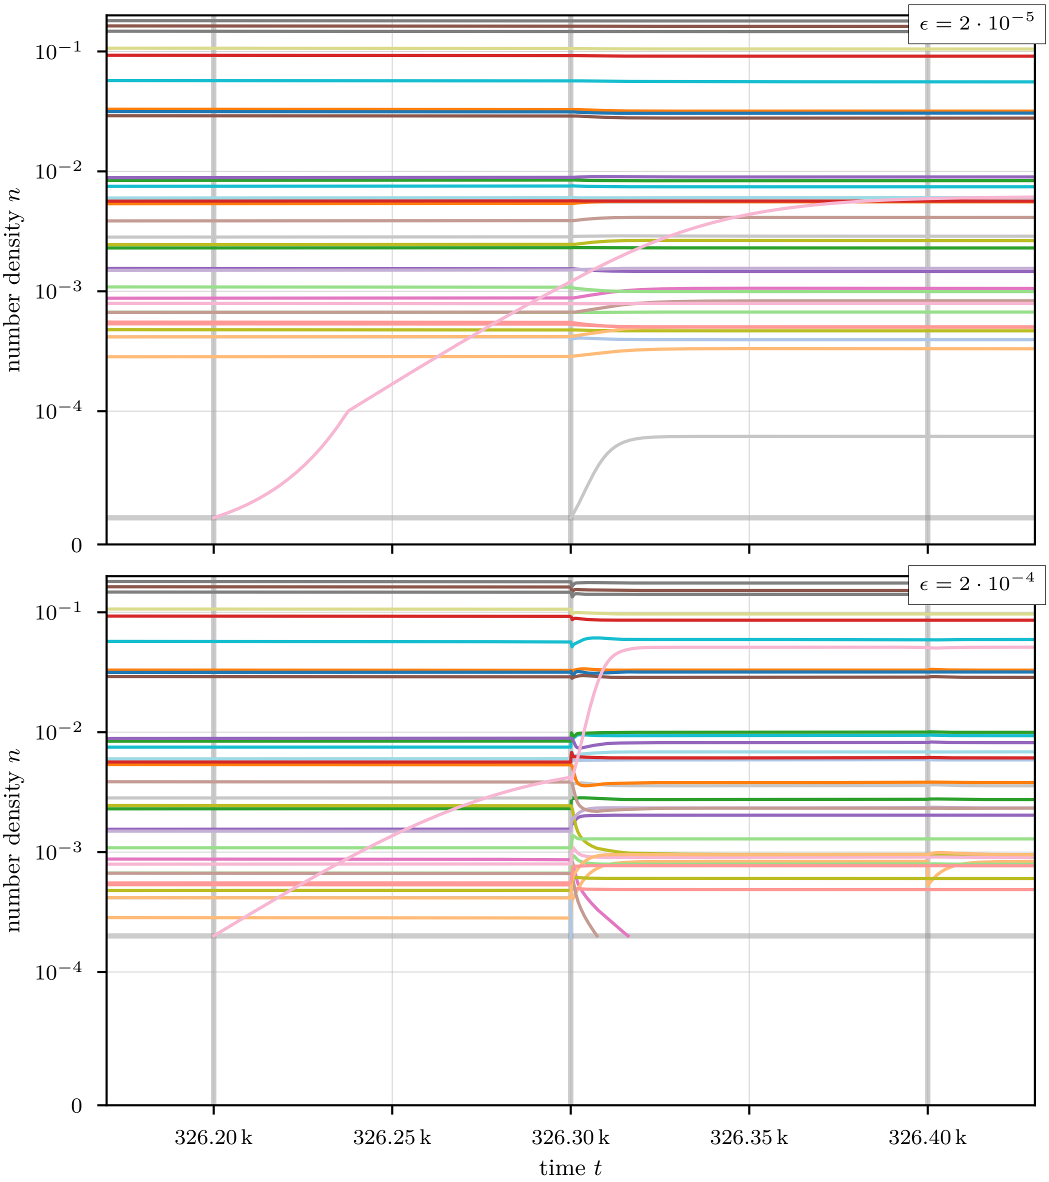

Figure B.1 Number density development for a small section of two long simulations (not shown), based on Figure II.3.11. The two simulations are identical up to $t = 326.2\,\mathrm{k}$, from which point on they differ in the value of their extinction threshold ε.
Vertical lines indicate the times at which new mutants are added; depending on ε, these new mutants are able to survive or not. Furthermore, the biomass transfer from parent to offspring is higher in the high ε simulation, leading to a larger perturbation of the system and subsequent extinctions that do not occur in the simulation with the low extinction threshold.

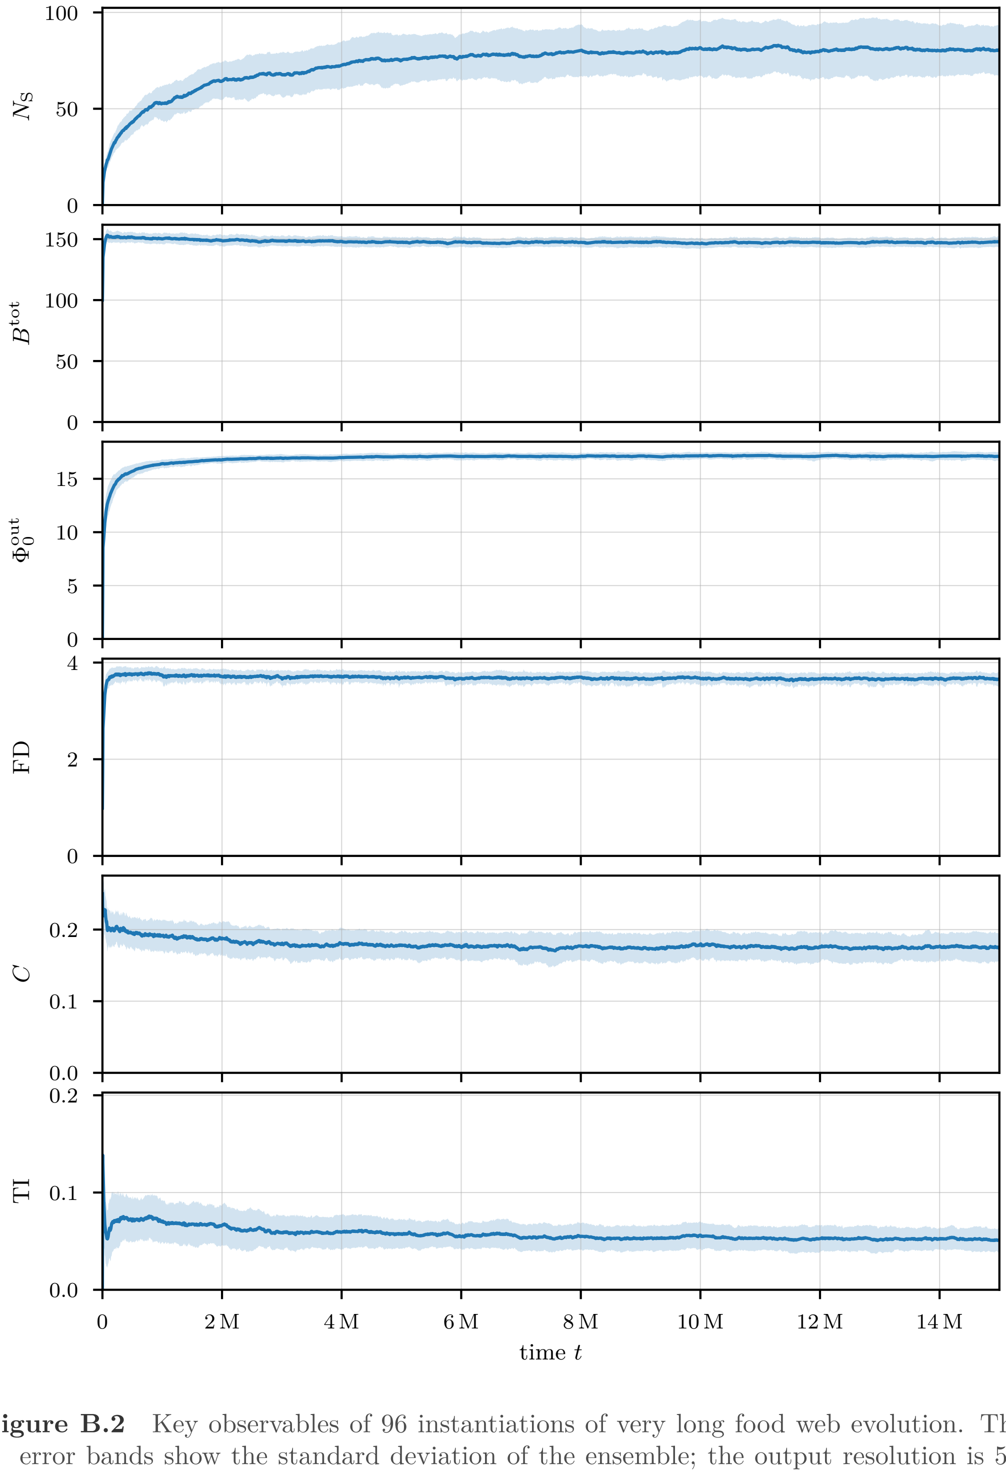

Figure B.2 Key observables of 96 instantiations of very long food web evolution. The error bands show the standard deviation of the ensemble; the output resolution is 5 k steps, equivalent to 50 mutation events.

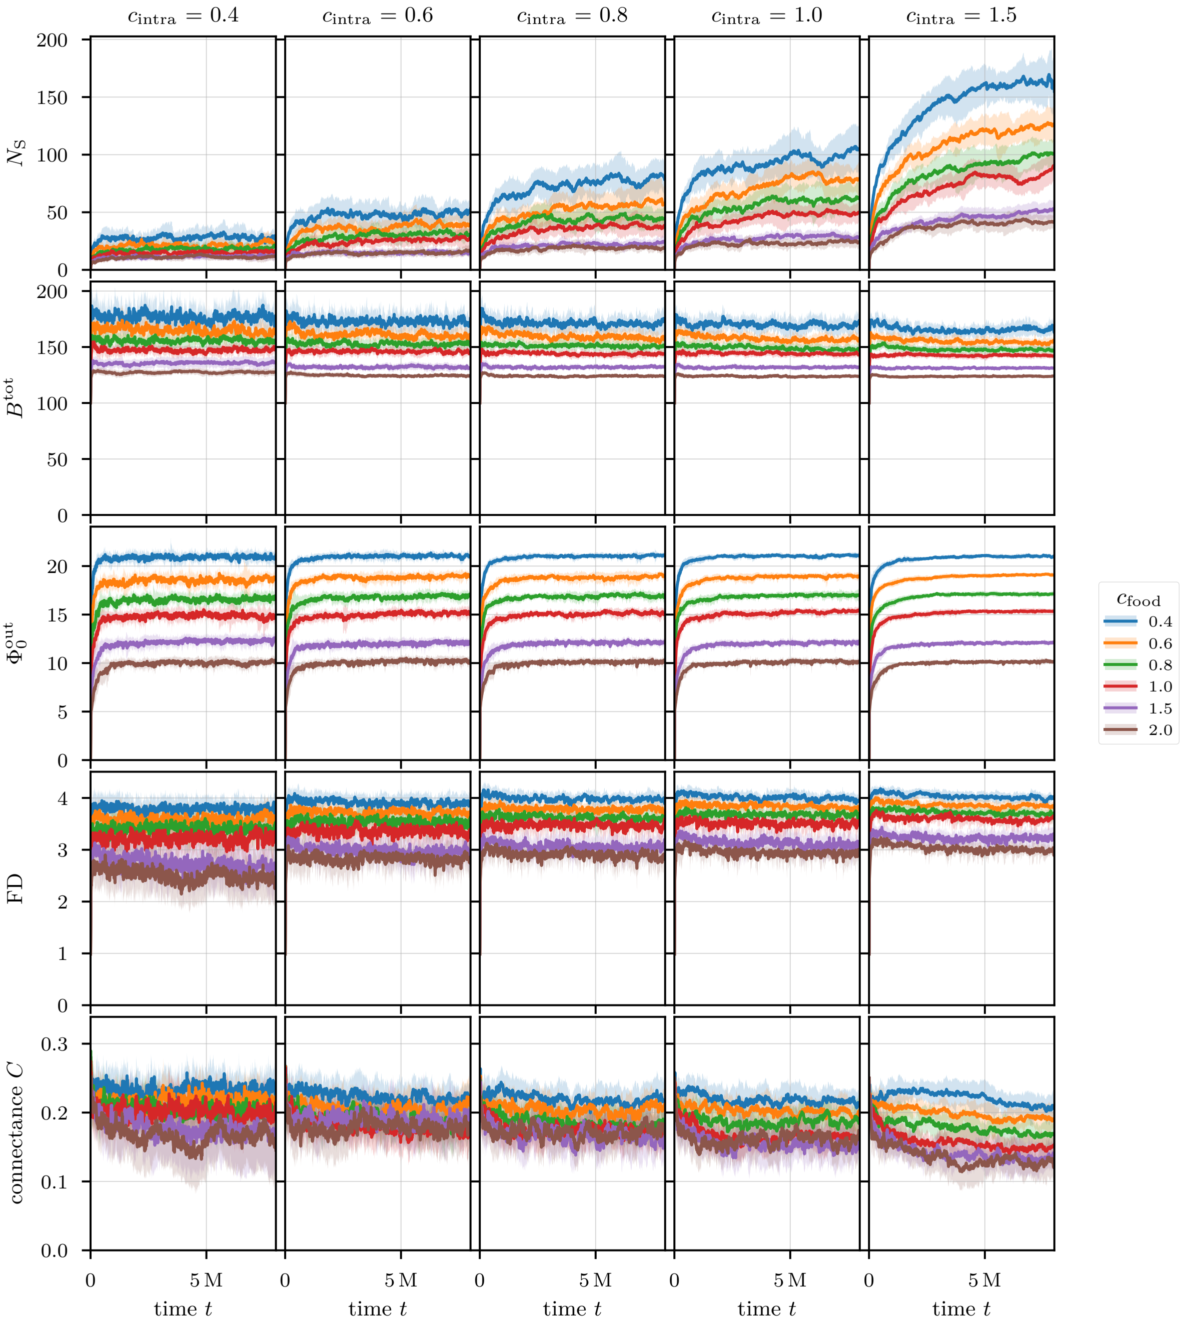

Figure B.3 Key observables for a sweep over competition parameters c_{food} and c_{intra} with 16 instantiations per parameter combination. The bands show
The competition parameters have a strong effect on almost all metrics of the system. Higher food competition c_{food} in particular leads to all these measures being reduced, mainly due to the higher denominator of the functional response g_{ij}. Intra-specific competition influences the number of species most: due to the further reduced functional response, species number densities remain small (not shown), allowing a wider variety of species to coexist, hence driving up the number of species. Notably, these species are not unique, as can be seen in FD not being affected by a higher c_{intra} (apart from the reduced variance, which)

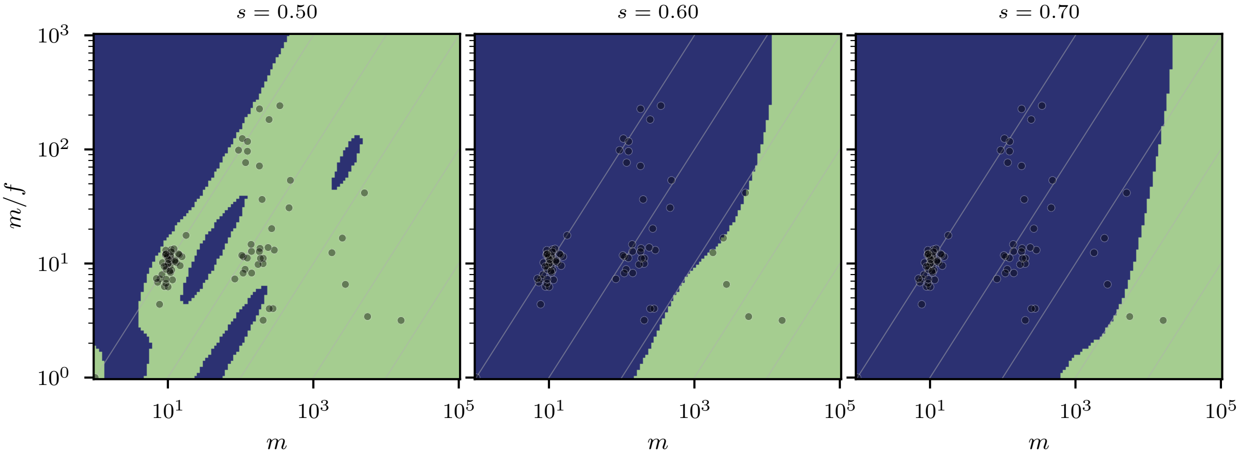

Figure B.4 Whether the invader survived (green) for a wider range of niche width values, corresponding to Figure II.4.1a.

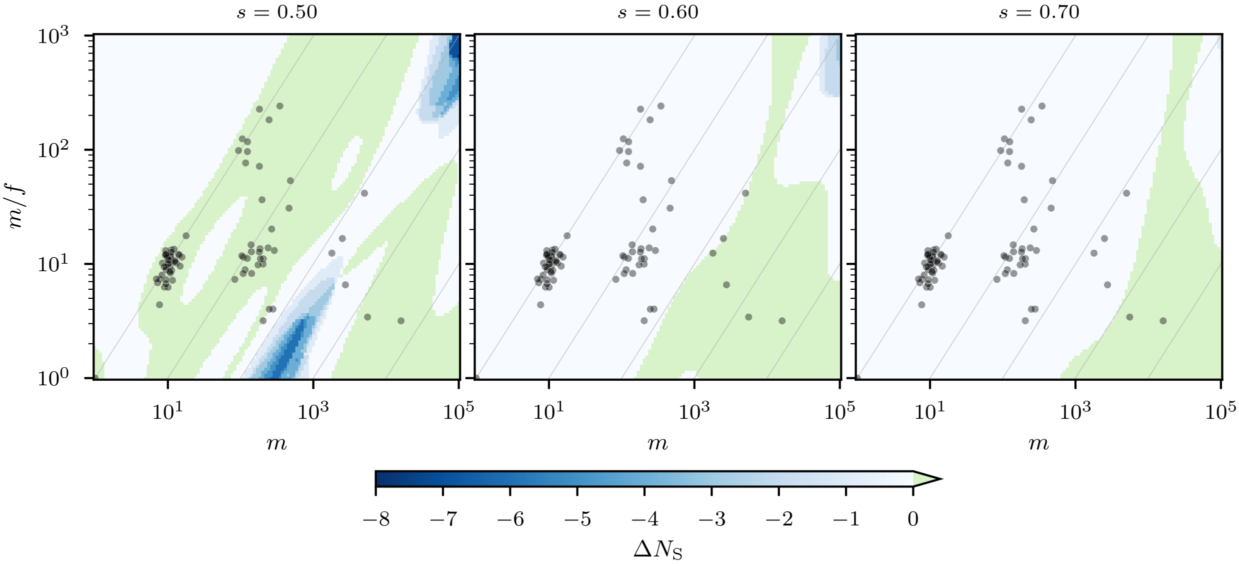

Figure B.5 Absolute change in population size induced by the invader species. This extends Figure II.4.1b with additional niche width values.

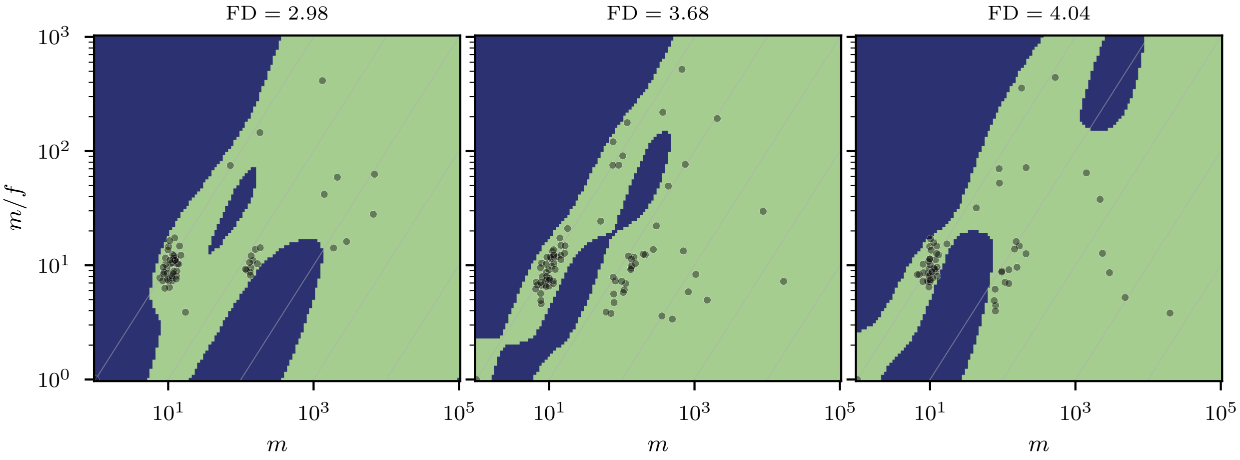

Figure B.6 Trait combinations in which the invader species survived (green), corresponding to Figure II.4.3.

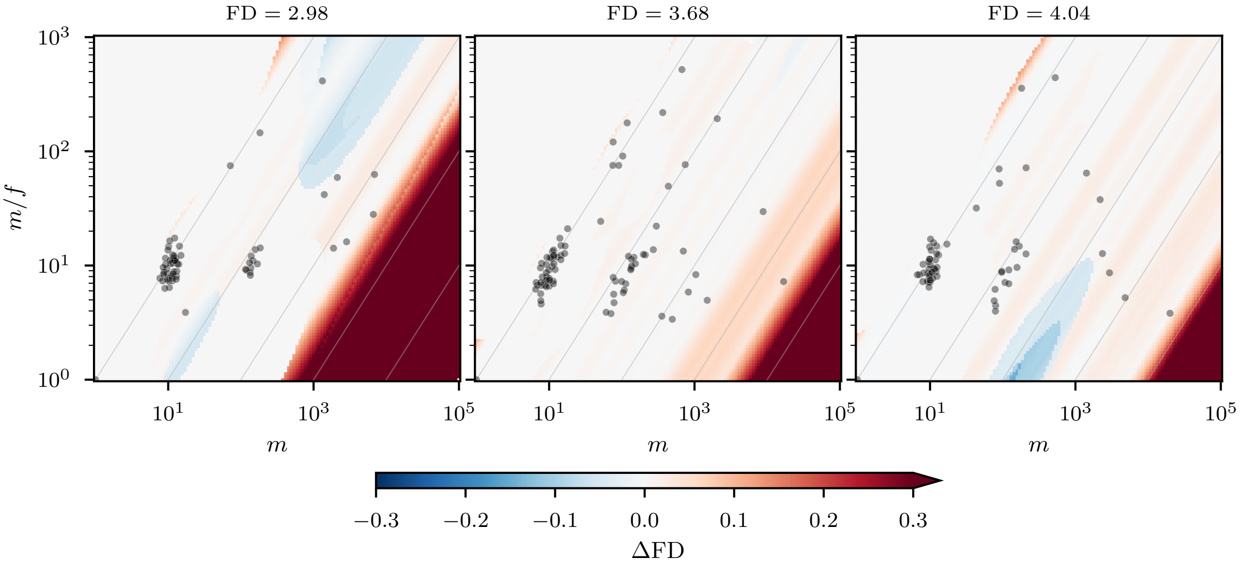

Figure B.7 Absolute change in functional diversity FD induced by the addition of new species into evolved food webs with different values of functional diversity FD, corresponding to Figure II.4.3.

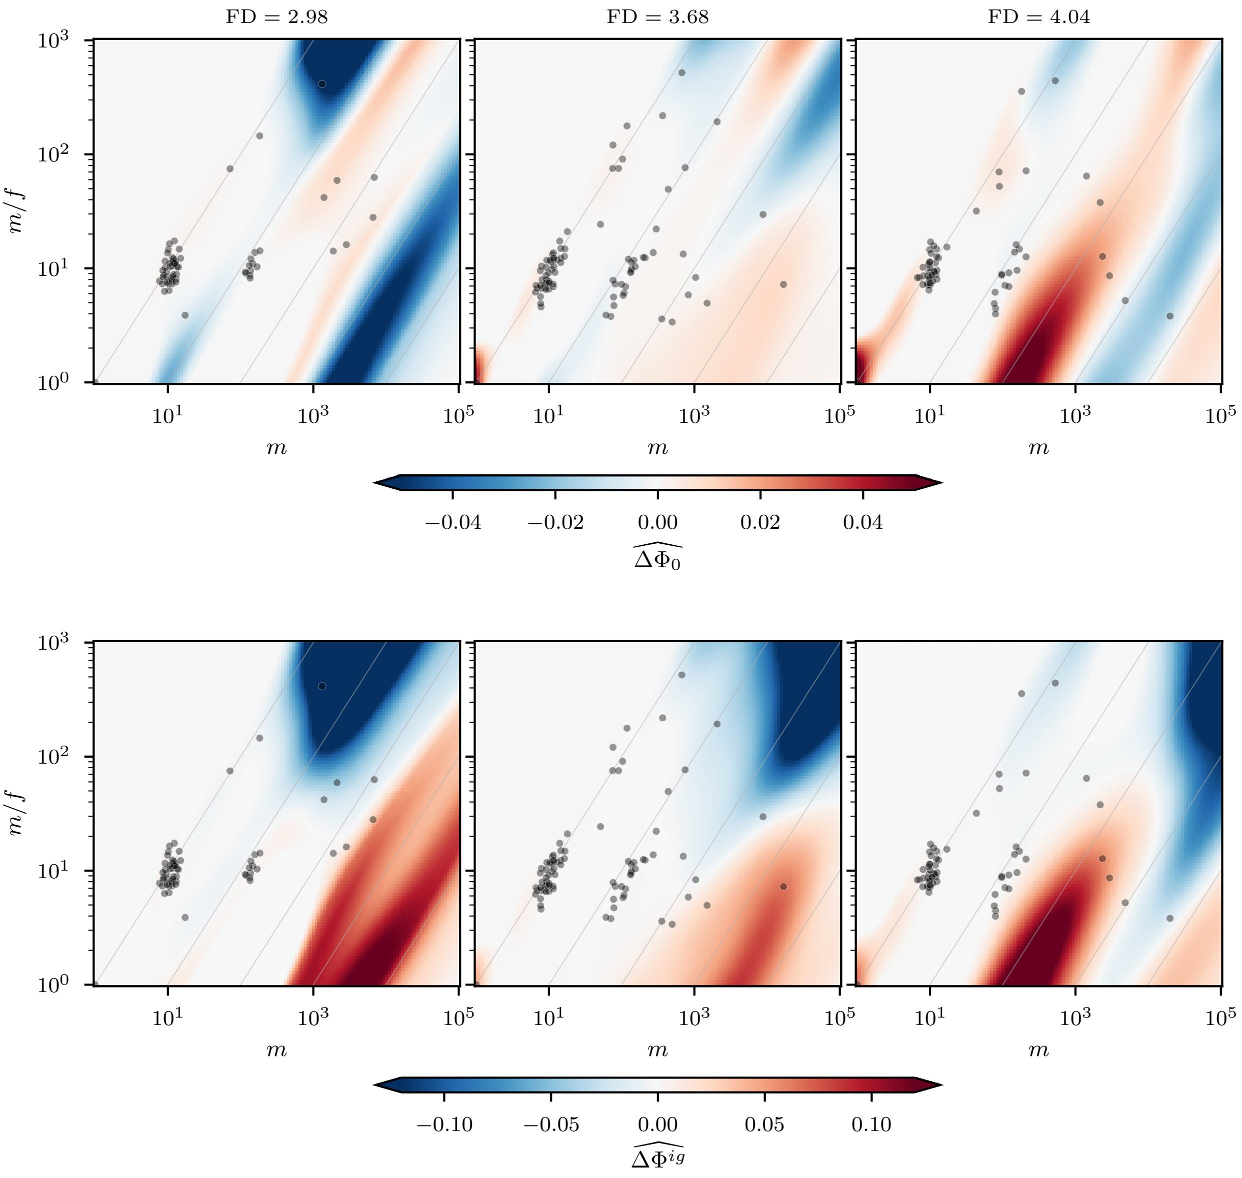

Figure B.8 Relative change in biomass flow from resource $\widehat{\Delta\Phi_0^{\mathrm{out}}}$ **(top)**, and intra-guild flow $\widehat{\Delta\Phi^{\mathrm{ig}}}$ **(bottom)** for the invasion experiments shown in Figure II.4.6. The colour range is capped to make low-valued structures more visible; scales are not the same between the subplots.

Note the alternating regions of positive and negative change in $\widehat{\Delta\Phi_0^{\mathrm{out}}}$ and how this is affected by the depth of the trophic network (increasing for higher FD) and species extinctions.

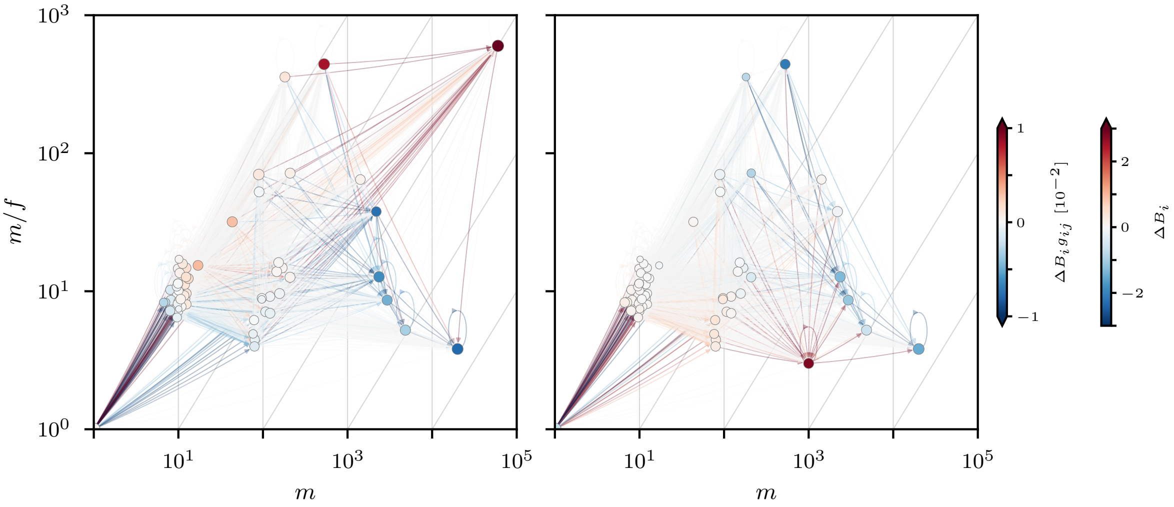

Figure B.9 Response to species invasion visualised using a fully connected network representation. The invader species are those at $(6 \cdot 10^4, 600)$ **(left)** and $(1000, 3)$ **(right)**. Unlike in the thresholded network representation in Figure II.4.7, these plots show changes to *all* feeding interactions and hence offers a more complete picture – at the cost of reduced clarity. From the thresholded networks alone, the changes in biomass flow (cf. Figure B.8) cannot be explained, because the thresholding algorithm can lead to weaker links not being included into the network representation; these weaker links are precisely those that change a lot and lead to the overall reduction (left scenario) or increase (right) in intra-guild biomass flow.

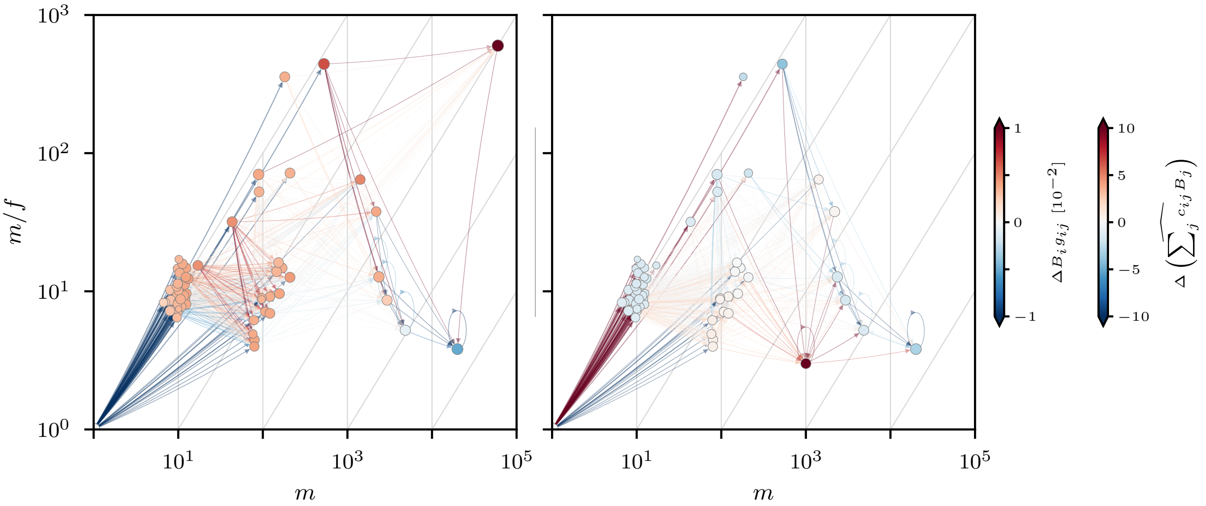

Figure B.10 As supplementary information to Figure II.4.7, this plot shows the relative change in perceived competition for each prey species. This value appears in the denominator of the functional response, Equation (II.2.4) and is the same for each link to a predator – it quantifies the competition between all its predator species *for* a particular prey species.

Studied center-of-mass dependencies

For detecting changes in trait space position of invasion effects depending on food web properties, center-of-mass (COM) analyses as done for Figure II.4.11 were carried out for the following observables:

- invader survival
- coexistence
- largest cascade
- cascades ≥ 2 (‡)
- cascades ≥ 5 (‡)
- ΔN_{S} (†)
- $\Delta B_{\mathrm{S}} > 0$ (‡)
- $\Delta B_{\mathrm{S}} < 0$ (‡)
- $\widehat{\Delta B}_{\mathrm{S}} > +0.05$ (‡)
- $\widehat{\Delta B}_{\mathrm{S}} < -0.05$ (‡)
- $\widehat{\Delta B}_{\mathrm{S}} > +0.1$ (‡)
- $\widehat{\Delta B}_{\mathrm{S}} < -0.1$ (‡)
- $\widehat{\Delta B}_{\mathrm{S}}$ above 95 % quantile
- $\widehat{\Delta B}_{\mathrm{S}}$ below 5 % quantile
- $\left|\widehat{\Delta B}_{\mathrm{S}}\right|$ (†)
- $\left|\widehat{\Delta B^{\mathrm{tot}}}\right|$ (†)
- $\left|\widehat{\Delta \Phi_0^{\mathrm{out}}}\right|$ (†)

Here, † denotes that a weighted COM computation was used and ‡ denotes that *both* weighted and unweighted COM computations were studied. As stated in the main text, no dependencies on functional diversity were detected.

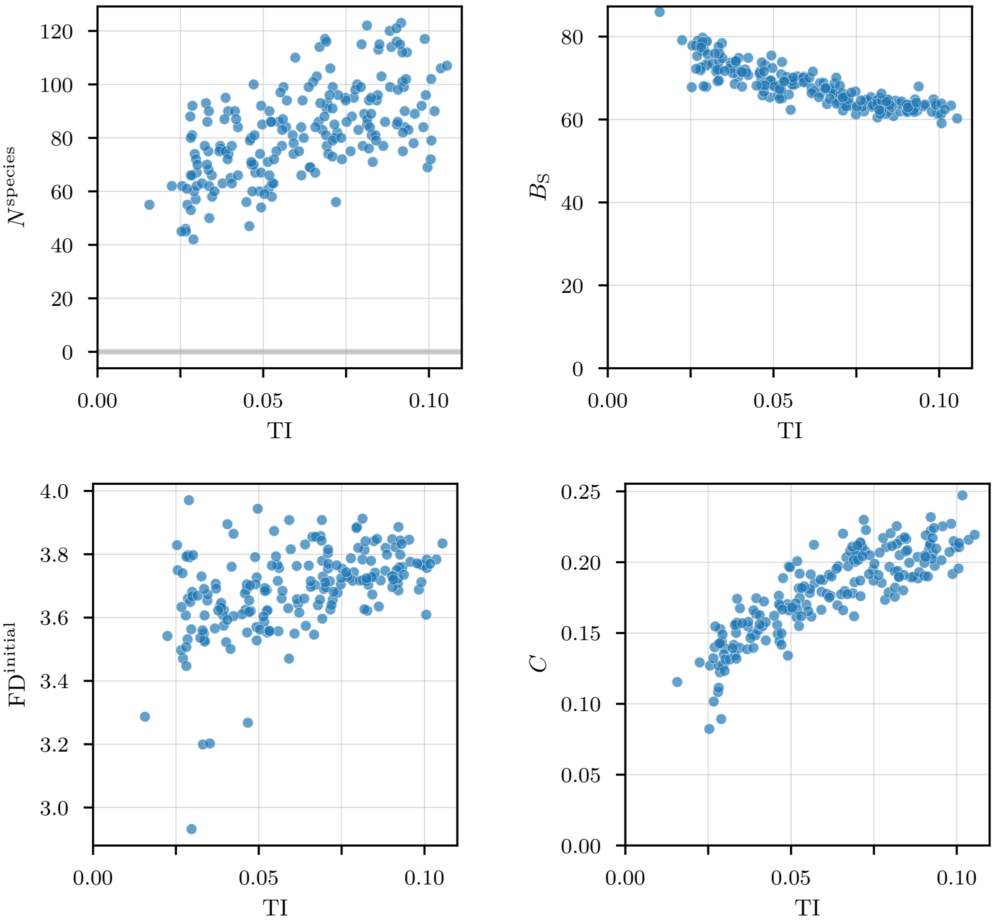

Figure B.11 Correlation between properties of the initial states of systematic invasion experiments in dependence of trophic incoherence. This is supplementary information to the results presented in Section II.4.1.4.

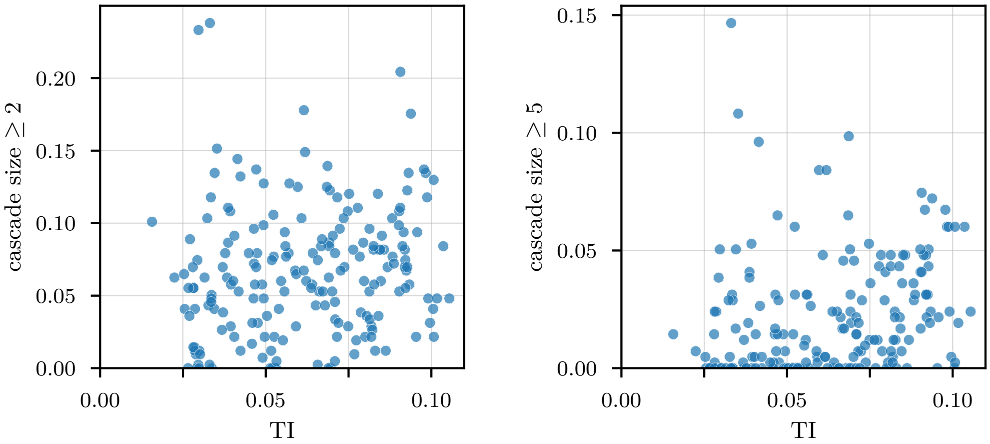

Figure B.12 Ratio of states with large or very large extinction cascades after species invasion. This is supplementary information to Figure II.4.13.
Note the differently scaled y-axes.

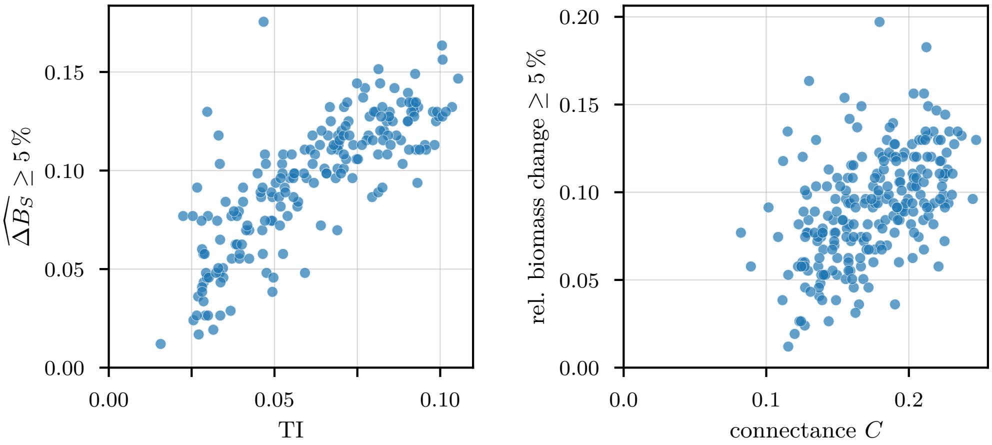

Figure B.13 Comparison between the ratio of states with strong positive biomass change for the invasion experiments depending on the trophic incoherence **(left)** and the connectance **(right)**.
Like for the invasion experiments for the functional diversity and trophic incoherence of the initial state, the initial states for the connectance-related effects were chosen by binned sampling (bin width $C = 0.01$, ≤ 20 samples per bin), yielding 237 food webs.
Note the differently scaled (and labelled) y-axes.

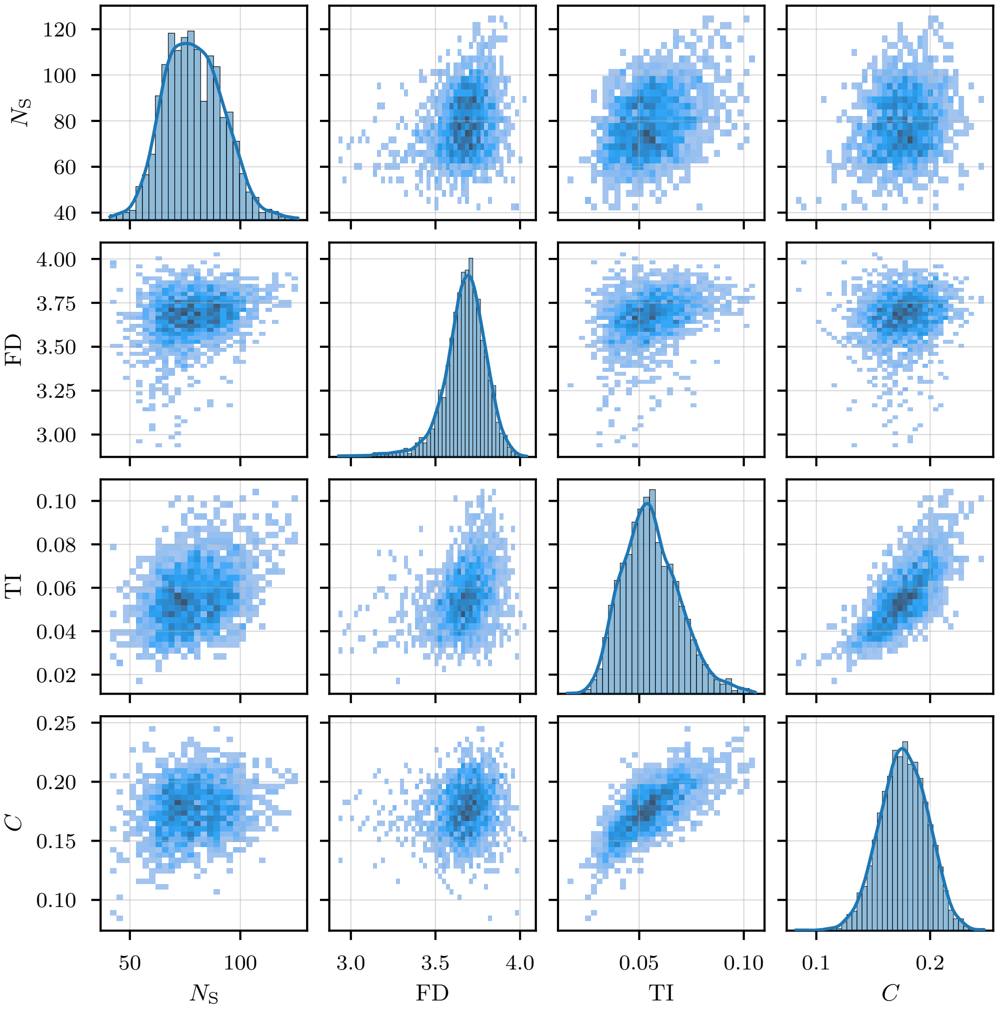

Figure B.14 Characteristics of the pool of 2048 evolved food webs, which are used to study perturbation response (see Section II.4.2) and part of which is used for systematic analysis of invasion response (see Section II.4.1.4). Note the skew in the functional diversity distribution, which is caused by the upper limit for FD given by the available energy. Regarding the correlation of the network measures, connectance C and trophic incoherence TI, see the remarks made in Section II.4.1.4.

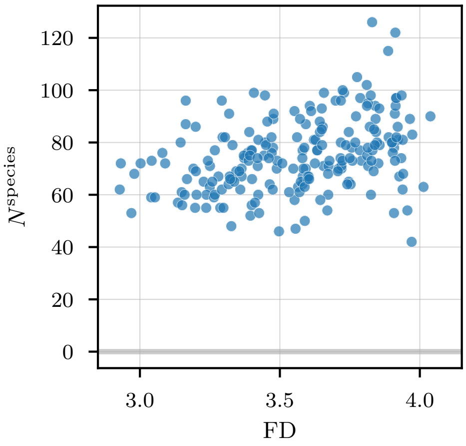

Figure B.15 Initial number of species for the same simulation as shown in Figure II.4.8.

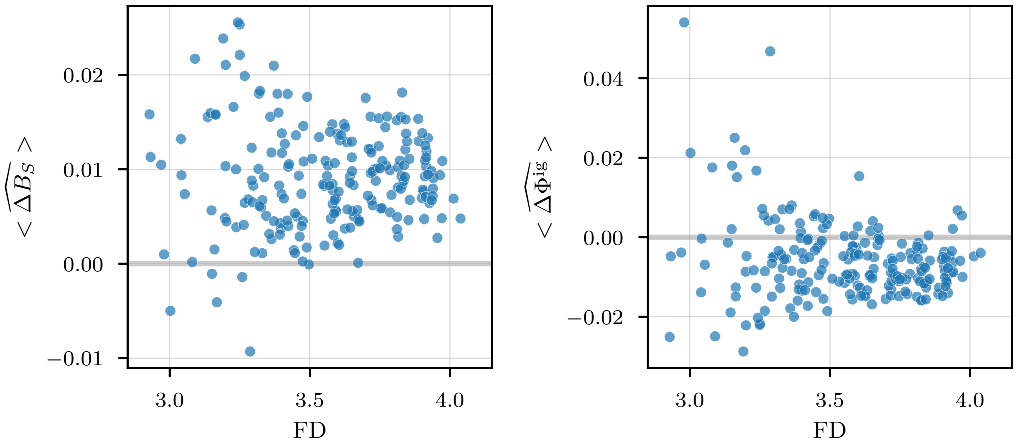

Figure B.16 Mean relative changes in species biomass B_s **(left)**, and intra-guild biomass flow Φ^ig **(right)**. Owing to the regions with strongly positive and negative biomass changes, see Figure II.4.10, the standard deviation over trait space is very large (≈ 0.03) and hence omitted here.

The overall positive value of $\widehat{B_\text{s}}$ and the negative value for $\widehat{\Phi^\text{ig}}$ are driven by the changes in the high $^m/_f$ region, which are stronger in absolute values than the biomass losses in the low $^m/_f$ regions.

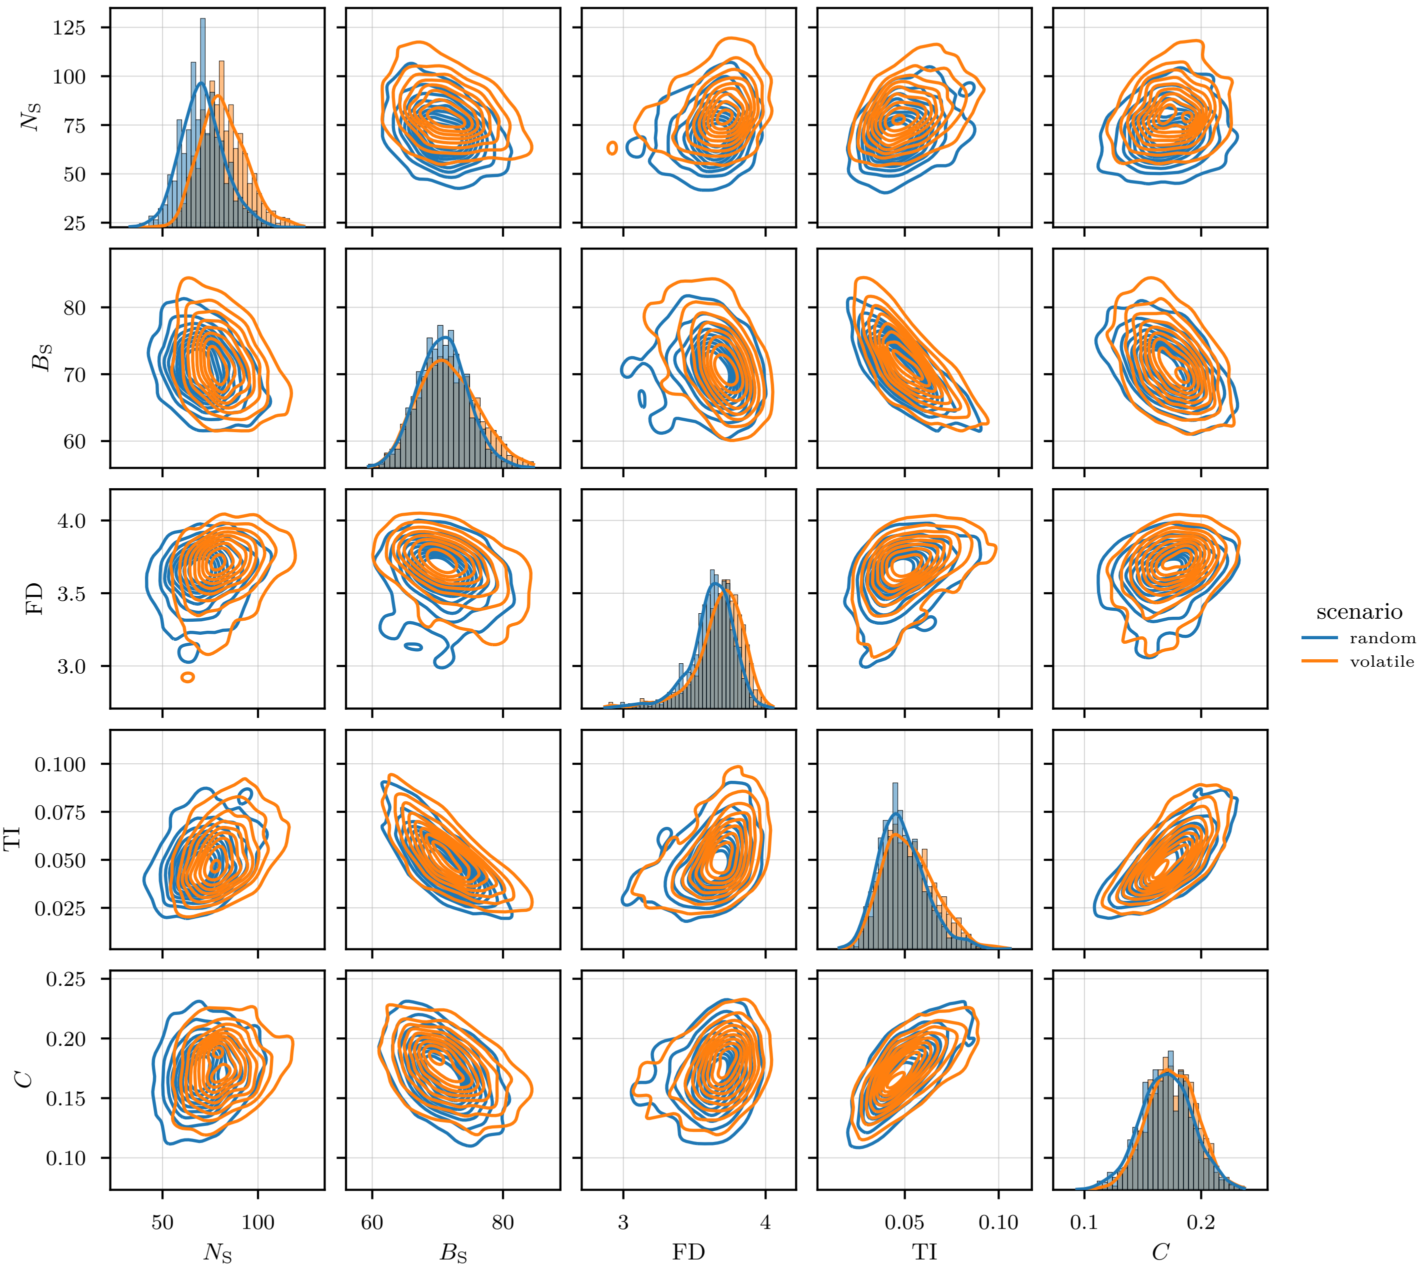

Figure B.17 Characteristics of two pools of 1249 evolved food webs each, which are used to study perturbation response, see Section II.4.2.

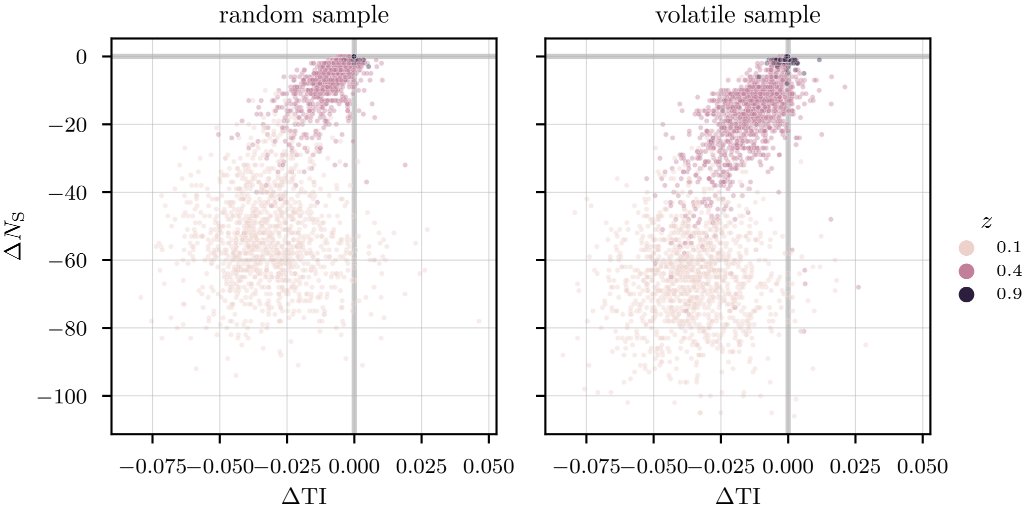

Figure B.18 Perturbation response in species number and change in trophic incoherence. For intermediate perturbation strengths, ΔTI seems to correlate to ΔN_S more strongly in the volatile than in the random sample.
See Figure II.4.16 for more information.

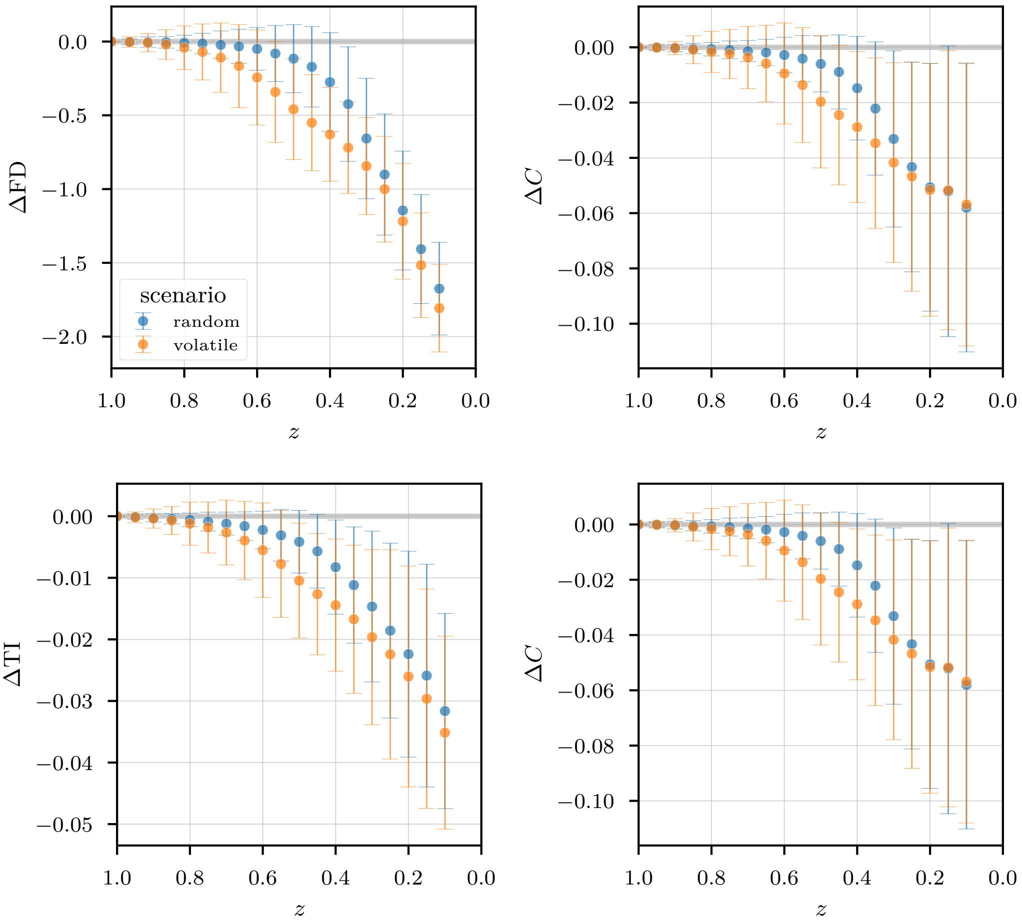

Figure B.19 Absolute change in FD **(left)** and connectance C **(right)**, depending on the perturbation factor z.
These observations are from the same simulation as shown in Figure II.4.17.

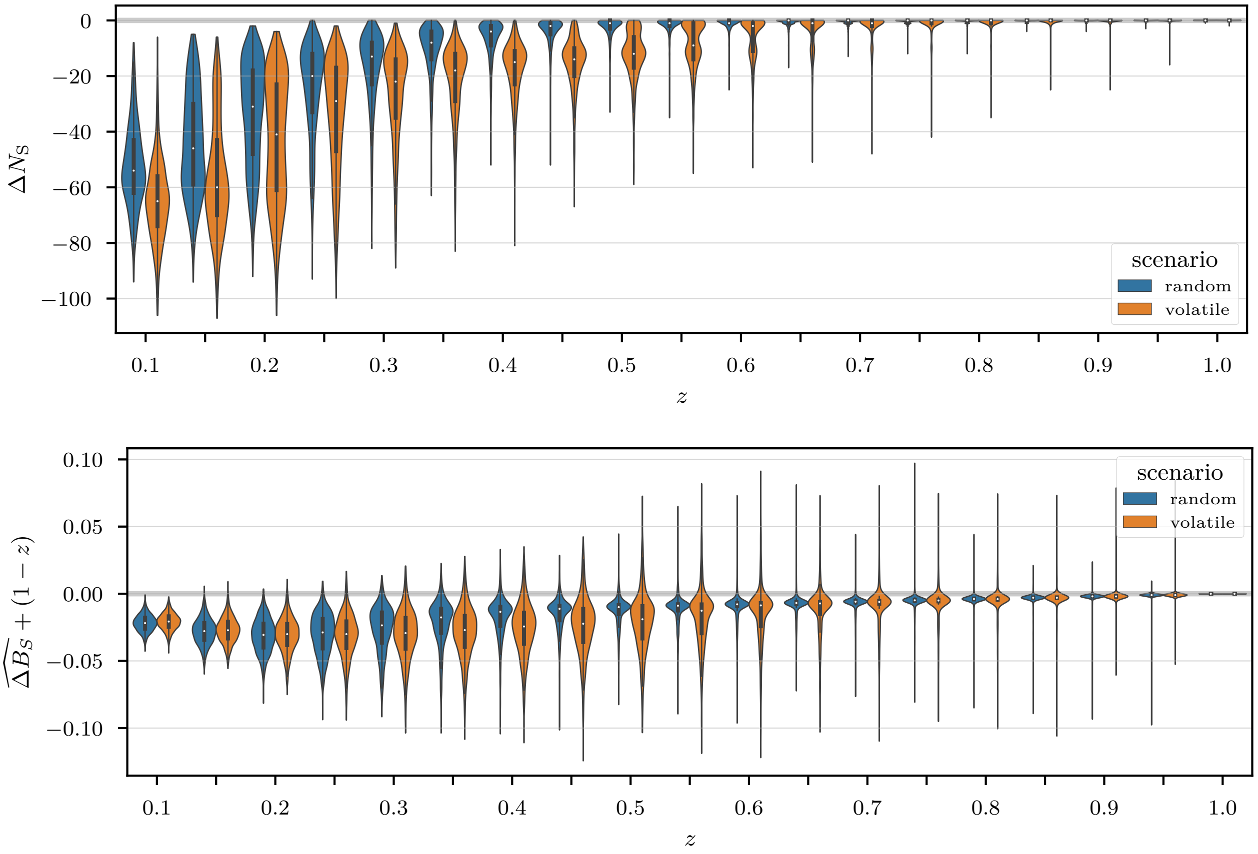

Figure B.20 Same analysis as done in Figure II.4.17 and Figure II.4.18 but showing the respective distributions using a violin plot representation (KDE & inset box plot). KDEs are not extended beyond the data range and are not normalised, hence allowing only comparison of the overall shape and extremal values.
Note that, unlike in the referenced figures, the x-axis is *not* reversed here.

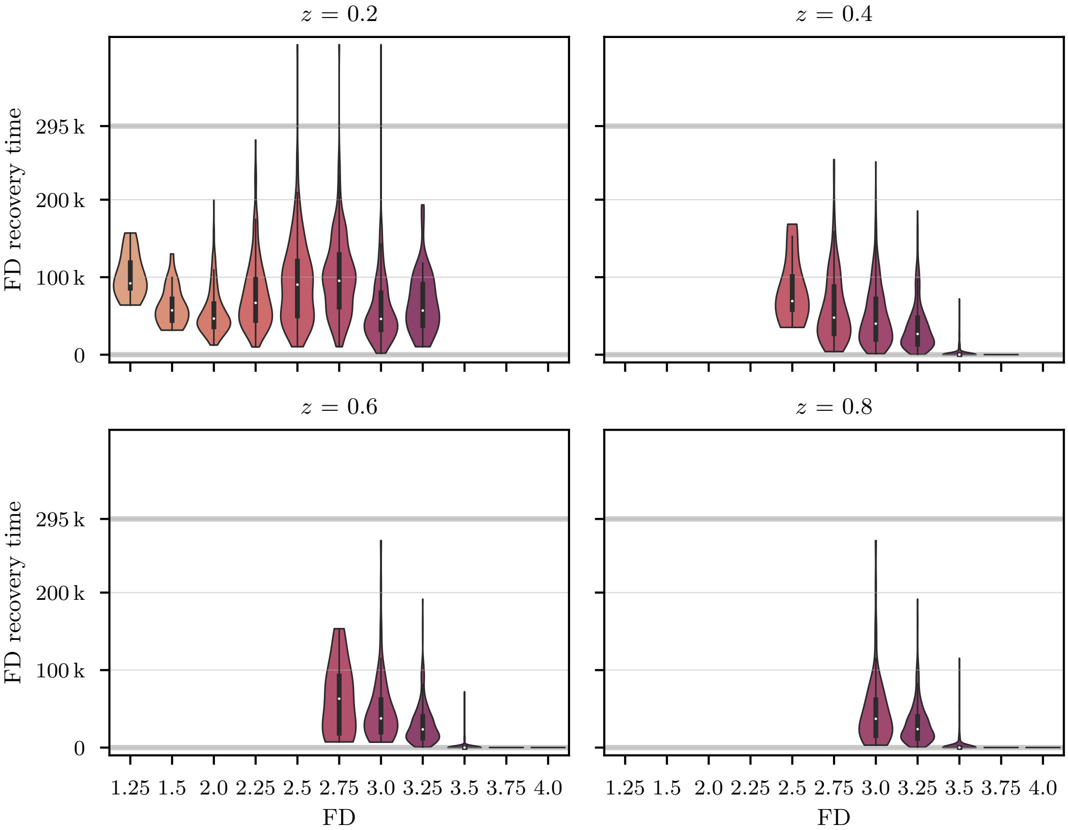

Figure B.21 Recovery time of food webs after perturbations: simulation times until the rolling-time-averaged functional diversity reached FD $\geq$ 3.4. The grouping on the x-axis is by the FD of the food web *after* the perturbation, with values denoting the bin centre.
See Figure II.4.19 for more informtaion.

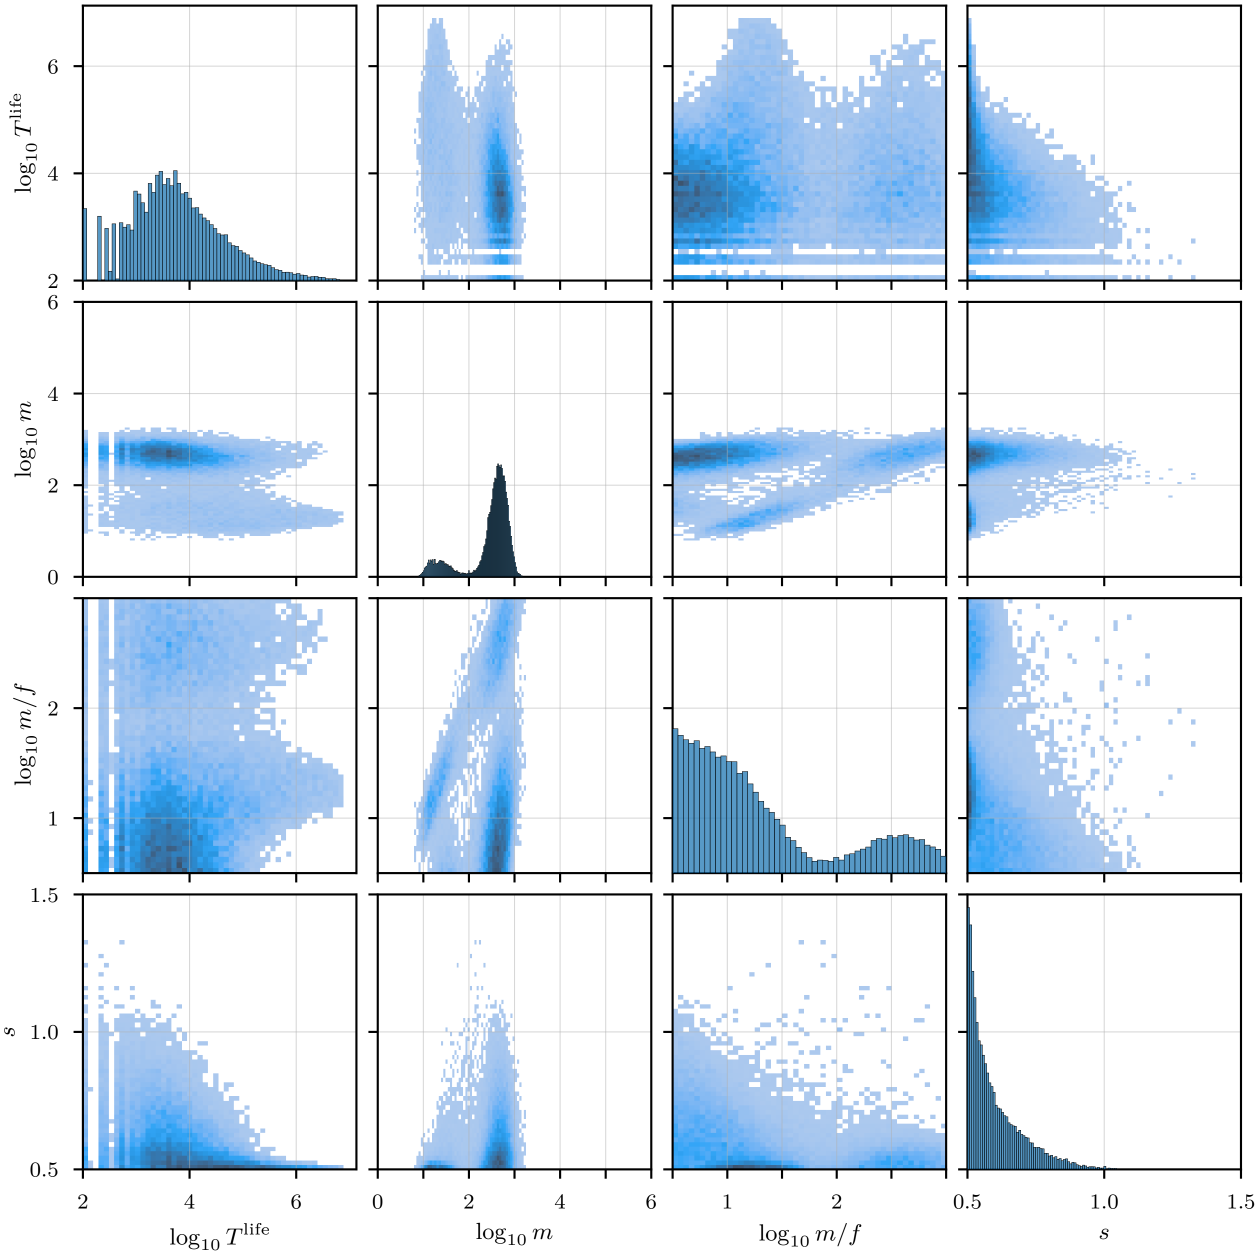

Figure B.22 Species life time and traits distributions resulting from 24 simulations with 8 M steps (as in Figure II.5.3) for a reduced carrying capacity of $K = 5$. Only viable species are shown.

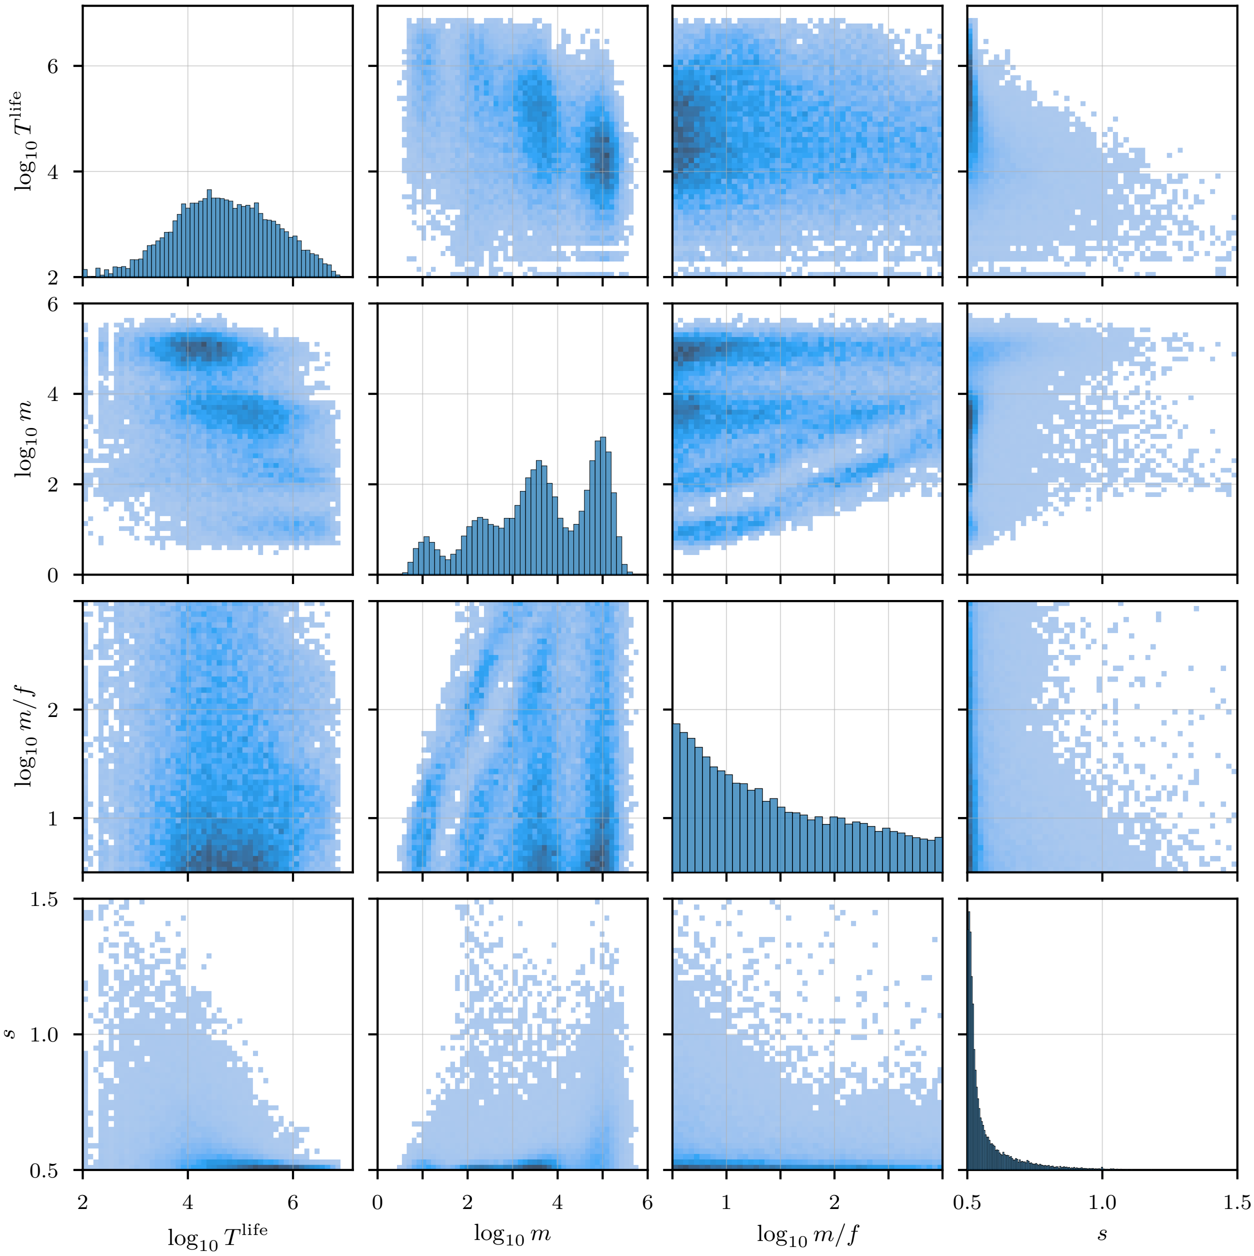

Figure B.23 Species life time and traits distributions as in Figure B.22 but for $K = 1000$.

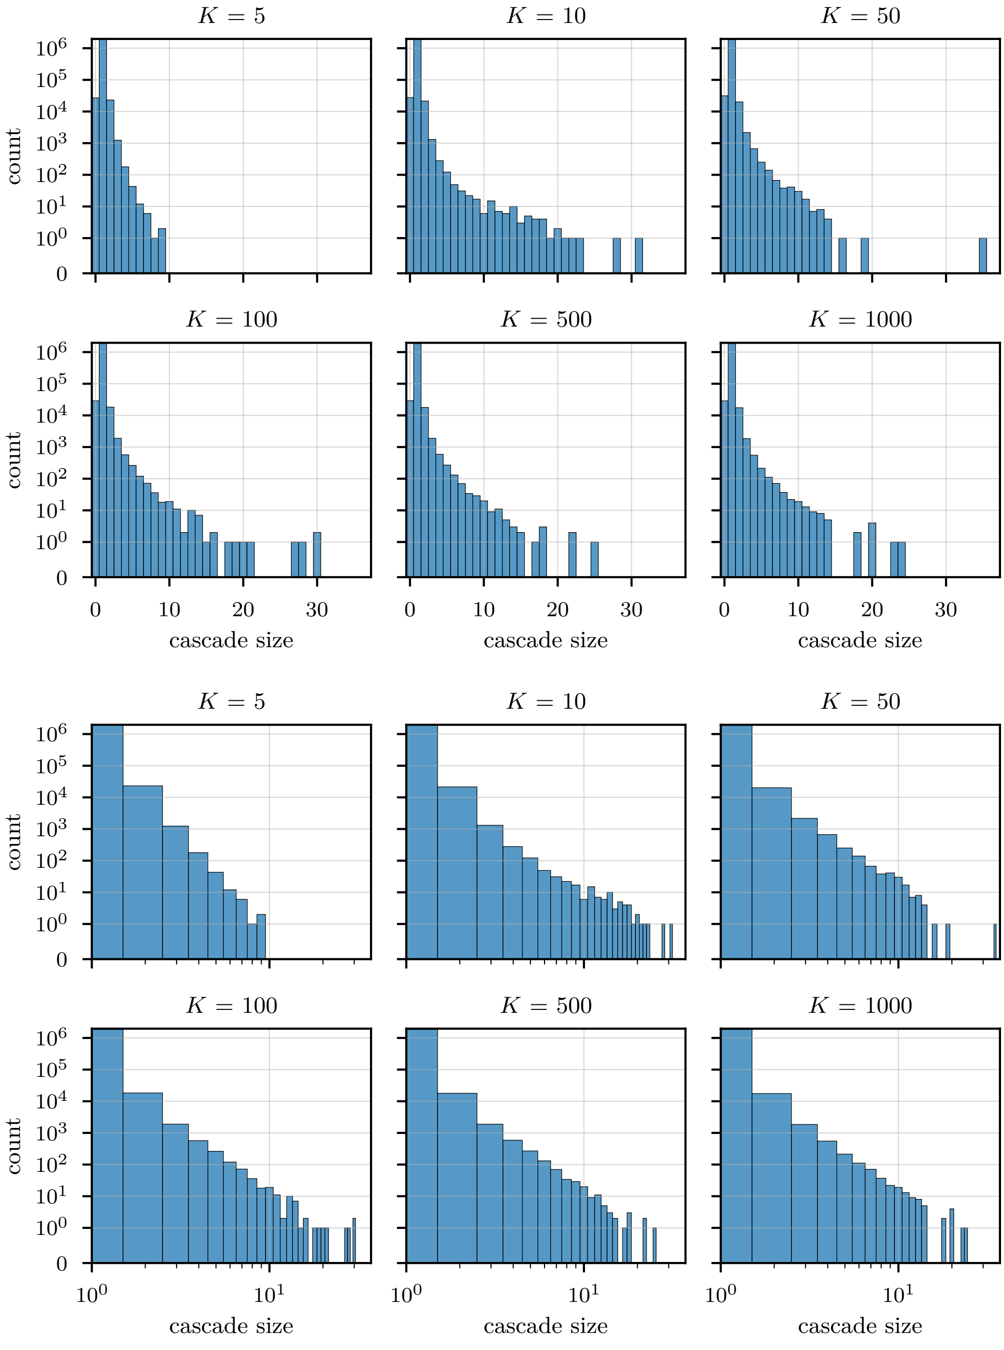

Figure B.24 Cascade size distribution depending on K, in linear **(top)** and logarithmic **(bottom)** representation.

Acknowledgements

The years of my doctoral studies have been an inspiring and exhilarating time. I am deeply grateful to everyone who joined me on that path and supported me throughout this journey, in one way or another – thank you!

To Kurt Roth for the opportunity and freedom to pursue this project, for intense and inspiring discussions, for widening my horizon in many directions – and for a great deal more that I cannot put into words.

To Jonathan Donges for co-supervising me, reflecting on ideas together, and initiating interesting discussions with the COPAN group.

To Björn Malte Schäfer for taking the time to be a referree of this thesis and for the uncomplicated and supportive communication in the final weeks.

To Korinna Allhoff for kindly answering my questions about the ECST model and providing her model implementation to me for reference.

To Benjamin Herdeanu, Harald Mack, and Lukas Riedel for sharing this journey with me, for working on and brooding over so many questions together, and for putting up with me even in more challenging situations.

To Thomas Gaskin, Julian Weninger, Jeremias Traub, Lukas Siebert, Narek Baghumian, Laila Schmidt, Paul Wiesemeyer, Daniel Lake, Josephine Westermann, Fabian Krautgasser, Hendrik Leusmann, and many others from the TS-CCEES group for creating an inherently joyous and thought-provoking atmosphere – be it during the group retreats, in spontaneous discussions, in prolonged coffee breaks, or in the many online meetings we had.

To all Utopia and dantro collaborators for jointly developing the research software used for this thesis and learning many lessons about collaborative software development in the process.

To Angelika Gassama and the IUP's secretariat for their friendly and competent help in all administrative matters.

To the state of Baden-Württemberg for support through bwHPC and the German Research Foundation (DFG) through grant INST 35/1134-1 FUGG.

To the open source software community for doing what they do.

To Elke Kuno for generously housing me during my visits to Heidelberg; these last years would have been far more challenging without her support.

To my family and friends for always being there for me, even from afar.

To the two furry companions who, coincidentally, were always in the same room as me when I was working on this thesis.

And to Maria. For everything.